Fair Scheduling and Buffer Management in Internet Routers

Nan Ni
IBM Corporation
Austin, TX 78758

Laxmi N. Bhuyan
Computer Science and Engineering
University of California
Riverside, CA 92521

Abstract—**Input buffered switch architecture has become attractive for implementing high performance routers and expanding use of the Internet sees an increasing need for quality of service. It is challenging to provide a scheduling technique that is both highly efficient and fair in resource allocation. In this paper, we first introduce an iterative fair scheduling(*i*FS) scheme for input buffered switches that supports fair bandwidth distribution among the flows and achieves asymptotically 100% throughput. The *i*FS is evaluated both under synthetic workload and with Web traces from the Internet. Compared to the commonly used synthetic input, our simulation results reveal significant difference in performance when the real network traffic is employed. We then consider fair scheduling under various buffer management mechanisms and analyze their impact on the fairness in bandwidth allocation. Our studies indicate that early packet discard in anticipation of congestion is necessary and per-flow based buffering is effective for protecting benign users from being adversely affected by misbehaved traffic. Buffer allocation according to bandwidth reservation is especially helpful when the input traffic is highly bursty.**

Keywords—**quality of service, fair bandwidth allocation, switch scheduling, buffer management, decongestion mechanism, web traffic.**

I. INTRODUCTION

The exponential growth of the Internet has put increasing demands on the routers and switches in the network for high bandwidth and low latency. In addition, as networks provide new services supporting multicast, voice, security, and bandwidth reservation, quality of service (QOS) is becoming a major issue in design of routers [25], [19]. Fairness in resource allocation is very important to support the need for diverse applications. Fair queuing algorithms have been developed [8], [38] to schedule packets at an output link of a router. However, little research has been done to address QoS issues inside the router operation itself. The purpose of this paper is to develop fair scheduling and buffer management schemes for Internet routers and to demonstrate their superiority using actual Web traces from NLANR [27] and UCB [37].

A router consists of three parts, namely, (a) line cards that connect to datalinks (b) a router processor that runs routing protocols and (c) a backplane crossbar switch that actually transfers the packets or cells from inputs to outputs. In this paper, we are mainly concerned with QoS issues of the backplane switch design in a router. The switch has buffers either at the input or output to store the packets temporarily during transmission. Although output buffering can achieve better throughput, it is known to suffer from poor scalability [18]. The reason is that the output port of an $N \times N$ switch has to operate N times faster than the input in order to accommodate requests on all possible inputs in a cycle. Consequently, most high-performance routers employ input queues with their crossbar backplanes [24], [30].

In this paper, we develop fair scheduling schemes for input-buffered switches. High throughput and fairness in resource allocation are contradictory goals in a switch design. Due to the fact that only a single cell can be transmitted from each input of the crossbar in a given slot, cells forwarded based on maximal set of input-output match may not coincide with those satisfying fairness. On the other hand, passing cells based on fair resource allocation may not produce the highest throughput and may give rise to under utilization of the crossbar switch. The aim of this research is to find scheduling and buffer management techniques that are both fair and highly efficient for link utilization.

This research has been supported by NSF Grants CCR 9622740 and CCR 9810205. This work was done while the authors were at Texas A&M University.

According to [8], bandwidth allocation of a link is fair if for each flow, the received bandwidth is proportional to its share of reservation. We know that the fluid flow queueing(or Generalized Processor Sharing(GPS)[29]) algorithm which sends packets in a bit-by-bit round robin fashion is absolutely fair but unrealistic. Numerous fair queueing algorithms have been developed to approximate the GPS algorithm in units of packets(see [38] for a survey). However, most work on fair queueing has been conducted in the context of output queueing due to its conceptual simplicity. Situation becomes complicated in the case of input buffered switch where flows not only contend for output link bandwidth but also have to compete to access the crossbar. Nevertheless, in order to achieve high performance as well as to provide quality of service, it is imperative to maintain fairness in input buffered switch scheduling.

In the first part of this paper, we propose an iterative fair scheduling(*i*FS) scheme which provides high throughput, low latency as well as fair bandwidth distribution among the contesting flows. As with other iterative approaches, *i*FS tries to increase the number of input-output matches during each iteration, but the major difference is that instead of picking a cell from each input in a probabilistic([1], [14]) or round robin([22], [34]) fashion, a cell is chosen based on the bandwidth requirement of the flows. In this way, we are able to allocate bandwidth to various flows in proportion to their reservation and prevent misbehaved flows from taking advantage at the expense of others. The *i*FS is compared with *i*SLIP [22] and WPIM [36] for throughput and fairness.

In a real deployment of network with finite buffering space, packets are dropped in the presence of congestion. In order to support fair bandwidth allocation, it is crucial to determine when to drop and which packet to drop. Thus, the fairness in scheduling issue should be considered in conjunction with the decongestion mechanism. In this paper, four decongestion mechanisms are studied and their impact on the fair sharing of bandwidth is examined.

The effectiveness of our fair scheduling approaches is demonstrated by simulations using a cycle-based simulator. Experiments are conducted with different types of input workload. Investigation has shown that the Internet traffic is very bursty in nature [31], [7] and therefore can not be characterized very well by commonly employed Bernoulli or geometrically distributed ON/OFF traffic model. While distributional models have been developed for the Internet workload [4], traffic traces are widely applied for evaluating web servers [2], proxy caches [6], [21] and packet forwarding methods [17]. The use of real traffic traces can offer direct validation of the network components under study. Hence, we incorporate the traffic pattern directly from the Internet in addition to employing synthetic workload and we are not aware of any existing application of traces in the context of router design or switch scheduling.

To summarize, we have the following original contributions in the paper:
- We propose an iterative fair scheduling(*i*FS) scheme for unicast traffic. It supports fair bandwidth allocation and achieves asymptotically 100% throughput with uniform traffic.
- We analyze various buffer allocation policies along with fair scheduling and pinpoint their effect on fair bandwidth allocation.
- Ours is the first attempt to evaluate switch scheduling schemes using

real traffic from the Internet. It gives some insight into the difference in performance compared to the common practice of evaluating through synthetic workload.

The rest of the paper is organized as follows. Section II gives an overview of the related work on scheduling for input buffered switches. In Section III we propose the iterative fair scheduling(*i*FS) for unicast traffic. Extensive simulation results with synthetic workload as well as traces from the Internet are presented in Section IV. Buffer management schemes are examined in Section V. Finally, Section VI concludes the paper.

II. RELATED WORK

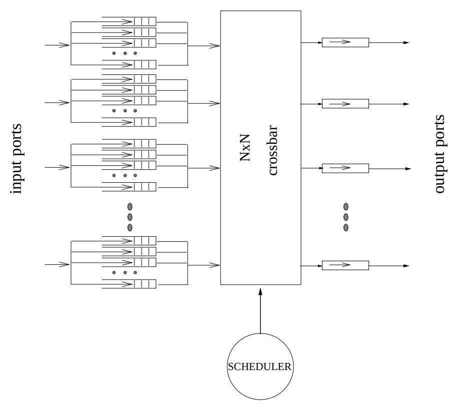

Fig. 1. Block diagram for an input buffered switch architecture

Input queueing is subject to head of line(HOL) problem with maximum throughput of 58.6% using FIFO input queue[18]. Solution has been found which constructs separate queues at each input so that a packet will not be blocked by packets going to other destinations. Major issues in the input buffered switch architecture design include buffer management and scheduling. We have studied various queueing and buffer allocation schemes in our previous research [9], [33]. In this paper, we seek to address the issue of fairness in scheduling(i.e., input to output matching) for quality of service. Figure 1 gives a block diagram of the cell-based switch architecture we consider throughout the paper. It consists of four components: input ports, output ports, crossbar switching fabric and a scheduler. At any given switch cycle(slot), at most one cell from each input can be routed to the output side and each output can only accept at most one cell. As a result, the scheduling is more intricate than that for output queued switch. Most existing work on scheduling for input buffered switch attempts to achieve high throughput by looking for maximum bipartite matching between inputs and outputs. Schemes such as PIM[1], *i*SLIP[22] and Shakeup[14] repeatedly search for matches at each time of scheduling. Some approaches can achieve 100% throughput asymptotically[23].

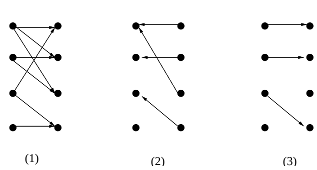

Fig. 2. 3-step procedure in iterative matching

An iterative scheduling algorithm essentially consists of three major steps in each iteration, as illustrated in Figure 2.

1. *Request stage*: An input may have several requests. Each unmatched input sends requests to the outputs for which it has cells for.
2. *Grant stage*: There may be several requests to an output. Each unmatched output chooses one from several received requests and sends a grant signal to one of the inputs.
3. *Accept stage*: Each unmatched input may receive grant signals from several outputs. Upon receiving grant signals, each input sends accept signal to only one of the outputs offering the grants.

Probabilistic Iterative Matching(PIM) suggested by Anderson *et al.*[1] is the first switch scheduling algorithm that employs an iterative approach. At grant stage, each output sends out a grant signal to a randomly selected requesting input. Again at accept stage, each input accepts one output from several grant signals at random. It takes an average of $O(logN)$ iterations for the PIM to converge. The consequence of random selection in grant and accept stage of PIM is that an input could possibly end up in not being served for a long time. Thus, this scheme is not starvation free. Moreover, implementation of random selection among the member of a time-varying set is not an easy task.

The *i*SLIP scheme proposed by McKeown [22] uses rotating priority(round-robin) arbitration to schedule active inputs and outputs in turn. A grant pointer is kept for each output to track the input with the highest priority. Similarly, there is an accept pointer for each input which tells the output with the highest priority. Whenever a match is found, the corresponding grant/accept pointers are incremented(modulo the number of ports). Compared to PIM, *i*SLIP is simpler in implementation[16] and achieves higher throughput[22].

Matches are irrevocable in both PIM and *i*SLIP, which means that later iterations can only add upon previously made matches but can not change them even if better matches can be found. To escape this "local maximum", Goudreau *et al.* offer a Shakeup technique in [14] where each unmatched input is allowed to force a match for itself randomly even though an existing match has to be knocked off. In other words, Shakeup attempts to find the global maximum, but the feasibility of its implementation in real system is unknown because it takes more iterations to converge.

The issue of fair resource allocation is not considered in the above mentioned approaches. Iterative schemes that do address the fairness issue include statistical matching [1] and Weighted Probabilistic Iterative Matching(WPIM)[36]. Statistical matching is similar to PIM, except that the matching process is now initiated by the outputs, each generating a grant signal to a randomly selected input based on the reserved proportion of bandwidth. It can happen that an input is picked when its queue is empty and the switch is poorly utilized. WPIM is also built upon PIM where based on the reservation, every input flow is assigned a quota that can be used in a frame of a constant number of slots. During each frame, flows that have not reached their quotas secure equal share of bandwidth by random selection as done in PIM. To accomplish this, an additional masking stage is added to the 3-step procedure to exclude those inputs that have consumed their quotas in the current frame.

Scheduling fairness is also addressed in other work under assumptions not relevant to this paper. Stephens and Zhang [35] considered a switch that has effectively fully connected crossbar [33] and processes packets of variable length. There, the inputs and outputs are decoupled so that each output independently reads from any input with a packet for it. Li and Ansari gave end-to-end delay bound provided that incoming flows conform to the (r, T) traffic model[20].

Among the packet discarding schemes, some assume a queueing discipline of FCFS, others are designed with per-flow queueing. The most basic one among the FCFS disciplines is the "do nothing" policy, which is a complete buffer sharing with a *drop-tail* (DT) mechanism, where cells are dropped when they arrive to find the buffer full. Generally, network systems do not support cell level retransmission, so a partially received packet is of no value. In *partial packet discard* (PPD) [3], after a cell from a packet is dropped, all subsequent cells of the same packet are dropped as well. In *early packet discard* (EPD) [11], the entire packet (i.e., all the cells constituting a new packet) is dropped whenever congestion is anticipated because the buffer occupancy exceeds certain threshold. A more sophisticated policy by Floyd and Jacobson is *random early discard* (RED) [10], whose primary goal is to avoid performance degradation and unfairness caused by DT. It does so by maintaining average buffer occupancy at a level significantly below the total number of buffers. To achieve this, packets are dropped with a certain probability when the average buffer occupancy reaches certain level. Drop probability increases with the average queue occupancy, and once the queue occupancy exceeds a maximum buffer threshold, all arriving packets are discarded. All these schemes assume a FCFS scheduler, each arriving packet is treated identically and all the flows see the same loss rate.

Example of per-flow discarding include *longest queue drop* (LQD) [12] and *fair buffer allocation* (FBA) [26]. The justification behind LQD is that if flows are given equal weights, the ones that use the link more tend to have longer queues. Hence, biasing packet discarding such that flows with longer queues have higher drop rate should make the bandwidth sharing more fair. In FBA, buffer share for each flow is computed as a function of the number of free buffers and the number of active connections. Once the aggregate occupancy is above the specified threshold, packets arriving for flows that have occupied more than the fair shares are thrown away. The CHOKe algorithm in [28] provides an approximate fair dropping mechanism to control unresponsive UDP flows without using per-flow information.

III. Iterative Fair Scheduling Scheme

In this section, we first introduce a definition of fairness in input buffered switch scheduling and then propose an iterative fair scheduling(*i*FS) scheme that can be used to achieve fair bandwidth allocation.

Let $f(i,j)$ denote the jth flow from input i. It goes to output $d_{i,j}$ and reserves a bandwidth of $B_{i,j}$. The number of flows from input i is n_i. Let $t_{i,j}(t_1,t_2]$ be the amount of traffic(in bits) reaching the output from flow $f(i,j)$ in time interval $(t_1,t_2]$. We say that two flows $f(i_1,j_1)$ and $f(i_2,j_2)$ are in contention if $i_1 = i_2$ or $d_{i_1,j_1} = d_{i_2,j_2}$. Note that we do not define precisely what a *flow* is because we would like our *i*FS scheme to be applied in a broad sense: Depending on the desired granularity of QoS and scalability concern, a flow can be as fine as a TCP connection or as coarse as a class of aggregated connections with similar bandwidth requirement.

On the analogy of the definition of the GPS server for a single shared resource in [29], we consider a scheduling scheme for input buffered switches to be fair as follows.

Definition 1: For any two back logged flows $f(i_1,j_1)$ and $f(i_2,j_2)$ that are in contention, a scheduling scheme is fair in $(t_1,t_2]$ if

$$\frac{t_{i_1,j_1}(t_1,t_2]}{t_{i_2,j_2}(t_1,t_2]} = \frac{B_{i_1,j_1}}{B_{i_2,j_2}}$$

This definition specifies the ideal situation, but if the output link bandwidth is to be best utilized (i.e., work conserving), it is possible that the equation may not be held under all the combinations of bandwidth reservations. The discrepancy of this definition from the one of GPS[29] is that here the flows are contesting to access the crossbar as well as to access the output links. In other words, we are trying to allocate correlated resources because a cell must first make its way out of the input buffer before going to the intended output.

Nevertheless, we find that the existing fair queueing algorithms, which have been successful in allocating a single shared resource, can be applied here to facilitate our effort to support fair bandwidth distribution. In a fair queueing algorithm, a virtual system time $V(t)$(corresponding to the number of rounds made till time t in GPS) is maintained. Every incoming packet is assigned a virtual starting time and a virtual finishing time depending on its bandwidth requirement. Virtual starting time and virtual finishing time denote the virtual time when a packet should begin and finish sending if GPS were used. The transmitting order is then regulated according to non-decreasing starting time[15], non-decreasing finishing time[8] or a combination of both[5]. Naturally, flows with larger bandwidth reservations will take greater proportion of bandwidth since they tend to have smaller virtual starting time(or finishing time).

Our basic idea of *i*FS is to enhance the iterative approaches described in Section II by giving out grant signals from the individual outputs based on virtual time and then try to resolve input contention in the accept stage.

For each output link we maintain a fair queueing engine, which assigns a virtual time to every incoming cell based on bandwidth reservation of the flow. Note that the arriving cells are queued in the input buffer first-in-first-out on a per flow basis. This is required for implementing a fair queueing algorithm where each output must keep track of the active flows to compute the virtual time. It is different from other iterative approaches where input queues are arranged on per output basis.

In each iteration, the first step in our scheme is for every unmatched output to independently send a grant signal to one of the unmatched inputs which has the cell with the minimal virtual time corresponding to that output. Such a cell is marked as a “candidate” with respect to its input. It may so happen that an input receives multiple grants from different outputs and have multiple candidates, which means that several flows from this input have minimal value of virtual time with their relevant outputs. Then, each unmatched input selects among its candidates a cell(called “winner”) with the oldest age at the switch and sends accept signal to its desired output. That is, in accept stage, an input resolves the contention on a first-come-first-serve(FCFS) basis. The justification is that since these candidate cells are from the same upstream node, the fact that the cell arrives first among the contending cells implies that it has the smallest virtual time among them at the previous node and naturally should be the first one to depart from the current node.

In summary, the *i*FS scheme can be formalized as the following:

- Initially, all inputs and outputs are considered as unmatched and none of the inputs have any candidates.
- Then in each iteration:

1. *Grant stage*: Each unmatched output selects a flow with the smallest virtual time for its head-of-line cell and marks the cell as a candidate for the corresponding input. Grant signal is then given to the input.
2. *Accept stage*: Each unmatched input examines its candidate set, selects a winner according to age and sends an accept signal to its output. The input and output are then considered as matched. Reset the candidate set to empty.

- At the end of each switch cycle, the winning cells are transferred from the input side to the output side.

Comparing *i*FS to other iterative scheduling schemes, we see that grants in *i*SLIP are given in round-robin manner without any respect to bandwidth reservation. WPIM guarantees fair bandwidth sharing, but in a much coarser granularity. Before running out of quota in WPIM, every flow has equal access to the bandwidth. Consequently, the flows with larger reservations get more bandwidth only after other flows run out of their quotas, which usually happens towards the end of a frame. If we look at the bandwidth distribution within a frame, the bandwidth share is disproportional. In statistical matching, grants are given by outputs to inputs randomly, therefore it is possible that an input port is selected to receive grant signal when its queue is empty. It has been shown in [36] that statistical matching has inferior performance than WPIM. Unlike Shakeup where matches in prior iterations can be modified, matched input-output pairs are ruled out in later iterations in *i*FS. Although Shakeup is attractive theoretically, its feasibility for high performance switch is unclear because one "knock-off" of an existing match could possibly trigger a chain of "knock-offs" which would take more iterations to converge to global maximum matches.

The issue we do not consider in detail here is the implementation feasibility. Efficient implementation of N fair queueing engines is crucial to the *i*FS for an $N \times N$ switch. It is to our advantage that the switch under discussion is cell based. The calculation of virtual time can be simplified due to fixed packet length. The switch proposed in [35] needs $3N$ fair queueing engines to deal with variable length packets for an $N \times N$ switch. A second advantage is that, unlike PIM, WPIM or Shakeup, no random number generator is needed for *i*FS, which again greatly reduces the implementation complexity. Hardware design for *i*SLIP is described in [16] where three iterations are made possible within a cycle of $51ns$. We do not have the expertise to pinpoint the timing for *i*FS, but we conjecture that it is feasible for high performance switches supporting OC-12, OC-48 line rate or even higher.

IV. Performance Evaluation of *i*FS

Following the criteria in [32] for assessing a resource allocation scheme, we evaluate the proposed *i*FS in two aspects: efficiency and fairness. The principal metrics for efficiency are throughput and delay. A cell based simulator is developed and the simulations are conducted with the assumption of infinite buffer size. The offered load refers to the probability that a cell arrives at an input in a given slot and the underlying fair queueing algorithm for *i*FS is self-clocked[13]. Our methodology of evaluation is as follows. We start by examining *i*FS and other approaches using synthetic workload under cases where destinations are uniformly and non-uniformly distributed among the outputs. Results are presented to show that *i*FS can achieve average cell latency and overall throughput close to the existing schemes. Then we compare the ability of various approaches to support fair bandwidth distribution. In the second part of this section, traces from Internet traffic are applied to the simulator to study how *i*FS works in the "real world". Our results demonstrate that under both synthetic and real network traffic, *i*FS can achieve high throughput like *i*SLIP and at the same time supports fair bandwidth allocation.

A. Measurement from synthetic workload

We begin by evaluating the throughput and the average delay of *i*FS under benign i.i.d. Bernoulli traffic, where at any given slot, a cell arrives with the probability determined by the offered workload. The simulation is performed on a 16×16 switch with 16 flows per input each destined for a different output, for a total of 256 flows. In the first scenario, each input port sends cells with destinations uniformly distributed among all the output ports. Figure 3 shows the delay versus the offered load for *i*FS, *i*SLIP and WPIM with the number of iterations equal to four. The *i*SLIP is known to offer low delay for an input buffered switch. The delay using output buffered switch is also shown for comparison. Under this circumstance, we observe that the average cell delay for *i*FS is almost identical to *i*SLIP and both are very close to output buffered switch, which is the lower bound for the delay. Hence, *i*FS is capable of achieving asymptotically 100% throughput for uniform traffic.

Next we consider a non-uniform Bernoulli traffic model as used in [36]. The assumption is that four of the switch ports are connected to servers and the remaining twelve to clients. Each client sends 10% of its generated traffic to each of the four servers, and the remainder is uniformly distributed among the other clients. Similarly, each server directs 95% of its traffic to the clients and the remaining 5% to the other servers. Figure 4 indicates that *i*FS is very close to *i*SLIP and WPIM in terms of average cell latency and can reach a throughput of 78.5%.

Now we turn to examine the effectiveness of *i*FS to support the fair bandwidth sharing when a link is overloaded. This time, we simulate a 4×4 switch so that we can plot the results with clarity. Assume that every input has four flows each going to a different output. Without loss of generality, let the jth flow from link i go to output j and denote it as $f(i, j)$ following the notation in Section III. Also assume that the flows to output 1, $f(1, 1)$, $f(2, 1)$, $f(3, 1)$ and $f(4, 1)$, have reserved 10%, 20%, 30% and 40% of the bandwidth, respectively. But they always maintain the same actual arrival rate. Others are background flows with a load of 5% each. We vary the input rate of the flows to output 1 and plot the received bandwidth share in Figure 5 and 6 using *i*SLIP and *i*FS scheduling. For workload under 25%, the throughput for every flow is able to keep up with the input workload for both the schemes. However, for workload beyond 25%, all the four flows are still treated equally in *i*SLIP, and therefore they obtain the same share of bandwidth(each 25%) despite of the variance in bandwidth reservation. The *i*FS, on the other hand, is observed to differentiate the flows according to the promised share when the load is greater than 25%. Ill-behaved flows are prevented from influencing the well-behaved ones. For load beyond 40%, each flow receives its allocated bandwidth. Let us look closely at the sharing when input load is between the range of 25% and 40%. At 30%, for instance, flow $f(4, 1)$ does not consume its share of 40% of bandwidth. The unused part is distributed to flow $f(1, 1)$ and flow $f(2, 1)$ so that they acquire a larger fraction of bandwidth than their reservations. Flow $f(1, 1)$ receives 13.3%(versus a reservation of 10%) and flow $f(2, 1)$ receives 26.6%(versus a reserva-

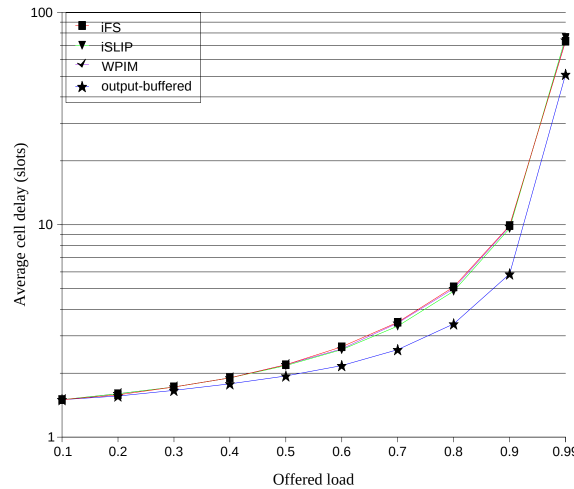

Fig. 3. Performance of *i*FS, *i*SLIP and WPIM under uniform traffic

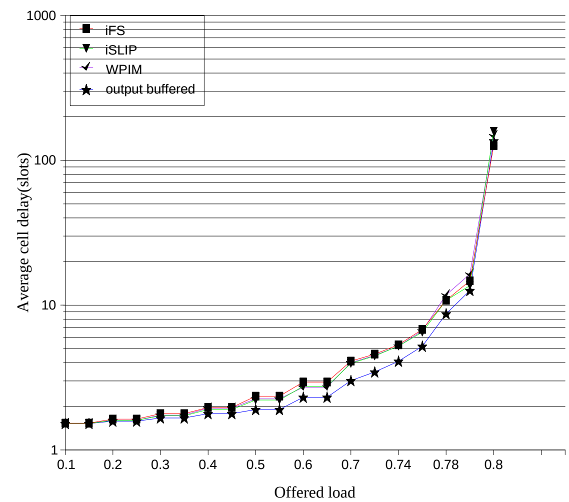

Fig. 4. Performance of *i*FS, *i*SLIP and WPIM under non-uniform traffic

tion of 20%). Such behavior also conforms to the fairness requirement in [8], which states that the unused portion should be assigned equally to other active flows.

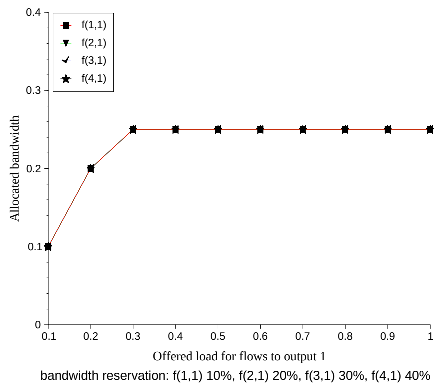

Fig. 5. Received bandwidth using *i*SLIP

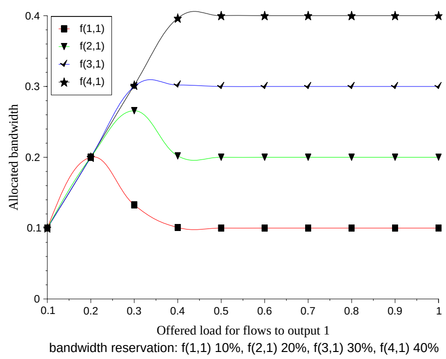

Fig. 6. Received bandwidth using *i*FS

To quantify the fairness, we define *fairness index* in the following to measure the fairness of a switch for allocating bandwidth during time $(t_1, t_2]$. Suppose there are N flows sharing a link and let λ_i, x_i and r_i denote the actual average arrival rate, received bandwidth and reserved bandwidth respectively for flow i during time $(t_1, t_2]$. Without loss of generality, we assume that the first n_1 $(0 \leq n_1 < N)$ flows honor their reservation, that is, $\lambda_i \leq r_i$ for flows 0 through $n_1 - 1$. The rest $n_2 = N - n_1$ flows have arrival rate greater than their reservation. For each over subscribing flow i $(n_1 \leq i < N)$, we denote

$$\beta_i = \frac{\Sigma_{j=0}^{n_1-1}(r_j - \lambda_j)r_i}{\Sigma_{k=n_1}^{N-1} r_k} \tag{1}$$

β_i is the bandwidth that can be spared to the over subscribing flow i, and the adjusted reservation for it is $r_i + \beta_i$. Ideally, the unused bandwidth is distributed proportionally to the back logged flows and this is exactly what is conveyed in Equation 1.

Now we define fairness index (α) for a link as

$$\alpha = \frac{\Sigma_{i=0}^{n_1-1}(|\frac{x_i - \lambda_i}{\lambda_i}|) + \Sigma_{i=n_1}^{N-1}(|\frac{x_i - min(\lambda_i, r_i+\beta)}{min(\lambda_i, r_i+\beta)}|)}{N} \tag{2}$$

α measures how close the received bandwidth is to the reservation for a link. The overall fairness index for a switch is calculated as the average α value of its output links:

$$\alpha_{switch} = \frac{\Sigma_{l=0}^{l=K-1}\alpha_{link-l}}{K} \tag{3}$$

where K is the number of outputs for the switch.

The smaller the α is, the better is the fairness. The fairness indices for the settings in Figure 5 and 6 are 0.25 for *i*SLIP and 0.0185 for *i*FS at the input rate of 0.3. The α values are 0.573 and 0 for *i*SLIP and *i*FS respectively under the input rate of 0.4 or above.

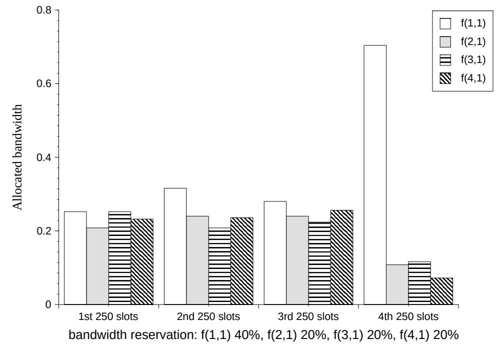

Fig. 7. WPIM: bandwidth distribution within a frame of 1000 slots

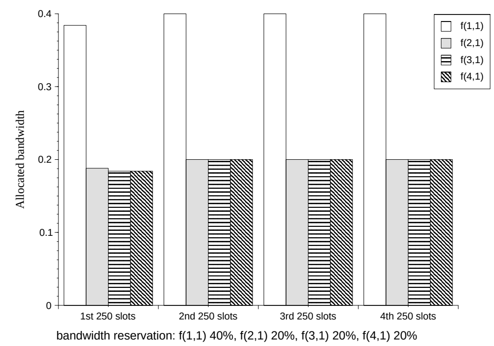

Fig. 8. *i*FS: bandwidth distribution within a frame of 1000 slots

TABLE I
FAIRNESS INDICES FOR WPIM AND *iFS* (SF FOR SUBFRAME)

α	1st SF	2nd SF	3rd SF	4th SF
WPIM	0.21	0.158	0.225	0.568
iFS	0.015	0	0	0

WPIM scheme also complies with the bandwidth requirement of each flow by restricting the number of transmitted cells during a frame within a limit determined by its reservation. Accordingly, bandwidth requirement can be met and the well-behaved flows are protected. However, a careful study reveals the drawback of WPIM whose mechanism rules that link bandwidth is evenly distributed among existing flows until some use up their quotas. Flows running out of their quotas are then excluded from accessing the link in the current frame, and the bandwidth is again allocated equally among the remaining flows. Let us inspect the flow with the largest bandwidth reservation. Under WPIM, this flow shares the link bandwidth equally with all the other flows at the beginning of each frame. Its share increases gradually as the quotas for other flows are exhausted. As an example, consider the four flows going to output 1 but with bandwidth reservations of 40%, 20%, 20% and 20% this time. Frame length is taken as 1000 slots as in [36]. We observe that every flow gets its fair share by the end of a frame under both WPIM and *i*FS. But if we break a frame into four 250-slot subframes, we notice that the bandwidth distribution for WPIM is not fair, as depicted in Figure 7. During the first 750 slots, bandwidth is almost equally distributed(about 25% each) regardless of the reservation. Towards the end of the frame, flows with less reservation($f(2, 1)$,

$f(3,1)$ and $f(4,1)$) have used up their quota. Flow with the highest reservation($f(1,1)$) then consumes almost all the capacity in the last 250 slots. Therefore, WPIM provides bandwidth guarantee at coarse granularity of 1000 slots. The *i*FS, on the other hand, can provide fair sharing both at coarse and fine grain levels. The received bandwidth in *i*FS complies with the reservations even within subframes as illustrated in Figure 8. Their ability to support fair bandwidth is also compared by fairness indices in Table I for individual subframes.

As for all the iterative approaches, the number of iterations required for convergence is a constraining factor because all the iterations must be done within a single switch cycle. In the case of 16×16 switch, 16 iterations are needed in the worst case. Figure 9 shows the effect of the number of iterations on the average cell latency under uniform traffic. We can see that with two iterations, a throughput of over 90% can be achieved and four iterations are adequate to obtain a throughput of nearly 100%. This is consistent with the findings in [1], [36] and [22] that $log(N)$ iterations are enough for the algorithm to converge.

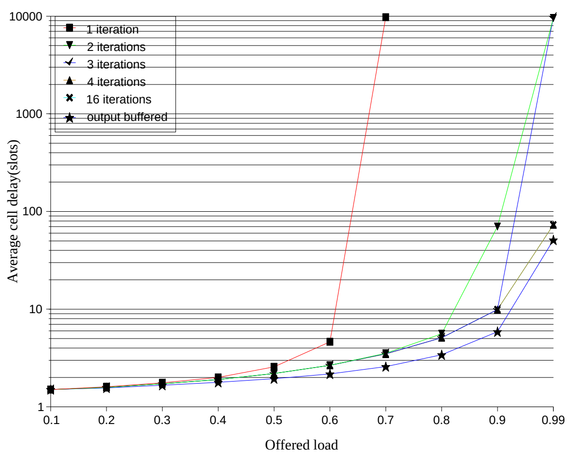

Fig. 9. Latency versus the number of iterations for *i*FS

In summary, the evaluation using uniform and non-uniform synthetic traffic models indicates that the *i*FS scheme is very promising in supporting fair bandwidth distribution as well as maintaining high overall throughput.

B. Measurement from real traffic

Studies using the Internet traffic traces have been extensively reported in network research. Yet, to the best of our knowledge, ours is the first attempt to incorporate such traces into evaluating switch/router scheduling schemes. We hope this practice can shed some light on the effective use of real traces in validating the performance of different approaches.

Traffic arrivals in the Internet have been shown to be highly correlated(self-similar) [31], [7]. Simulations using Poisson or Bernoulli distribution offer good judgment on how a scheme works, but may significantly underestimate the burstiness of the real traffic pattern and give rise to unrealistic performance results. Geometrically distributed ON/OFF traffic is used in [22] and [14] to model the burstiness, but we will show later that the approach of taking real traces is better in characterizing the impact of traffic on switch design.

B.1 Measurement from NLANR traces

The traces taken from National Lab of Applied Network Research(NLANR [27]) are collected using OC3mon, a traffic monitor on OC-3 link, at the ATM backbones on NSF vBNS. Traces from the same site were used in [17] for studying packet forwarding method. The traces considered in this paper were collected in May 2000 from facilities AIX, FRG and MRT. Every line in the trace file includes the following information: timestamp when a TCP header arrives at the OC3mon, source IP/port, destination IP/port and size of the packet. We use four traces as arriving traffic on four input links to our 4×4 switch. The overall average utilization for these input links is listed in the first row of Table II. Based on the destination IP address, we classify traffic on each link into four flows, one for each output port(i.e., flow $f(i,j)$ from link i going to output j). The link utilization for the individual flows are also listed in Table II. We derive each burst length from packet size as they are segmented into ATM cells. Each idle period is calculated as the interarrival time between two packet headers minus the burst length of the prior packet. Unfortunately, the traces carry no information about the bandwidth reservation so we impose our assumption of reserved bandwidth in the simulation.

TABLE II
INPUT LINK UTILIZATION FOR THE VARIOUS FLOWS(NLANR)

utilization	link 1	link 2	link 3	link 4
overall	19.5%	20.5%	19.3%	18.6%
$f(i,1)$	7.6%	8.5%	18.9%	13.1%
$f(i,2)$	4.2%	4.5%	0.2%	3.5%
$f(i,3)$	4.7%	3.6%	0.09%	1.4%
$f(i,4)$	3.0%	3.8%	0.07%	0.6%

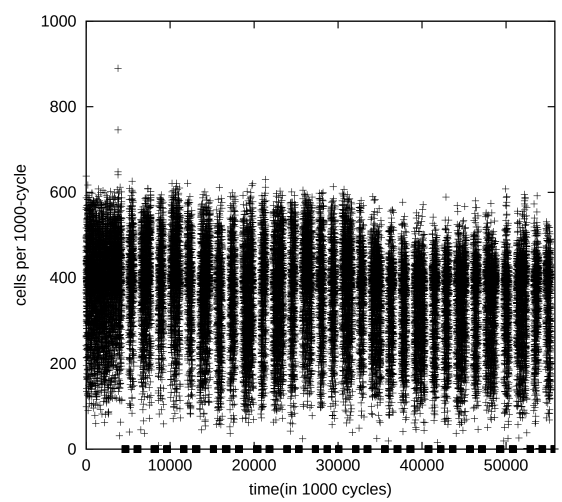

Fig. 10. The arrival pattern for traffic from link 1(NLANR)

To see how real traffic is different from geometrically distributed ON/OFF traffic, we give an example by plotting cell arrival patterns for input link 1 and its corresponding geometric counterpart in Figure 10 and Figure 11, respectively. The geometrical ON/OFF traffic is so generated that it has the same average burst length and idle period as those from the trace. We observe that the ON time is much more clustered in the trace and the typical number of cells coming within 1000 cycles ranges from 100 to 600. There are also non-negligible times when not a single cell shows up during a 1000-cycle period. For geometrically distributed ON/OFF arrival pattern, on the other hand, the interarrivals are more evenly spread out with only 100 to 350 cells incoming within every 1000 cycles. Studies from other links tell similar differences between traces and synthetically generated traffic.

To examine the impact of the distinct traffic pattern, we feed the flows from the traces and the geometrically distributed ON/OFF traffic to the simulator. Figure 12 presents the delay for flows to output 1 under both cases. The average cell delay is substantially greater for the real traffic. Geometrically distributed ON/OFF traffic underestimates the contention of the flows and therefore cannot be used very well as a

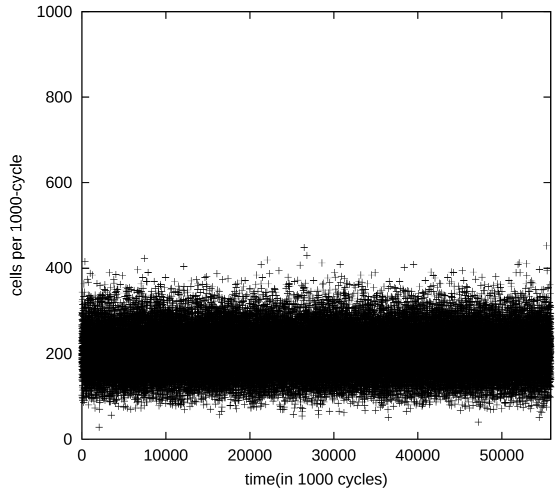

Fig. 11. The arrival pattern for geometric traffic from link 1

representative model.

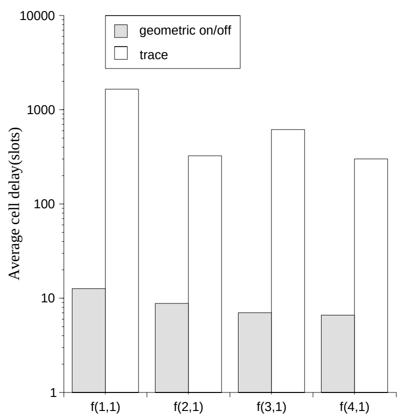

Fig. 12. Average cell delay for trace vs. geometrically ON/OFF traffic

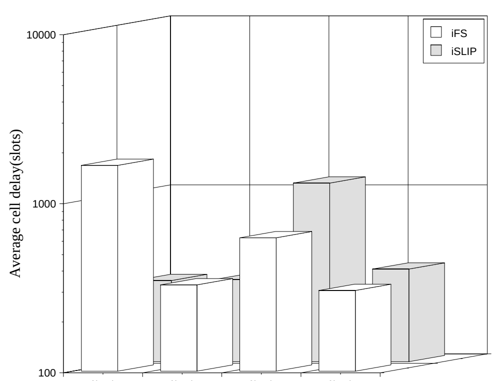

Fig. 13. Internet traffic: average cell delay for *i*FS and *i*SLIP(NLANR)

Now we study effectiveness of the *i*FS scheme for maintaining fair bandwidth allocation for real traffic. On top of the flows presented in Table II, we impose a bandwidth requirement arbitrarily. Let the bandwidth requirement for flows $f(1,1)$, $f(2,1)$, $f(3,1)$ and $f(4,1)$ be 10%, 20%, 30% and 40%, respectively. As the sum of the workload is well below the capacity of output link 1, throughput for each flow is equal to offered input rate. However, since some flows reserve a greater fraction of bandwidth than others, these flows should receive a larger

TABLE III

LINK UTILIZATION FOR THE FLOWS AFTER CUTTING DOWN IDLE TIME(NLANR)

utilization	link 1	link 2	link 3	link 4
$f(i,1)$	24.7%	27.1%	48.3%	37.6%
$f(i,2)$	15.0%	15.8%	0.78%	12.8%
$f(i,3)$	16.6%	13.2%	0.35%	5.6%
$f(i,4)$	10.9%	13.5%	0.28%	2.2%

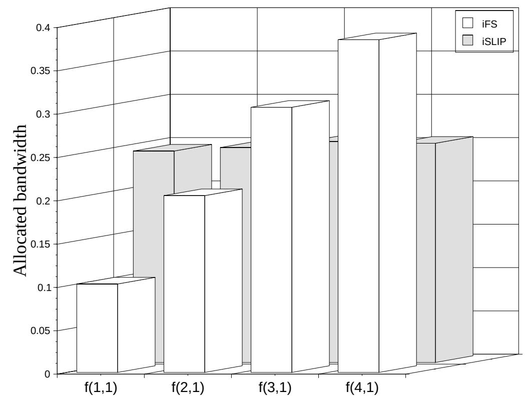

Fig. 14. Internet traffic: received bandwidth distribution(NLANR)

bandwidth and perceive less delay. Figure 13 shows the average cell delay for the flows to output 1 under *i*FS and *i*SLIP schemes. Note that *i*SLIP does not observe the bandwidth requirement. Flows $f(1,1)$ and $f(2,1)$ have lower delay because their input rate is only about 8%. Flow $f(3,1)$ experiences much longer delay than $f(4,1)$ since the it has higher input rate than the latter(18.9% vs. 13.1%). In contrast with the situation for *i*SLIP, the delay for the individual flows in *i*FS depends not only on the input rate, but also on the bandwidth requirement. It is shown in Figure 13 that flow $f(1,1)$ has much greater average cell latency than that of $f(2,1)$ even though their input rates are close. This is because $f(1,1)$ reserves only 10% of the bandwidth whereas 20% is set aside for $f(2,1)$. For the same reason, $f(3,1)$ suffers larger delay than $f(4,1)$ since the it has higher input rate(18.9% vs. 13.1%) but reserves less bandwidth(30% vs. 40%). However, a greater reservation itself cannot guarantee lower delay. Look at $f(2,1)$ and $f(3,1)$ for a case in point. Compared to $f(2,1)$ with a reservation of 20%, $f(3,1)$ has larger average cell delay even if it reserves 30% of the bandwidth. The actual input workload for $f(3,1)$ is so much higher than $f(2,1)$ that its reservation is not great enough to bring in a lower delay.

In order to demonstrate how *i*FS enforces fair bandwidth assignment, we need flows to send packets at a higher rate than reservation. However, the highest link utilization from the traces available at NLANR is only around 20%, and it becomes even lower after being split over four outputs. We get around this problem by cutting down each idle period to 25% of the original time in the simulation. The resulting average link utilization of various flows are listed in Table III. We believe this manipulation on the traces will not compromise the quality of the real traffic. Assume the same bandwidth reservations of 10%, 20%, 30% and 40% for $f(1,1)$, $f(2,1)$, $f(3,1)$ and $f(4,1)$, respectively. The achieved bandwidth of the flows is plotted in Figure 14 for *i*FS and *i*SLIP schemes. It is observed that *i*FS is capable of supporting fair bandwidth allocation when the workload exceeds the capacity of the shared link and the received bandwidth for the individual flows is proportional to the reserved share. The *i*SLIP fails to do so and distributes the bandwidth evenly among the various flows due to the mechanism of rotating priority.

TABLE IV

LINK UTILIZATION FOR THE FLOWS AFTER CUTTING DOWN IDLE TIME(UCB)

utilization	link 1	link 2	link 3	link 4
$f(i,1)$	23.3%	25.9%	36.1%	35.9%
$f(i,2)$	8.0%	22.1%	21.1%	14.0 %
$f(i,3)$	5.1%	19.9%	20.5%	15.8%
$f(i,4)$	28.5%	9.3%	20.5%	16.4%

B.2 Measurement from UCB traces

Next, we consider Web proxy traffic from HTTP traces gathered by UC Berkeley in November 1996 from its Home IP service [37]. The community gains IP connection across about 600 modems (with speed from 2.4Kb/s to 28.8Kb/s), and all their traffic ends up going through a single 10Mb/s shared Ethernet segment, on which a network monitoring computer is placed. Interested fields from the traces include the time when the first byte of a HTTP response data file was seen, the time when the last byte of data was seen, anonymized client and server addresses and the size of response data file. We ignore the effect of HTTP requests because their consumed bandwidth is negligible.

We do the following processing to the trace. According to the maximum transmission unit of the Ethernet, we segment each data file into a series of 1500-byte packets. Since our simulator is cycle-based, each packet is further divided into cells which are sent back to back as a burst. The idle time between two consecutive bursts is determined as if the packets belonging to a file are evenly spaced between the time of the first byte seen and the last byte seen. The idle period between the files are calculated directly from the trace. The interleaving of packet transmission from multiple files is also taken into account. Again, we assume a 4×4 switch and classify traffic depending on the IP address and then feed the traces into the four input links. Due to the same reason for NLANR traces, the idle time is cut down to obtained higher link utilization, as shown in Table IV.

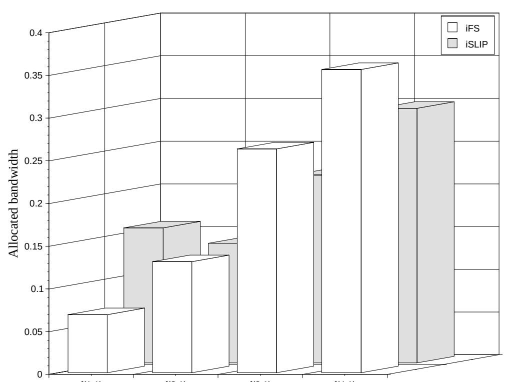

Fig. 15. Internet traffic: received bandwidth distribution(UCB)

Assume the same bandwidth reservations of 10%, 20%, 30% and 40% for $f(1,1)$, $f(2,1)$, $f(3,1)$ and $f(4,1)$, respectively. The achieved bandwidth of the flows from link 1 is plotted in Figure 15 for *i*FS and *i*SLIP schemes. As expected, with *i*FS scheduling, the more a flows reserves, the more bandwidth it entitles to. But one thing brought to our attention is that, although $flow(1,1)$ has an average link utilization of 23.3% and has reserved 10% of the link bandwidth, its perceived bandwidth is only 6.8%. We examine the traffic pattern for this flow and plot it in Figure 16. It is observed that in approximately the first one third of the time, there are almost no packet arrivals from $flow(1,1)$. So $f(1,1)$ has nothing to send during that period. After that, its traffic increases dramatically. However, due to the reservation of 10% only, $flow(1,1)$ is restricted by the *i*FS from transmitting too much because

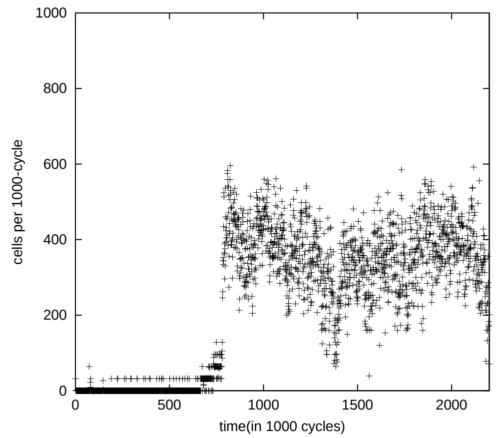

Fig. 16. The arrival pattern for traffic from flow(1,1)

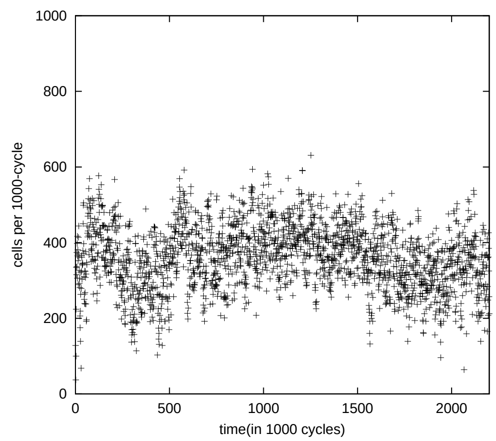

Fig. 17. The arrival pattern for traffic from flow(4,1)

other completing flows are also over subscribed. Therefore, the resulting *average* received bandwidth over the entire monitored interval is only 6.8%. With *i*SLIP, $flow(1,1)$ grasps 25% of the bandwidth in its active period and attains 15.8% of the link bandwidth on average, well above the preserved 10%. According to the trace collector, the uncharacteristically low activity in the traces corresponds to network outrages from Berkeley's ISP, rather than from trace failures [37]. For comparison, we plot an example of a normal flow arrival pattern in Figure 17. We can see that even under such skewed scenario, *i*FS is able to distribute resources more fair than *i*SLIP.

V. BUFFER MANAGEMENT FOR FAIR SCHEDULING

Section III focuses on the switch scheduling scheme itself without considering the buffer size. Yet , in practice the input buffer is finite. With rate-based flow control, which is the common choice for supporting bandwidth distribution, excessive packets are dropped when buffer is full or congestion is anticipated. In the following, we study four selective packet discarding mechanisms and examine their impact on fair bandwidth allocation.

A. Decongestion Mechanisms

While messages from a source node are fragmented into fixed cells, which are transmitted individually across the network and reassembled at the destination, cell level retransmission is not supported. Thus, loss of a single cell in a packet forfeits the whole packet and it has to be retransmitted. Four decongestion mechanisms schemes are investigated: *drop tail* (DT), *early packet discard* (EPD), *equal size per flow* (ESPF) and *rate based size per flow* (RSPF). DT and EPD are stateless whereas

ESPF and RSPF are on per-flow basis.

Drop tail: This is the basic dropping strategy: Cells are first-in-first-out and an incoming cell is dropped if it arrives to find the input buffer full. After a cell is shredded, the switch still makes effort to transmit the remaining part of the packet even if they turn out to be worthless at the destination and the entire packet has to be retransmitted. Hence, DT scheme is poor in performance despite its simplicity in implementation without keeping status for each flow.

Early packet discard [11]: EPD overcomes the the drawback of DT by dropping all the cells constituting a new packet when congestion is predicted. Upon receiving a header cell of a new packet, the switch checks to see whether the buffer occupancy exceeds certain threshold. If so, the header cell is dropped and so are the upcoming cells from the same packet. Otherwise, the header is inserted into the buffer and subsequent cells are allowed into the buffer as long as the it is not full upon their arrival. The justification behind this *early* discard is that because the buffer is almost full and the congestion is likely to occur, the upcoming cells of the packet are most probably dropped. So it is better to give up sooner than later the packet that cannot be received successfully any way. A good side effect is that, once the packet is dropped, it makes room in the buffer for other upcoming packets so that they are less likely to be discarded. Like drop tail mechanism, EPD is also a stateless scheme. But we need to watch for buffer occupancy constantly and keep track of the packets that have been selected for discard so that their upcoming cells are taken care of. Therefore, EPD is more intricate in implementation compared to DT. It will been seen later with simulation that the extra effort pays off.

Equal size per flow: Both drop tail and early packet discard are stateless, and they are very much handicapped to the degree to achieve flow isolation, which is important in order to prevent ill-behaved flow to take advantage of others. An EPD-based per flow queueing approach called ESPF secures equal share of buffer space after certain threshold is reached. In ESPF, the share is calculated as the total buffer size divided by the number of concurrent active flows. Thus, this is a *dynamic* reservation policy in which the quota for each flow changes with the number of active flows. When buffer occupancy is below certain threshold, all incoming cells are accepted. After that, a header cell is allowed into the buffer only if the flow's quota has not been used up. Otherwise, the entire packet is discarded. Once the header cell is admitted, subsequent cells can follow in provided that the buffer is not full. One concern of ESPF in terms of implementation complexity is the necessity to compute the quota for the active flows on the fly because the flows are on and off.

Rate-proportional size per flow: ESPF attempts to assign equal share of the buffer space to the active flows without taking their desired bandwidth into consideration. Intuitively, some flows may deserve larger fraction of buffer than others because they reserve a greater portion of link bandwidth. In this sense, it is not necessarily fair to treat all the flows equally when allotting the buffer space. With the scheme of rate-proportional size per flow (RSPF), the quota is assigned in proportion to the fraction of bandwidth reservation. Same as ESPF, if buffer occupancy is low, all incoming cells are admitted. After the threshold is reached, a packet is allowed into the buffer only if its quota has not been exhausted. In contrast to ESPF, quota for each flow is fixed by the time the reservation is made in RSPF. Therefore, it is less computational intensive in implementation .

B. Performance Analysis

The decongestion mechanisms are evaluated based on the simulation with a 4×4 switch. From each of the four inputs there are four flows going to different outputs, amounting to a total of 16 flows. We first examine the various schemes using geometric on/off traffic for average cell delay, packet loss ratio and the ability to support fair bandwidth allocation. Then we extend our study to employ real network traffic as input. The scheduling scheme used throughout this section is *i*FS and the threshold is chosen to be 80% of the buffer size.

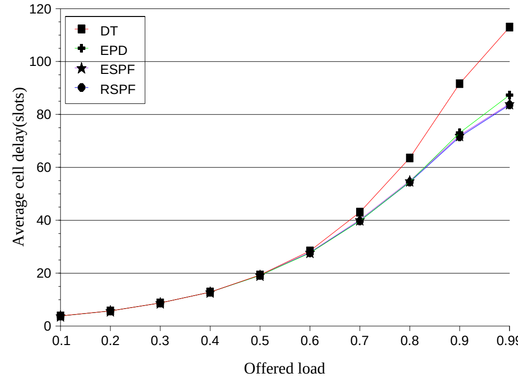

Fig. 18. Average cell delay versus offered load

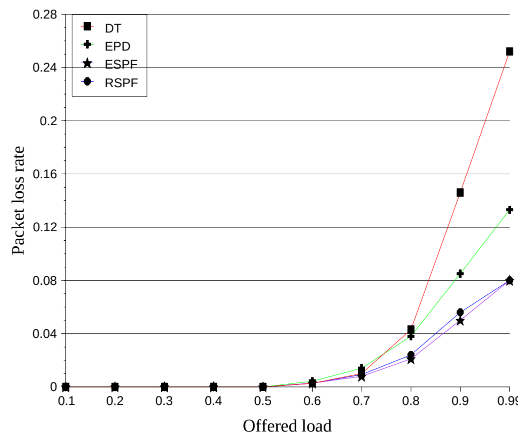

Fig. 19. Packet loss ratio versus offered load

Let us consider a scenario with a set of benign flows, where each of the 16 connections reserves an equal share of its intended link bandwidth and all the flows keep the same actual average traffic rate all the time. The workload is geometrically distributed on/off traffic and the average *on* period is 20 consecutive cells in a burst. The buffer size for each input block is set to 100 slots. We vary the offered input rate by changing the mean *off* time and examine the average cell delay and packet loss ratio, which is defined to be the ratio of the number of lost packets to total number of packets sent by the source. From the results in Figure 18 and Figure 19, it is observed that under the input rate of 0.6, there is little difference among the four schemes. As the input rate further increases, the delay for DT is far more worse than the others. In DT, incoming cells are dropped only when the buffer is full. Therefore, even if a header cell has already been discarded, constituent cells of the same packet may still clog in the buffer and compete to pass the crossbar. Such cells, eventually thrown away at the destination, worthlessly obstruct the way of useful cells from other packets and delay their transmission. The other three schemes all employ early packet discard technique where useless cells are detected at early stage and dropped so that delivery of other packets is expedited. Since all the flows reserve equal share of bandwidth, RSPF becomes almost identical to ESPF under this circumstances, so it is not a surprise that ESPF and RSPF give nearly the same performance. However, careful study reveals that ESPF offers slightly lower packet loss ratio. We trace this to the fact that buffer quota is dynamically assigned to the flows in ESPF. The active flows in ESPF can take advantage of the *off* flows and allow more packets in the buffer than in the case of RSPF.

Both ESPF and RSPF outperform EPD under heavy workload. With EPD, a new packet is discarded when the threshold is reached, whereas in ESPF and RSPF, an incoming packet can still be admitted provided that the quota for its flow has not been exhausted. Therefore, with the same threshold, ESPF and RSPF allow more packets into the buffer. And it explains why EPD has much higher packet loss ratio than ESPF and RSPF in Figure 19. It is possible to set a higher threshold in order for EPD to achieve better buffer utilization, but again the same problem with DT may occur. Simulation with longer burst length (not shown here due to space limitation) indicates similar performance trend. Although longer burst results in larger cell delay and greater loss ratio for all the four schemes, DT is more sensitive to burstiness than others.

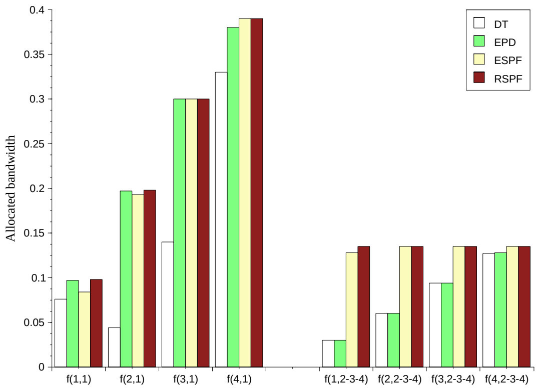

Fig. 20. The ability for fair bandwidth allocation under synthetic workload

TABLE V
EFFECT OF DECONGESTION MECHANISM ON FAIRNESS UNDER SYNTHETIC WORKLOAD(NLANR)

α	DT	EPD	ESPF	RSPF
link 1	0.420	0.024	0.055	0.014
link 2,3,4	0.480	0.480	0.110	0.100
switch overall	0.465	0.366	0.096	0.079

The performance of the various decongestion mechanisms are inspected in supporting fair bandwidth distribution using geometric on/off workload. Following the notation of the tagged flows in Section IV, it is assumed $f(1,1)$, $f(2,1)$, $f(3,1)$ and $f(4,1)$ reserve 10%, 20%, 30% and 40% of the output link 1 bandwidth respectfully. Each of the rest flows reserves a fraction of 15% of their intended link bandwidth. Suppose that the tagged flows fail to keep the contract and each sends at a rate of 0.4 but others abide their promises. The perceived bandwidth for the flows are presented in Figure 20. Without flow isolation, DT performs the worst and cannot support bandwidth distribution according to the reservation. For example, compared to their respective reservation of 10% and 20%, $f(1,1)$ receives 7.6% and $f(2,1)$ gets 4.4%. Yet, the figure shows that the various flows do receive somewhat different bandwidth using DT. This is attributed to the fair scheduling scheme being used, which attempts to favor flows with greater reservation. EPD improves over DT by supporting fair bandwidth distribution among the tagged flows. However, if we look at untagged flows from input link 1 and 2, they receive bandwidth well below the reservation. The reason is that the actual sending rate from tagged flows of input 1 and 2 is so much higher than their reservation that many of their packets are backlogged. Since EPD is not per flow based, such packets clog the buffer and prevent cells from the untagged flows to get in. Eventually packets from the untagged flows are discarded, resulting in underutilization of their allocated bandwidth. With per-flow queueing, ESPF and RSPF sucessfully protect the benign flows in this scenario. The fairness indices for the four policies are given in Table V. One may wonder why ESPF, without taking any reservation, works almost equally well as RSPF. Intuitively, because fair scheduling is employed, flows reserving greater fraction of bandwidth also receive higher service rate and cells are drained more quickly. As a result, such flows do not *necessarily* require buffer size in proportion to the arrival rate.

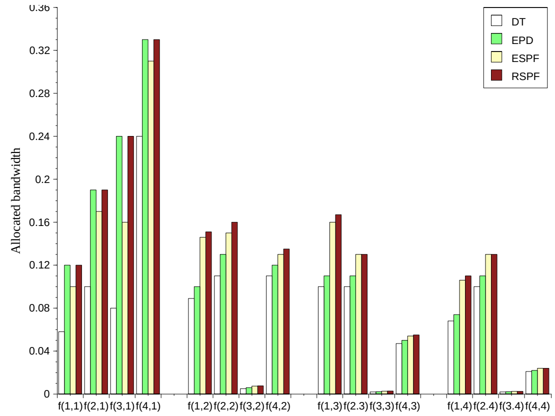

Fig. 21. The ability for fair bandwidth allocation for Internet traffic(NLANR)

TABLE VI
EFFECT OF DECONGESTION MECHANISM ON FAIRNESS UNDER REAL WORKLOAD(UCB)

α	DT	EPD	ESPF	RSPF
link 1	0.500	0.140	0.220	0.140
link 2	0.293	0.200	0.023	0.000
link 3	0.310	0.247	0.079	0.058
link 4	0.240	0.180	0.043	0.036
switch overall	0.336	0.192	0.086	0.059

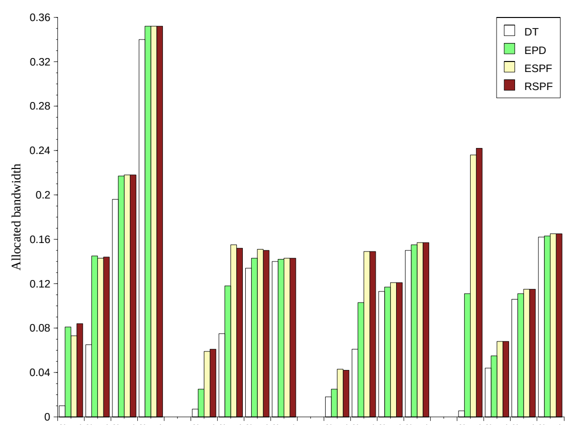

Fig. 22. The ability for fair bandwidth allocation for Internet traffic(UCB)

TABLE VII
EFFECT OF DECONGESTION MECHANISM ON FAIRNESS UNDER REAL WORKLOAD(UCB)

α	DT	EPD	ESPF	RSPF
link 1	0.400	0.270	0.240	0.250
link 2	0.440	0.330	0.170	0.174
link 3	0.460	0.320	0.160	0.160
link 4	0.508	0.240	0.260	0.230
switch overall	0.468	0.290	0.207	0.203

Finally, the various decongestion mechanisms are compared using the Web traces as in Section IV-B. For NLANR packet header traces, the link utilization can be found in Table III. The bandwidth reservation imposed is assumed to be the same with the last setting, i.e., 10%, 20%, 30% and 40% for the tagged flows and 15% otherwise. The perceived bandwidth for the individual flows under different schemes is plotted in Figure 21. Significant difference between ESPF and RSPF is seen, especially for flow $f(3,1)$, whose traffic is the most intensive and clustered. Rate based buffer space allocation gives better results under real traffic when the incoming packets are highly bursty. Although flows with greater bandwidth reservation drain their cells faster under fair bandwidth scheduling, it may not be fast enough to offset the fact their cells arrive in a dramatically clustered fashion. The fact that EPD provides a more fair bandwidth allocation than ESPF for output link 1 (α value 0.14 vs. 0.22 in the first row of Table VI) also proves this to be true since flows with higher input rate are allowed to take more buffer space in EPD but not in ESPF. Therefore, equal share of buffer is not capable of maintaining fairness for such flows and it is beneficial to provide larger buffer for those with highly bursty arrival pattern.

We assume the same bandwidth requirement for the flows in UCB Web proxy traces and rerun the simulation using various decongestion policies. The results are presented in Figure 22 and Table VII. We again found that both RSPF and ESPF offer better performance than EPD. RSPF outperforms ESPF for some flows and is little bit inferior for others. We attribute this to the traffic characteristics of the individual flows.

The above experiments indicate that decongestion mechanism can significantly affect the switch performance, and fair scheduling scheme alone cannot guarantee fairness. Early packet discard, which drops worthless packets at an earlier stage and relieves the network congestion, is a must. Flow isolation is important to support fair link bandwidth distribution. Under the situations when the traffic is not highly bursty, equal sharing of buffer space is sufficient if fair scheduling scheme is used. However, when cell arrival is highly clustered, equal buffer sharing will force the flows with high input rate to drop more packets. This can be avoided to a large extent if buffer allocation is in proportion to the bandwidth reservation.

VI. Concluding Remarks

In this paper, we explored the scheduling schemes for input buffered switches to support fair bandwidth allocation. We first proposed an iterative fair scheduling(*i*FS) algorithm capable of scheduling cells so that each flow receives bandwidth proportional to its reservation under heavy traffic. We showed that the fairness support does not compromise the average cell latency when compared with other iterative scheduling schemes. We examined four decongestion mechanisms and studied their impact on supporting fair bandwidth scheduling. With the extensive experimental results, we showed that early packet discard is necessary to relieve congestion. Our simulation also tells that fair scheduling scheme alone cannot ensure fair bandwidth allocation and per flow buffering is needed to protect well-behaved flows. Buffer allocation based on the bandwidth reservation offers better fairness especially when the workload is highly clustered. It is worth mentioning that we explored the issues of switch design by incorporating the Web traffic traces in the study. Although it has been applied in other aspect of network research, it has not been used in the study of this area before. Results from real trace workload provides further validation in addition to commonly employed traffic model.

References

[1] T. E. Anderson, S. S. Owicki, J. B. Saxe, and C. P. Thacker. High speed switch scheduling for local area networks. *ACM Transctions on Computer Systerms*, 11(4):319–352, November 1993.

[2] M. F. Arlitt and L. Williamson, Carey. Internet web servers: Workload characterization and performance implications. *IEEE/ACM Transactions on Networking*, 5(5):631–645, October 1997.

[3] G. Armitage and K. Adams. Packet reassemly during cell loss. *IEEE Network Magazine*, 7(5):26–34, September 1993.

[4] P. Barford and M. Crovella. Generating representative web workloads. In *Proceedings of the 1998 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 151–160, Madison, WI, July 1998.

[5] J. Bennett and H. Zhang. WF^2Q: worst-case fair weighted fair queueing. In *Proceedings of IEEE INFOCOM'96*, pages 120–128, San Francisco, CA, March 1996.

[6] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web caching and Zipf-like distributions: Evidence and implications. In *Proceedings of IEEE INFOCOM'99*, pages 126–134, March 1999.

[7] M. E. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.

[8] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. *Journal of Internetworking Research and Experience*, 1(1):3–26, September 1990.

[9] J. Ding and L. N. Bhuyan. Evaluation of multi-queue buffered multistage interconnection networks under uniform and non-uniform traffic patterns. *International Journal of Systems Science*, 28(11):1115–1128, 1997.

[10] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1:397–413, August 1993.

[11] S. Floyd and A. Romanov. Dynamics of TCP traffic over ATM networks. *IEEE Journal on Selected Areas in Communications*, 13(4):633–641, 1995.

[12] L. Georgiadis, I. Cidon, R. Guerin, and A. Khamisy. Optimal buffer sharing. *IEEE Journal on Selected Areas in Communications*, 13:1229–1240, September 1995.

[13] S. J. Golestani. A self-clocked fair queueing scheme for broadband applications. In *Proceedings of IEEE INFOCOM'94*, pages 636–646, Toronto, Canada, June 1994.

[14] M. W. Goudreau, S. G. Kolliopoulos, and S. B. Rao. Scheduling algorithms for input-queued switches: Randomized techniques and experimental evaluation. In *Proceedings of IEEE INFOCOM'00*, volume 3, pages 1634–1643, Tel Aviv, Israel, March 2000.

[15] P. Goyan, H. Vin, and H. Chen. Start-time fair queueing: A scheduling algorithm for integrated services. In *ACM SIGCOMM*, pages 157–168, Palo Alto, CA, August 1996.

[16] P. Gupta and N. McKeown. Design and implmentation of a fast crossbar scheduler. *IEEE Micro*, pages 20–28, Janurary/Febrary 1999.

[17] F. Hoymany and D. Mosse. A simulation study of packet forwarding methord over ATM. In *Proceedings of IEEE INFOCOM'98*, pages 401–408, San Francisco, CA, March 1998.

[18] M. J. Karol, M. G. Hluchyj, and S. P. Morgan. Input versus output queueing on a space-division packet switch. *IEEE Transactions on Communications*, COM-35(12):1347–1356, December 1987.

[19] S. Keshav and R. Sharma. Issues and trends in router design. *IEEE Communications Magazine*, pages 144–151, May 1998.

[20] S. Li and N. Ansari. Input-queued switching with QoS guarantees. In *Proceedings of IEEE INFOCOM'99*, pages 1152–1159, New York City, NY, March 1999.

[21] A. Mahanti, C. Williamson, and D. Eager. Traffic analysis of a web proxy caching hierarchy. *IEEE Network Magazine*, pages 16–23, May/June 2000.

[22] N. McKeown. The iSLIP scheduling algorithm for input-queued switches. *IEEE/ACM Transactions on Networking*, 7(2):188–201, April 1999.

[23] N. McKeown, V. Anantharam, and J. Walrand. Achieving 100% throughput in an input-queued switch. In *Proceedings of IEEE INFOCOM'96*, pages 296–302, San Francisco, CA, March 1996.

[24] N. McKeown, M. Izzard, A. Mekkittikul, W. Ellersick, and M. Horowitz. Tiny tera: a packet switch core. *IEEE Micro*, pages 26–33, January/February 1997.

[25] C. Metz. IP routers: New tool for gigabit networking. *IEEE Internet Computing*, pages 14–18, November/December 1998.

[26] P. Misha and M. Saksena. Designing buffer management policies at an IP/ATM gateway. In *Proceedings of IEEE ATM Workshop*, pages 144–153, 1998.

[27] National Lab for Applied Network Research (NLANR) packet header traces, http://moat.nlanr.net/traces/.

[28] R. Pan, B. Prabhakar, and K. Psounis. CHOKe, a stateless active queue management scheme for approximating fair bandwidth allocation. In *Proceedings of IEEE INFOCOM'00*, pages 942–951, Tel Aviv, Israel, March 2000.

[29] A. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in intergrated services networks: The single-node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, June 1993.

[30] C. Partridge, P. P. Carvey, E. Burgess, I. Castineyra, T. Clarke, L. Graham, M. Hathaway, P. Herman, A. King, S. Kohalmi, T. Ma, and J. Mcallen. A 50-Gb/s IP router. *IEEE/ACM Transactions on Networking*, 6(3):237–248, June 1998.

[31] V. Paxson and S. Floyd. Wide area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.

[32] L. L. Peterson and B. S. Davie. *Computer Networks — A System Approach*. Morgan Kaufmann, San Francisco, CA, 2000.

[33] M. Pirvu, N. Ni, and L. N. Bhuyan. Exploring the switch design space in a CC-NUMA multiprocessor environment. In *Proceedings of International Parallel and Distributed Processing Symposium*, pages 703 –710, Cancun, Mexico, May 2000.

[34] D. Serpanos and P. Antoniadis. FIRM: A class of distributed scheduling algorithm for high-speed ATM switches with multiple input queues. In *Proceedings of IEEE INFOCOM'00*, volume 2, pages 548–555, Tel Aviv, Israel, March 2000.

[35] D. C. Stephens and H. Zhang. Implementing distributed packet fair queueing in scalable switch architecture. In *Proceedings IEEE INFOCOM'98*, pages 282–290, San Francisco, CA, March 1998.

[36] D. Stiliadis and A. Varma. Providing bandwidth guarantees in an input-buffered crosbar switch. In *Proceedings of IEEE INFOCOM'95*, pages 960–968, Boston, MA, April 1995.

[37] UC Berkeley Home IP Web Traces, http://ita.ee.lbl.gov/html/contrib/UCB.home-IP-HTTP.html.

[38] H. Zhang. Service displines for guaranteed performance service in packet-switching networks. In *Proceedings of the IEEE*, volume 83, pages 1374–1396, October 1995.