

On the Computational Complexity of Maintaining GPS Clock and Providing Tight Delay Bounds in Packet Scheduling

Jun Xu Qi Zhao
 College of Computing
 Georgia Institute of Technology
 {jx, qzhao}@cc.gatech.edu

Abstract—Packet scheduling is an important mechanism for providing QoS guarantees in data networks. A scheduling algorithm in general consists of two functions: one estimates how the GPS (General Processor Sharing) clock progresses with respect to the real time, and the other decides the order of serving the packets based on the estimation of their GPS start/finish times. In this work, we answer important open questions concerning the computational complexity of performing both functions. The first part of our work systematically studies the complexity of computing the GPS virtual start/finish times of the packets, which is long believed to be $\Omega(n)$ per packet but has never been proved or explicitly refuted. It also answers several other related open questions such as “whether the complexity can be lower if we only want to compute the relative order of the GPS finish times of the packets rather than their exact values?” The second part of our work studies the inherent complexity for scheduling algorithms to guarantee tight delay bounds. We extend the prior work by Xu and Lipton to a stronger and more practical computational model and explore related issues. We show rigorously that existing methodologies used in prior work will not be suitable for establishing lower bound results under the new model.

Index Terms— Quality of Service (QoS), Packet Scheduling, Delay Bound, Computational Complexity, Linear Decision Tree, General Processor Sharing

I. INTRODUCTION

Service discipline (i.e., packet scheduling) is an important mechanism for providing QoS guarantees such as end-to-end delay bounds and fair bandwidth allocation, in computer networks [4], [8],

[15]. The perfect service discipline for providing delay and fairness guarantees is General Processor Sharing (GPS) [4], [9]. In a GPS scheduler, all backlogged sessions are served simultaneously in a weighted fashion as follows. In a link of rate r served by a GPS scheduler, each session i is assigned a weight value ϕ_i . Each backlogged session i at every moment t is served simultaneously at rate $r_i = r \phi_i / (\sum_{j \in B(t)} \phi_j)$, where $B(t)$ is the set of sessions that are backlogged at time t . However, GPS is not a realistic service discipline since in a packet network, service is performed packet-by-packet, rather than “bit by bit”. Nevertheless, GPS serves as a reference scheduler that real-world packet-by-packet service disciplines (e.g., *WFQ* [4]) can be compared with in terms of delay and fairness.

A real-world packet-by-packet service disciplines typically consists of the following two functions.

- 1) **Tracking GPS time:** This function is to track the progress of GPS virtual time (described later) with respect to the real time. Its main objective is to estimate the GPS virtual start and finish times of a packet, which is the time that a packet should have started and finished service if served by a GPS scheduler.
- 2) **Scheduling according to GPS clock:** This function is to schedule the packets based on the estimation of their GPS virtual finish/start times. For example, *WFQ* selects, among the packets currently in queue, the one with the lowest GPS virtual finish time, to serve next.

In this paper, we study two sets of open prob-

lems concerning the computational complexity of performing each of the aforementioned two functions. Through this endeavor, we will also exhibit the deep structural connections between them. Our work consists of the following two themes, which will be introduced in greater detail in Sec. III and IV, respectively.

- 1) We study the open problem concerning the complexity lower bound for computing the GPS virtual finish times of the packets. This complexity is long believed to be $\Omega(n)$ per packet [11], [12], [3], [9], where n is the number of sessions. For this reason, many scheduling algorithms such as WF^2Q+ [2], SFQ [7], FFQ [11], and $SCFQ$ [6] only approximate the GPS clock (with certain error), trading accuracy for lower complexity. However, it has never been carefully studied whether the complexity lower bound of tracking GPS clock perfectly is indeed $\Omega(n)$ per packet. Our work not only settles this question, but also leads to other surprising discoveries.
- 2) We investigate the computational complexity of providing tight delay bounds (e.g., $O(1)$) in service disciplines. Our initial goal was to extend a prior lower bound result on this problem [14] to a stronger computation model. This extension is practically important since the stronger model accurately captures the computation power of a router. We instead establish an interesting “negative” result (the prior results of [14] viewed as “positive”), showing that the methodology used in [14] for proving the lower bound of $\Omega(\log n)$ per packet will not work for the stronger model.

In the sequel, we will distinguish between a *service discipline* and a *scheduling algorithm*. The former refers to the policy as to which packet should be served next, whereas the latter refers to actual mechanism to carry out this policy. For example, WFQ service discipline says that packets should be scheduled in the sorted order of their GPS virtual finish times. A WFQ algorithm specifies the data structure and algorithm that is used to track the GPS clock and to sort the GPS timestamps. This distinction is critical in the context of this paper, since we will pose and answer questions such as “Is this scheduling algorithm asymptotically optimal in computa-

tional complexity for carrying out the said service discipline?”

A. Complexity of Estimating GPS clock

Tracking the GPS clock is an important task in service disciplines. How well this task is performed directly affects the QoS guarantees provided. For example, it is shown in [14] that, among all the service disciplines, only WFQ and WF^2Q provides a tight GPS-relative delay (defined in Sec. II.A) of $O(1)$, because they keep perfect track of GPS clock. All other disciplines, such as [11], [6], [2], [7], only approximate GPS clock (with certain error), resulting in much less tight delay bounds.

This prompts us to study the following open problem: “What is the complexity lower bound for computing the GPS virtual finish times of the packets?” This complexity has long been considered as $\Omega(n)$ per packet [11], [12], [3]. For this reason, WFQ and WF^2Q are considered expensive and impractical when the number of sessions are large, and computationally less expensive algorithms such as [11], [12], [3] are proposed to approximate the GPS clock. Therefore, it is very interesting to find out whether $\Omega(n)$ per packet is indeed the complexity lower bound for tracking GPS clock, which decides whether such an unfortunate tradeoff between accuracy and complexity is justified.

This, in turn, leads to another interesting question. Note that, in WFQ and WF^2Q , one only needs to establish the relative orders among GPS virtual finish times. So we wonder whether we can establish this order relationship without explicitly computing their exact values, and hopefully the complexity can be smaller this way. Unfortunately, we prove that such a “shortcut” does not exist: establishing such order relationship is at least as asymptotically expensive as computing their exact values.

Then we established the complexity lower bound of $\Omega(\log n)$ per packet for establishing the relative order of the GPS virtual finish times, under the linear decision tree model (described later). It seems that this lower bound is not tight since we were suspecting that it is $\Omega(n)$ per packet. However, we discovered an existing algorithm [8] that matches the lower bound of $\Omega(\log n)$ per packet. The algorithm is seldom mentioned later in literature partly because it is an integrated part of a theoretical exploration of

a related yet different topic. Our discovery has an immediate practical implication: WFQ and WF^2Q implementations based on this algorithm are optimal among all implementations of WFQ and WF^2Q service disciplines.

Unfortunately, we discovered an important subtlety in this $\Omega(\log n)$ algorithm, which has never been discussed in [8] or any other open literature. A computation timing constraint we refer to as “mandatory lazy evaluation” has to be satisfied for the algorithm to have $\Omega(\log n)$ complexity. We show that theory and practice clash on this constraint: in theory the complexity of this algorithm is strictly $\Omega(\log n)$ per packet, but in practice the worst case complexity could be as high as $\Omega(n)$ per packet due to this “mandatory laziness”. This discovery completely clarifies and answers the question “ $\Omega(\log n)$ or $\Omega(n)$?” (per packet) concerning the complexity of tracking the GPS clock.

Finally, we establish the most important result of this paper: the complexity lower bound of tracking the GPS clock is $\Omega(\log n)$ per packet. This result is shown by establishing an interesting trade-off between the number of comparison operations and the number of other computations (e.g., additions). We show that, for tracking the GPS clock, one has to pay either $\Omega(\log n)$ comparisons or $\Omega(\log n)$ other computations, per packet. Therefore, the overall complexity of the algorithm, which is the sum of both, has to be $\Omega(\log n)$ per packet. This work, combined with the aforementioned $\Omega(\log n)$ algorithm [8], completely settles the long-standing open problem concerning the computational complexity of tracking the GPS clock, in theory¹.

B. On the Minimum Complexity of Achieving $O(1)$ Delay Bounds

The original goal of the second part of our work is to extend the the prior work by Xu and Lipton [14] to a stronger computational model: *linear decision tree*. Among other things, they show that the complexity lower bound is $\Omega(\log n)$ per packet for a scheduling algorithm to guaranteeing $O(1)$ or $O(r^a)$ ($0 < a < 1$) GPS-relative delay bounds. Their result is established under a much weaker computational

model than linear decision tree. We show that the linear decision tree is more practical since the GPS virtual start/finish times in general are *linear functions* of the inputs, and many algorithms require the comparisons among such values².

Instead, we obtained an interesting negative result (prior result in [14] viewed as positive): the methodology (reduce to sorting) used in [14] in proving lower bounds will not work for the stronger model. We show that given a packet arrival scenario used in [14], there is an $O(\log \log n)$ per packet algorithm to guarantee $O(1)$ GPS-relative delay bound under the linear decision tree. This demonstrates that, under such scenarios, scheduling packets to guarantee tight delay bounds is much easier than sorting the GPS virtual finish times. Therefore, it is impossible to reduce the problem of scheduling to sorting. This negative result indicates that a new methodology is needed in order to possibly establish this complexity lower bound, which remains an open question.

The rest of the paper is organized as follows. In Section II, we introduce the computation models and assumptions we will use in proving our results. The aforementioned two major research themes are established in Section III and IV respectively. Section V concludes the paper.

II. RELATED WORK AND BACKGROUND

In this section, we introduce the background knowledge and related work on packet scheduling and computational complexity. Their relevance to the objective of this work will be explicitly stated.

A. Background on GPS and packet-based algorithms

We have introduced GPS, an idealized service discipline that other service disciplines can be compared with in terms of QoS guarantees. One important property of GPS, proved in [9], is that it can guarantee tight end-to-end delay bounds to traffic that is leaky-bucket [13] constrained. It has been shown that scheduling algorithms can provide similar end-to-end delay bounds if their packet service does not lag behind that of GPS too much [9]. This

¹In practice, we will always be “haunted” by the aforementioned subtlety of “mandatory lazy evaluation”.

²Note that [14] also contains results that is based on the linear decision tree. However, they only succeed for results concerning a weaker type of delay, referred to as *disadvantage delay*.

motivates the definition of GPS-relative delay (introduced in [9] and [8] but formally defined in [14]) and the study of existing service disciplines on this property.

For each packet p , the *GPS-relative delay* of a packet p under a particular scheduler ALG is defined as $\max(0, F_p^{ALG} - F_p^{GPS})$, where F_p^{ALG} and F_p^{GPS} are the times when the packet p finishes service in the ALG scheduler and in the GPS scheduler, respectively. It has been shown in [9] and [1] respectively that WFQ and WF^2Q schedulers both have a worst-case GPS-relative delay of $\frac{L_{max}}{r}$, where L_{max} is the maximum packet size and r is the total link bandwidth. That is, for each packet p ,

$$F_p^{WFQ} - F_p^{GPS} \leq \frac{L_{max}}{r} \quad (1)$$

$$F_p^{WF^2Q} - F_p^{GPS} \leq \frac{L_{max}}{r} \quad (2)$$

Using the convention of [14], we say that the delay bound is $O(1)$ since L_{max} and r can be viewed as constants independent of the number of sessions n . WFQ and WF^2Q achieve this $O(1)$ GPS-relative delay bound by (a) keeping perfect track of the GPS clock and (b) picking among all (in WFQ) or all eligible (in WF^2Q) head-of-session (HOS) packets, the one with smallest GPS virtual finish time to serve next. Here HOS packets are the first packets from each session that are currently in queue. This paper essentially studies whether steps (a) and (b) are necessary and if yes, how to carry out both steps in an optimal way.

B. On tracking the GPS clock

The GPS virtual time $V(t)$, as a function of real time t is calculated as follows.

$$V(0) = 0 \quad (3)$$

$$V(t_{j-1} + \tau) = V(t_{j-1}) + \frac{\tau}{\sum_{i \in B_j(t)} \phi_i} \quad (4)$$

$$\tau \leq t_j - t_{j-1}, j = 2, 3, \dots$$

Here $\{t_j\}_{j=1,2,\dots}$ are the times at which two types of events happen: (1) the service starts for a new packet, and (2) the service finishes for a packet currently in queue. Let A_i^k be the real time that the k^{th} packet of the i^{th} session arrives and L_i^k be its length.

Let S_i^k and F_i^k be the virtual times when it should have started and finished service under GPS, respectively. Then one can show that S and F are calculated as follows.

$$S_i^k = \max\{F_i^{k-1}, V(A_i^k)\} \quad (5)$$

$$F_i^k = S_i^k + \frac{L_i^k}{\phi_i} \quad (6)$$

Note that the GPS virtual finish time of a packet is calculated as soon as it arrives as above. However, the real time corresponding to this virtual finish time is not determined until later on (explained later). In service disciplines, the tracking of the GPS clock is to find the GPS *virtual* and *real* finish times of each packet. Although GPS virtual finish times are all that are needed in WFQ for scheduling, the GPS real finish times are equally important since they determine the GPS virtual finish times of the future packets, as shown in formula (3) and (4). In this paper, we solve the open problem concerning the complexity lower bound for computing the GPS virtual and real finish times.

C. On Computational Complexity

Computation complexity of a computational problem is defined under a computational model that specifies what operations are allowed in solving the problem and how they are “charged” in terms of complexity. Same as in [14], we adopt a standard and commonly-used computational model in proving lower bounds: the *decision tree*. A decision tree in general takes as input a list of real variables $\{x_i\}_{1 \leq i \leq n}$. Each internal and external (leaf) node of the tree is labeled with a predicate of these inputs. The algorithm starts execution at the root node. In general, when control is centered at any internal node, the predicate labeling that node is evaluated, and the program control is passed to its left or right child when the value is “yes” or “no” respectively. Before the control is switched over, the program is allowed to execute *unlimited number* of sequential operations such as data movements and arithmetic operations. In particular, the program is allowed to store all results (i.e., no constraint on storage space) from prior computations. When program control reaches a leaf node the predicate there is evaluated

and its result is considered as the output of the program. The complexity of such an algorithm is defined as the depth of the tree, which is simply the number of predicates that needs to be evaluated in the *worst case*. Fig. 1 shows a simple decision tree with four nodes. Each P_i ($1 \leq i \leq 4$) is a predicate of the inputs.




Fig. 1. Decision tree computational model

Allowing different types of predicates to be used in the decision tree results in models of different computational powers. *Throughout this paper, we consider the decision tree that allows the comparisons between linear functions of inputs, known as linear decision tree (LDT)* [5]. In this model, given inputs $\{x_i\}_{1 \leq i \leq n}$, each predicate allowed by the decision tree is in the form of “ $h(x_1, x_2, \dots, x_n) \geq 0$?”, where h is a linear function (defined below) of the inputs $\{x_i\}_{1 \leq i \leq n}$.

Definition 1—Linear Function: A linear function f of the variables $\{x_i\}_{1 \leq i \leq n}$ is defined as $f(x_1, x_2, \dots, x_n) = \sum_{i=1}^n a_i x_i + a_0$, where $\{a_i\}_{0 \leq i \leq n}$ are real constants.

In the context of this work, the inputs will be the lengths and the arrival times of the packets. Note that the linear decision tree model is quite generous: functions like f in the above definition may take up to $O(n)$ steps to compute, but the model charges only “1” for each of them. Due to this generosity, for certain computation problems, we may not be able to obtain a tight lower bound using this model. We stick to this model, however, since it provides mathematical tools (convexity arguments) for solving lower bounds. In addition we will show that this model captures all computation powers of a router, both for tracking the GPS clock and for scheduling packets.

D. Related work on the complexity lower bound of packet scheduling

The first work in studying the tradeoff between delay bound and computational complexity is due to Xu and Lipton [14]. In [14], they have established that the complexity lower bound for a service discipline to provide $O(1)$ or $O(n^a)$ ($0 < a < 1$) GPS-relative delay bound is $\Omega(\log n)$ per packet. This bound is established under a weaker decision tree model that only allows direct comparisons between its inputs. In other words, its comparisons are all in the forms of “ $x_i - x_j \geq 0$?”. This model is strictly weaker than LDT since $x_i - x_j$ is clearly a linear function. However, for the particular class of instances that are used in establishing the lower bounds in [14], the weaker computational model is reasonable in the sense that existing scheduling algorithms are able to provide $O(1)$ GPS-delay bounds using only such predicates.

We study whether the result can be extended to a much stronger and more practical computational model: linear decision tree. We show that this extension is practically important since GPS virtual finish times are all linear functions of the inputs. While trying to establish the lower bound of $\Omega(\log n)$, we instead find a “negative result”. We show that there is no way to use the method of “reducing to sorting” suggested in [14], since under the scenarios used in [14], the scheduling complexity becomes $\Omega(\log \log n)$ while sorting remains $\Omega(\log n)$, per packet.

Whether the complexity lower bound of tracking the GPS clock is $\Omega(n)$ per packet was posed as an open problem in Sigcomm 2002 by Xu (one of the authors in [14]). Since their focus is on the computational complexity of scheduling packets, they elegantly avoid answering this question by introducing the concept of CBFS (Continuously Backlogged Fair Sharing) condition³. They showed that under the CBFS condition, the GPS finish times can be computed easily in $O(1)$ and becomes a nonissue in studying the complexity of scheduling.

³A less elegant way to avoid dealing with this question is to use the relativization arguments in theoretical computer science [10].

E. Related theorems in prior work

In the proofs of this paper, we will use a lemma from [14], which says the complexity lower bound for determining whether a given input vector (x_1, x_2, \dots, x_n) belongs to the following set \mathcal{S} is $n \log_2 n - o(n \log_2 n)$, under the linear decision tree. Here $\mathcal{S} = \{(y_1, y_2, \dots, y_n) \in R^n : \text{there exists a permutation } \pi \text{ of } 1, \dots, n \text{ such that } \frac{iL_{max}}{n+1} - \delta < y_{\pi(i)} < \frac{iL_{max}}{n+1} + \delta, i = 1, 2, \dots, n\}$. Here $L_{max} > 0$ is the maximum packet size and δ ($0 < \delta < \frac{L_{max}}{3(n+1)}$) is a “small” real constant. Here is the lemma from [14].

Lemma 1: Under the linear decision tree model, given the inputs $\{x_i\}_{1 \leq i \leq n}$, determining whether $(x_1, x_2, \dots, x_n) \in \mathcal{S}$ requires at least $n \log_2 n - o(n \log_2 n)$ linear comparisons.

III. ON THE COMPLEXITY OF TRACKING GPS CLOCK

The focus of this section is to clarify and establish the fundamental algorithmic complexity of tracking the GPS clock perfectly in packet scheduling. Recall that both WFQ and WF^2Q service disciplines track GPS clock perfectly, and both are able to achieve a tight GPS-relative delay bound of $O(1)$. All other service disciplines, on the other hand, only estimate GPS clock (with certain error), and it was shown in [14] that none of them can provide the same tight GPS-relative delay bound. Intuitively there is a causal relationship between the perfect tracking of GPS clock and the tightness of the delay bound.

A. GPS clock tracking and its complexity question

The computational complexity lower bound of tracking the GPS clock has long been considered $\Omega(n)$ per packet [11], [12], [3]. For this reason, both WFQ and WF^2Q are considered impractical, and several algorithms that approximate GPS clock at much lower complexity, such as [11], [6], [2], [7], are proposed. However, as mentioned above, none of these algorithms can achieve a tight GPS-relative delay bound of $O(1)$.

We ask the following fundamental question: is $\Omega(n)$ per packet the complexity lower bound for tracking the GPS clock perfectly? Settling this question is important in two aspects.

- First, if this complexity is indeed $\Omega(n)$ per packet, then the complexity lower bound of $\Omega(\log n)$ per packet established in [14] for scheduling the GPS timestamps to achieve tight GPS-relative delay bounds becomes less relevant, since the combined complexity of scheduling is $\Omega(n)$ per packet anyway. Note that [14] elegantly avoids answering this question, as their results are established under a CBFS condition (shown later) in which GPS tracking cost becomes $O(1)$ per packet. This complexity question leads to two other interesting questions. Is GPS clock tracking a necessary step for implementing WFQ and WF^2Q (i.e., Is there a “shortcut”)? Or is this step necessary for providing $O(1)$ GPS-relative delay? We will answer the first question in full. The second question remains an open problem.
- Second, if there is an efficient algorithm of complexity $O(\log n)$ per packet, for tracking GPS time, then WFQ and WF^2Q may become practical since their overall cost will be $O(\log n)$ per packet!

Our findings turn out to be a mixed blessing: in theory (i.e., the strict computational complexity sense) the complexity lower bound is $\Omega(\log n)$ per packet and there is an algorithm that matches the bound. However, in practice it can be as expensive as $\Omega(n)$ per packet in the worst-case due to a “mandatory lazy evaluation” subtlety associated with the algorithm. Our contribution is to discover and clarify this subtlety.

The official reason for tracking the GPS clock to be considered $\Omega(n)$ in WFQ and WF^2Q is that there can be n near-simultaneous arrivals and departures in a tiny interval. However, we found out that this is not the real issue since the algorithms in [11], [6], [2], [7] suffer from at least the problem of near-simultaneous arrivals as well. The real issue is the aforementioned subtlety.

B. Is there a “shortcut”?

Note that in WFQ and WF^2Q , we only care about the relative order of GPS finish virtual/real times rather than their exact values. This leads us to ask the following question: can the complexity of establishing such order relationships be smaller than

```

1. class Packet {
2.     double rt_arrival, rt_finish, vt_start, vt_finish;
3.     /* rt_arrival and rt_finish are the real times of the packet's arrival and departure
4.     (under GPS) respectively. vt_start and vt_finish are its GPS virtual start and
5.     finish times respectively. Initially, only rt_arrival field is filled by the value of
6.     arrival time of the packet and the other fields are initialized to 0. */
7.     int id, len; /* the session id of the packet and the packet length */
8. };

9. class Packet_Queue {
10.    Packet peek(); /* return the first packet in queue without deleting it */
11.    Packet pop(); /* return the first packet in queue and delete it */
12.    int empty(); /* return 1 if queue is empty, otherwise return 0 */
13. };

14. external Packet_Queue P, Q; /* explained in the text and Fig. 3 */

15. total_weight = 0; /* the total weight of all the backlogged sessions between two events,
16. i.e.,  $\sum_{i \in B_j(t)} \phi_i$  */
17. for (i=0; i < n; i++) backlog[i] = 0;
18. /* backlog[i] indicates the amount of backlog at session i queue (in GPS virtual time). */

19. while (!Q.empty()) {
20.     q = Q.peek(); p = P.peek();
21.     if (total_weight == 0) {
22.         /* the start of a new busy period */
23.         total_weight = weight[p.id]; /* weight[i] is  $\phi_i$  */
24.         vt_last = 0; /* the virtual time when the last event happens. */
25.         rt_last = p.rt_arrival; /* the real time of the last event */
26.         continue;
27.     }
28.     vt_temp = vt_last + (p.rt_arrival - rt_last) / total_weight;
29.     if (q.vt_finish > vt_temp) {
30.         /* p starts earlier than q finishes */
31.         p.vt_finish = max(vt_temp, backlog[p.id]) + p.len / weight[p.id];
32.         vt_last = vt_temp; rt_last = p.rt_arrival;
33.         if (backlog[p.id] < vt_temp)
34.             total_weight = total_weight + weight[p.id];
35.         backlog[p.id] = p.vt_finish;
36.         P.pop();
37.     }
38.     else {
39.         /* q finishes first before p starts */
40.         q.rt_finish = rt_last + (q.vt_finish - vt_last) * total_weight;
41.         rt_last = q.rt_finish; vt_last = q.vt_finish;
42.         if (backlog[q.id] ≤ vt_last)
43.             total_weight = total_weight - weight[q.id];
44.         Q.pop();
45.     }
46. }

```

Fig. 2. Algorithm *B* for GPS time tracking.

that of computing the GPS virtual/real finish times? In other words, is there a “shortcut” that allows us to establish such orders without explicitly computing their values? Such shortcuts may be possible in some computations. For a trivial example, to know whether $x + z < y + z$ holds, we may just compare x with y without adding up $x + z$ and $y + z$. Real-world examples can be much more sophisti-

cated since comparisons can be nested deeply in a expression.

Our discovery is surprising: establishing the relative orders among the GPS finish times is at least asymptotically as expensive as computing them, proved in the following theorem.

Theorem 1: The complexity lower bound of sorting the GPS virtual/real finish times is asymptotically at least as large as that of computing their exact




Fig. 3. Queue structure of Algorithm B

values.

Proof: Our proof proceeds as follows. Assume that we know the relative order in which the packets finish service under GPS, which is computed using a *magic algorithm* A with complexity $C(n)$ per packet. We show there is an algorithm B that uses such order to produce the correct GPS virtual/real finish times of the packets, with complexity $O(1)$ per packet. We denote as $D(n)$ the complexity lower bound of computing the GPS finish times per packet. Clearly $D(n) \leq C(n) + O(1)$. However, since $D(n)$ is at least $O(1)$ (just to read the input), this shows that $C(n)$ is at least *asymptotically* as large as $D(n)$.

The reduction algorithm B is shown in Fig. 2. The queue Q contains pointers to these packets sorted in the increasing order of their GPS finish times. This captures our assumption that we know their relative orders. The incoming packet FIFO queue P contains pointers to these packets in the sorted order of their arrival times. Note that pointers in both P and Q refer to the same pool of actual packets, as shown in Fig. 3. For each packet, initially only the real arrival time field (`rt_arrival`) is set. After the execution of the program, we show that its GPS virtual finish time (`vt_finish`) and its real finish time (`rt_finish`) will also be set. This way the mission of computing the GPS real/virtual finish times of the packets is accomplished.

We used a standard technique in programming language to make the algorithm B succinct. We assume that at the end of queue P , there is a pointer to a “dummy” packet with arrival time $+\infty$. Q , on the other hand, does not contain such a pointer. This “dummy” pointer will “flush” (line 41) all packets in Q before itself gets popped. This “flushing” considerably simplifies the inspections of boundary conditions.

We need to prove two things: (a) the complexity of

the algorithm is $O(1)$ per packet, and (b) it correctly computes the GPS virtual and real finish times of the packets.

Proof of (a): Note that during each “while” iteration except for the first, which goes through line 19 to 46, there are at least one element popped from either P (line 36) or Q (line 44). Eventually both will become empty and the program terminates. Each iteration clearly takes $O(1)$ time.

Proof of (b): We need to show that the algorithm computes the GPS real/virtual finish times as defined in formula (3) and (4). We prove this by induction on the number of events. The GPS virtual time when the first event (the first packet arrival) happens is 0 (line 24). Suppose the values of the GPS virtual/real times and other state variables corresponding to the k^{th} event is computed correctly. We need to show that the values corresponding to the $(k+1)^{th}$ event are correctly computed. The $(k+1)^{th}$ event can be one of the following two cases: (i) the finish of a packet from a session i , (ii) the start of a new packet from a previously unbacklogged session.

- In case (i), p arrives after q gets popped, since otherwise its arrival would have been the $(k+1)^{th}$ event. So line 38 through 45 will be executed. Depending on whether the session $q.id$ queue becomes unbacklogged (line 42) or not after the departure of q , the total weight (`total_weight`) is either decreased or unchanged (line 43). The real finish time of q can now be finalized (line 40) since there are no new arrivals before its departure.
- In case (ii), p arrives earlier than q 's departure, and therefore line 30 through 37 will be executed to update the total weight and compute p 's GPS virtual start and finish times.

In both cases, the advance of the GPS clock $V(t)$ from the time of the k^{th} event to that of the $(k+1)^{th}$ event conforms to the formulas (3) and (4). Therefore, the algorithm should compute the GPS virtual/real times and other state variables correctly for the $(k+1)^{th}$ event and by induction all events. ■

This result is important since it shows that there is no shortcut to computing the GPS virtual finish times for implementing WFQ and WF^2Q .

C. A tight lower bound

Next, we prove the lower bound of $\Omega(\log n)$ per packet for sorting GPS virtual finish times (Theorem 2) under the linear decision tree model. To prove this, we need a lemma showing that the complexity of sorting n numbers is $\Omega(n \log n)$ under the linear decision tree (Lemma 2). Note that the traditional sorting complexity we know only allows comparisons among inputs.

Lemma 2: The complexity lower bound for sorting n distinct positive numbers $x_i, i = 1, 2, \dots, n$, is $\Omega(n \log n)$ under the linear decision tree model.

Proof: Suppose sorting these numbers is not $\Omega(n \log n)$. Then we construct an algorithm A for deciding the membership for the aforementioned set \mathcal{S} , with a complexity that is not $\Omega(n \log n)$. This would contradict Lemma 1. Algorithm A first sorts these numbers using linear comparisons. Then it verifies if the sorted list is in \mathcal{S} , and this step clearly only needs $O(n)$ time (check if $\frac{iL_{max}}{n+1} - \delta < y_i < \frac{iL_{max}}{n+1} + \delta, i = 1, 2, \dots, n$). Since $O(n) = o(n \log n)$, the complexity of algorithm A is not $\Omega(n \log n)$ if the sorting part is not $\Omega(n \log n)$. ■

Theorem 2: The complexity lower bound for computing the relative order of the GPS finish times among n packets is $\Omega(n \log n)$ under the linear decision tree model.

Proof: We reduce the problem to sorting. Given a sorting instance $x_i > 0, i = 1, 2, \dots, n$ and $x_i \neq x_j$ when $i \neq j$, we convert it to a packet arrival instance as follows. Let there be n sessions sharing a link of rate 1 with equal weight. A packet of length x_i arrives at time 0 to session $i, i = 1, 2, \dots, n$. Note that the relative order of the GPS finish times is exactly the order of x_i . Then from Lemma 2 we know that to sort the lengths of n packets requires the cost of $\Omega(n \log n)$. So it is $\Omega(\log n)$ per packet. ■

Note that the lower bound is not tight unless we can find an algorithm that matches it. Fortunately and surprisingly, we found an existing algorithm that tracks GPS clock using only $\Omega(\log n)$ per packet complexity [8]. However, it was seldom mentioned or referenced in later literature, including the algorithms that try to reduce the complexity of GPS tracking by using approximation. Part of the reason is that [8] is a theoretical endeavor on a related but different topic.

The algorithm in [8] is involved and is omitted here due to lack of space. Its key idea, however, can be explained relatively easily based on our algorithm in Fig. 2. Note that in Fig. 2, once we have Q , we can figure out the GPS virtual finish times of the packets with complexity $O(1)$ per packet. The algorithm in [8] essentially keeps a priority queue of the HOS packets to obtain Q , which is $\Omega(\log n)$ per packet. Each session also maintains a per-session FIFO queue. When a packet in an HOS queue should have finished service in GPS, its next packet (if it exists) is moved from its corresponding per session FIFO queue to the HOS priority queue.

The following result shows that if the WFQ is performed by a priority queue and the computation of GPS clock is performed using the algorithm in [8], then it is the asymptotically optimal implementation of the WFQ service discipline.

Corollary 1: Algorithm in [8] for tracking GPS clock combined with a priority queue to sort GPS virtual finish times is the asymptotically optimal algorithm for implementing WFQ scheduling discipline, under the linear decision tree model.

Proof: Combine the algorithm in Fig. 2, and Theorem 2 on the lower bound of comparing GPS virtual finish times. ■

D. Mixed blessing

Now that we have an $O(\log n)$ per packet algorithm for WFQ , do we still need to approximate the GPS clock as in [11], [6], [2], [7]? The answer is yes. We found that the new $\Omega(\log n)$ algorithm is a mixed blessing in the following sense. It requires a subtlety we refer to as “mandatory lazy evaluation” for its theoretical complexity to be $\Omega(\log n)$ per packet. In practice, this mandatory laziness may lead to $\Omega(n)$ worst-case complexity. To avoid this worst case, certain “approximation measures” need to be taken anyway. Our contribution here is to discover and clarify this subtlety.

We discover that the algorithm in [8] implicitly requires a computation timing we refer to as “mandatory lazy evaluation”. Suppose a packet p 's real finish time is t , provided that there are no other packet arrivals before its departure. Then for an algorithm to be efficient, the computation of this time cannot be finalized until t has passed. The reason is that if this computation is performed at time $t' < t$, and

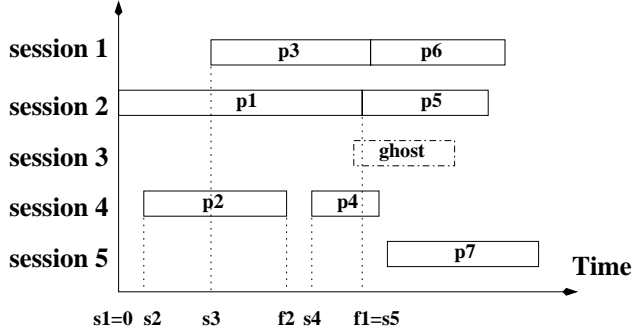


Fig. 4. Why lazy evaluation is necessary?

if there is a packet arrival between t' and t , the GPS real finish time of p would have to be recomputed. In practice, this “lazy evaluation” can be implemented as a hardware timer that triggers at time t the computation of its GPS real finish time. However, in the worst case, up to $\Omega(n)$ GPS virtual finish times may be “clustered together” in a small time interval, and a new arrival immediately after the “cluster” can not be processed until all $O(n)$ departures in that cluster get processed. The GPS real finish times of all these departure events need to be computed during this interval due to lazy evaluation. Therefore, the processing and service of the new arrival triggers the computation of $\Omega(n)$ GPS real finish times. This leads to the $\Omega(n)$ worst-case complexity.

We illustrate this “lazy evaluation problem” by a concrete example, shown in Fig. 4. The horizontal axis is the GPS timeline. We denote these packets as p_i and their GPS virtual finish and start times as f_i and s_i respectively. In addition, we define the function $R(t)$ that returns the real time corresponding to virtual time t . In Fig. 4, we show that $R(f_1)$, $R(f_3)$ and $R(f_4)$ are very close to each other and form a “cluster”, followed immediately by the arrival of p_7 . To be able to compute s_7 , we need to finish computing the GPS real finish times corresponding to f_1 , f_3 , and f_4 . However, these computations need to be finished “within” the tiny interval of $[R(f_1), R(s_7)]$ due to “mandatory lazy evaluation”, which may not be possible in real-world finite-speed CPU’s. On the other hand, if we compute these real times well in advance, then if the “ghost” packet in the figure does arrive, all these values have to be recomputed.

This is the classical irony between theory and practice. From the theoretical point of view it is

$\Omega(\log n)$ per packet, since we only count the number of operations. However, in practice, CPU has finite speed and a “cluster” of $O(n)$ packets may need to be processed during a short time interval. We are currently investigating how to alleviate this problem by using approximation, e.g., assuming that all the real finish times are the same as that of the last packet in the “interval”. It is not clear whether such approximations will be better than existing approximation schemes such as used in [11], [6], [2], [7]. This can be an interesting topic for further study.

E. Tracking GPS is $\Omega(\log n)$ per packet

In this section, we answer the most important and challenging question of this paper: “What is the complexity lower bound for computing the GPS virtual finish times of the packets?” We show that this lower bound is $\Omega(\log n)$ per packet.

Our lower bound is established based on the following set of instances that we will refer to later as $I(n)$. Suppose n sessions share a link of rate 1. There are n arrivals p_1, p_2, \dots, p_n at time 0 from sessions 1, 2, ..., n respectively. The weights of these sessions are w_1, w_2, \dots, w_n respectively. Also we assume that l_i/w_i are all distinct. We refer to the set of all such instances of size n mentioned above as $I(n)$. Note also that these weights w_i s are viewed as constants rather than as inputs.

We would like to compute their GPS real finish times. We let π be the unique permutation such that $l_{\pi(1)}/w_{\pi(1)} < l_{\pi(2)}/w_{\pi(2)} < \dots < l_{\pi(n)}/w_{\pi(n)}$. We can compute the GPS real finish times as the follows:

$$S(p_{\pi(1)}) = l_{\pi(1)} * \frac{\sum_{j=1}^n w_{\pi(j)}}{w_{\pi(1)}}$$

$$S(p_{\pi(i)}) = \frac{l_{\pi(1)}}{w_{\pi(1)}} * \sum_{j=1}^n w_{\pi(j)} + \sum_{m=2}^i \left[\left(\frac{l_{\pi(m)}}{w_{\pi(m)}} - \frac{l_{\pi(m-1)}}{w_{\pi(m-1)}} \right) \sum_{j=m}^n w_{\pi(j)} \right], i = 2, 3, \dots, n$$

Let $x_1 = l_{\pi(1)} * \frac{\sum_{j=1}^n w_{\pi(j)}}{w_{\pi(1)}}$, $x_i = \left(\frac{l_{\pi(i)}}{w_{\pi(i)}} - \frac{l_{\pi(i-1)}}{w_{\pi(i-1)}} \right) \sum_{j=i}^n w_{\pi(j)}$ for $i = 2, 3, \dots, n$, and $S_i = S(p_{\pi(i)})$. We need to show that to compute $\{S_i\}_{i=1, \dots, n}$ requires $\Omega(n \log n)$ computation.

Obviously, if we know the right permutation π , we can compute $\{S_i\}_{1 \leq i \leq n}$ in $O(n)$, since $S_{i+1} =$




Fig. 5. Example GPS computation tree

$S_i + x_{i+1}$ and x_i can be computed in $O(1)$. But computing this π requires $\Omega(n \log n)$ comparisons, as shown in Lemma 2. However, this is not a proof for the lower bound since it may be possible to compute S_i without computing π completely (perhaps only computing some partial orders captured in π). To be able to establish the lower bound of S_i 's, we need a computation model that captures *all possible ways to compute S_i and their costs*.

Our model is the following. We assume that each term x_i can be computed with $O(1)$ computation cost and each addition costs 1. Note we does not gain (rather lose) from this assumption since we are trying to establish the lower bound. If S_j is computed from S_i ($j > i$), i.e., S_j is computed as $S_i + x_{i+1} + x_{i+2} + \dots + x_j$, the amount of computation is counted as $j - i$. This models the savings we get by using prior computation results. Also, for the computation S_j to use the result of computation S_i , we “must have” learned about the partial order $S_i < S_j$. This model is very natural since it captures all conceivable ways of computing S_j . It also captures the tradeoff between comparison and computation: either one pays for “learning” about the partial orders between S_i 's and S_j 's in the form of comparisons or pays for the additions in the forms of $S_i + x_{i+1} + x_{i+2} + \dots + x_j$.

In the following, we show that no matter how the computation is carried out, either we have to pay $\Omega(n \log n)$ cost in comparison or $\Omega(n \log n)$ in computation, for the worst-case instance in $I(n)$. In other words, the total complexity (computation + comparison) has to be at least $\Omega(n \log n)$. Since there are n packets, the complexity per packet is $\Omega(\log n)$.

For the simplicity of stating and proving our result, we use a directed graph to represent the above computation model. Any instance (of computing

```

1. FIFO_PacketQueue  $Q \leftarrow \emptyset$ ;
2. Sorted_Array  $P \leftarrow \{0\}$ ;
3. while (! $P.empty()$ ) {
4.    $p \leftarrow P.min()$ ;
5.   delete  $p$  from  $P$ ;
6.   if  $p$  has children
7.     Sort  $p$ 's children and merge into  $P$ ;
8.   Append  $p$  to  $Q$ ;
9. }
```

Fig. 6. Algorithm B for getting the sorted sequence from the GPS computation tree

GPS virtual/real finish times) in $I(n)$ above corresponds to a graph containing $n + 1$ points, $0, 1, \dots, n$. Here points $1, 2, \dots, n$ correspond to the computation job of S_1, S_2, \dots, S_n respectively and the point 0 corresponds to a “null” computation job. There is a directed edge from i to j if and only if S_j is computed from S_i . The cost of this edge is $j - i$, which captures the number of additions for computing S_j from the intermediate result S_i , as explained above. A directed edge from 0 to i means that S_i is computed “from scratch” as $x_1 + x_2 + \dots + x_i$, and its cost is i . Note that, ignoring the direction, this directed graph is a tree because each computation job has a unique computation path from the point 0. In each such tree, there must be a directed path from 0 to any other point, since every S_i needs to be computed. We refer to this tree as the *GPS computation tree* corresponding to the instance. An example tree is shown in Fig. 5. In Fig. 5, S_3 and S_4 are computed from S_1 and S_2 respectively, and S_1 and S_2 are computed “from scratch”. The numbers associated with the links are the computation costs.

We denote as $GPS(n)$ the worst case complexity for computing S_i 's, among all instances in $I(n)$. Note that for each instance in $I(n)$, its computational complexity includes both additions and comparisons. With this graph model, our result can be formulated as the following theorem.

Theorem 3: $GPS(n) = \Omega(n \log n)$

Proof: We prove by contradiction. Suppose $GPS(n)$ is not $\Omega(n \log n)$. In other words, for any $N > 0$ and $\epsilon > 0$, there exists $n' > N$ such that, the worst case complexity (comparisons plus computations) for GPS computation trees of size n' is no more than $\epsilon n' \log n'$. We denote this as assumption (I).

We denote as $SORT(n)$ the worst-case complex-

ity of sorting the GPS virtual finish times for the aforementioned instance set $I(n)$. We know from Theorem 2 that $SORT(n)$ is $n \log_2 n - o(n \log_2 n)$, under the linear decision tree model. So there exists N_1 such that for all $n > N_1$, $SORT(n) > \frac{5}{6}n \log_2 n$. We also choose N_2 such that when $n > N_2$, the value of the function $2n$ is no more than $\frac{1}{6}n \log_2 n$. We let $N_3 = \max(N_1, N_2)$. We know from assumption (I) that there exists $n' > N_3$ such that $GPS(n') \leq \frac{1}{6}n' \log_2 n'$, denoted as fact (II). Note that since $n' > N_3 > N_1$, $SORT(n') > \frac{5}{6}n' \log_2 n'$. This is denoted as fact (III).

We show that we can construct an algorithm A that, given any instance of size n' as above, sorts the GPS virtual finish times of these packets using no more than $\frac{5}{6}n' \log n'$ comparisons. This clearly contradicts the fact (III) above. The algorithm A runs as follows. For each packet arrival instance of size n' , A first tries to compute the virtual finish times of the packets. This results in a GPS computation tree corresponding to this instance. Since $GPS(n') \leq \frac{1}{6}n' \log_2 n'$ (fact (II) above), neither the number of comparisons nor the number of computations in this computation tree can be more than $\frac{1}{6}n' \log_2 n'$. Then A calls another algorithm B , which takes as its input a GPS computation tree, and outputs the GPS finish times in the sorted order. We will show that the complexity of algorithm B is no more than $\frac{4}{6}n' \log_2 n'$ in the worst case. So the complexity of the algorithm A is no more than $\frac{1}{6}n' \log_2 n' + \frac{4}{6}n' \log_2 n' = \frac{5}{6}n' \log_2 n'$, which is exactly what we would like to prove above.

The algorithm B is shown in Fig. 6. It sorts the GPS finish times of these n' packets when it takes as its input a GPS computation tree corresponding to an instance in $I(n')$. In Fig. 6, Q is a FIFO queue of packets and P is an array sorted by S_i and $\min()$ returns the element with the minimum S_i value (the head element actually). If a set of packets S_{j_1}, S_{j_2}, \dots are computed from S_i , then they are all S_i 's children.

It remains to prove two things: (a) the output from the algorithm is indeed sorted, and (b) the complexity of the algorithm is no more than $\frac{4}{6}n' \log_2 n'$.

Proof of (a): Clearly the ‘‘while’’ iterations in Fig. 6 result in a traversal of all the nodes in GPS computation tree since every point is reachable from 0. Therefore, Q will contain a permutation of all the

points S_i after the execution. To prove that they are sorted, we only need to show the following invariant \mathcal{I} at each iteration of the ‘‘while’’ loop: the number S_i is deleted from P in the $(i + 1)^{th}$ iteration for $i = 0, 1, \dots, n$. We prove it by induction on i . In the first iteration, clearly number 0 (ie., S_0) is deleted from P . Suppose \mathcal{I} holds for the number $i < k$. In the $(k + 1)^{th}$ iteration, suppose $S_{k'}$ ($k' \neq k$) will be deleted from P instead of S_k . We know that $k' > k$ due to induction hypothesis. Clearly S_k cannot be in the queue P at the moment since otherwise ‘‘min()’’ will choose S_k instead of $S_{k'}$. Suppose the parent of S_k is $S_{k''}$. Clearly $k'' < k$ and $S_{k''}$ is deleted during the $(k'' + 1)^{th}$ iteration due to the induction hypothesis. So at the $(k'' + 1)^{th}$ iteration, S_k must have been inserted. Therefore, the only possibility is that S_k has been deleted earlier than the $(k + 1)^{th}$ iteration. This contradicts the induction hypothesis.

Proof of (b): Let E be the set of edges in the GPS computation tree and let e_{ij} be an edge from S_i to S_j ($j > i$). By the fact (II) above, $\sum_{e_{ij} \in E} (j - i) \leq \frac{1}{6}n' \log_2 n'$. Note that, in terms of the number of comparisons, the complexity of inserting an element into a sorted array of size N is no more than $\log_2 N$, and the cost of deletion is 1. We assume that right before number S_i is deleted (line 5), there are C_i elements in P . We also assume that S_i has D_i children. The complexity of inserting all these children is $C_i + D_i + D_i \log_2 D_i$, since sorting all the children up takes $D_i \log_2 D_i$ time and merging two sorted lists of size C_i and D_i each, takes time $C_i + D_i$. The total complexity is therefore $\sum_{i=0}^{n'-1} (1 + C_i + D_i + D_i \log_2 D_i)$. Note that the complexity of the last iteration is zero in terms of comparisons since at that time there is only one element left in the priority queue P . So in the following summations, we will not count the last iteration, and will only add up the index $n' - 1$ (instead of n').

Clearly, $\sum_{i=0}^{n'-1} D_i = n'$ since each of the n' nodes $0, 1, \dots, n - 1$ appeared in D_i exactly once by the aforementioned invariant \mathcal{I} . We claim that $(\sum_{i=0}^{n'-1} C_i)$ is no more than $|E| + \sum_{e_{ij} \in E} (j - i)$. This can be shown by the technique of *double counting* in combinatorics as follows. Note that an edge e_{ij} is counted for most $j - i + 1$ times in the above summation since e_{ij} is inserted in round $i + 1$ and deleted in round $j + 1$ (i.e., counted in C_i, C_{i+1}, \dots, C_j) by the aforementioned invariant \mathcal{I} .

So $(\sum_{i=0}^{n'-1} C_i) \leq \sum_{e_{ij} \in E} (j - i + 1) = (\sum_{e_{ij} \in E} (j - i)) + |E|$. So $(\sum_{i=0}^{n'-1} C_i) \leq \frac{1}{6}n' \log_2 n' + n'$ since $|E| = n'$ (property of tree).

Now we claim that $\sum_{i=0}^{n'-1} D_i \log_2 D_i \leq \sum_{i=0}^{n'-1} \frac{1}{2} D_i (D_i + 1) \leq \sum_{e_{ij} \in E} (j - i)$. The first inequality is due to the fact $\log_2 x \leq (x + 1)/2$ when $x > 0$. The second equality is because these D_i edge from node S_i are at least of lengths $1, 2, \dots, D_i$ respectively, and therefore the total length of edges coming out from S_i is at least $\frac{1}{2} D_i (D_i + 1)$. So the total complexity of the algorithm is no more than $n' + (\sum_{i=0}^{n'-1} C_i) + (\sum_{i=0}^{n'-1} D_i) + (\sum_{i=0}^{n'-1} D_i \log_2 D_i) \leq n' + (\frac{1}{6}n' \log_2 n' + n') + n' + (\frac{1}{6}n' \log_2 n') \leq \frac{4}{6}n' \log_2 n'$. ■

IV. THE COMPLEXITY OF SCHEDULING UNDER LINEAR DECISION TREE

In this section, we study the minimum complexity needed for a service discipline to provide tight delay bounds. It was shown in [14] that the complexity needed for scheduling algorithms to guarantee $O(1)$ or $O(n^a)$ GPS-relative delay bounds is $\Omega(\log n)$, under the aforementioned weaker model. Our initial goal was to extend this result to the stronger computation model of linear decision tree. This extension would be important as we will show in Sec. IV. A that the GPS virtual finish times are all linear functions of the inputs. We instead discover an interesting negative result. We show in Sec. IV. B. 1 and IV. B. 2 that methodologies in [14] will not work for the new model. The complexity lower bound for scheduling packets to achieve tight delay bounds under the linear decision tree model remains an open problem.

A. Why extension to the linear decision tree model is important?

In this section, we prove in the following theorem that the GPS virtual finish times are all linear functions of the inputs. Extending the results of [14] to the linear decision tree is therefore important since comparing between these timestamps is an important step in most of the existing scheduling algorithms. In the following proof, packet arrivals during a busy period are denoted as p_1, p_2, \dots, p_m in the sorted order of their GPS virtual start times. Let

$S(p)$ and $F(p)$ denote the GPS virtual start and finish time of the packet p respectively. Let $L(p)$ and $A(p)$ denote its length and real arrival time. As before, $V(t)$ denotes the GPS virtual time corresponding to the real time t . The inputs here are the real arrival times $A(p_i)$ and lengths of the packets $L(p_i)$, $i = 1, 2, \dots, m$.

Theorem 4: The GPS virtual finish times of the packets are all linear functions of the inputs.

Proof: Without loss of generality, we prove the result only for a busy period and assume that its start time is 0. We only need to show that the virtual start times $(\{S(p_i)\}_{1 \leq i \leq m})$ are linear functions, since $F(p_i) = S(p_i) + L(p_i)/\phi_i$ and $L(p_i)/\phi_i$ is a linear function. We prove this by induction on the index i . When $i = 1$, $S(p_1) = 0$ since it is the start of a busy period. Suppose $S(p_i)$ are all linear functions for $i < k$ ($k > 1$). When $i = k$, $S(p_k) = \max\{F(q), V(A(p_k))\}$ where q is the predecessor of p_k in the same session ($F(q)$ is defined as $-\infty$ if such q does not exist). We have two cases:

- Suppose $F(q) \geq V(A(p_k))$. Then since $q = p_j$ for some $j < k$, $S(p_k) = F(p_j)$ is a linear function by the induction hypothesis.
- Suppose $F(q) < V(A(p_k))$. Then $S(p_k) = V(A(p_k))$. We need to show that $V(A(p_k))$ is a linear function. Let e be the last event that happens before the arrival of p_k and t is the real time when it happens. We know that e is either the start event of p_j for some $j < k$ or the finish event of $p_{j'}$ for some $j' < k$. In either case, $V(t)$ is a linear function by the induction hypothesis. Then we know that $S(p_k) = V(t) + \frac{A(p_k) - t}{\sum_{i \in B(t)} \phi_i}$, where $\sum_{i \in B(t)} \phi_i$ is a constant between time t and $A(p_k)$. Therefore, $S(p_k)$ is a linear function. ■

B. A surprising “negative” result

In this section, we prove the “negative” result that the methodology used in [14] for establishing the minimum complexity for service disciplines to provide tight delay bounds will not work for the linear decision tree model. We first introduce in Sec. IV. B. 1 the methodologies used in [14]. Then we show

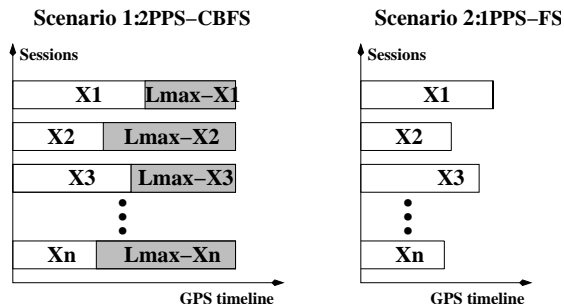


Fig. 7. Scenarios used in prior work

in Sec. IV. B. 2 why it does not work for the new model.

1) *Methodology used in prior work:* We first introduce a packet arrival scenario that is used throughout the proof of [14], shown in the left side of Fig. 7. In this scenario, each session has the equal weight, i.e., $\phi_i = \phi_j$ for $i \neq j$, and the link rate is r . Each session has two packet arrivals, both at time 0 but their relative order is distinguished. The length of the first packet is L_i and the length of the second packet is $L_{max} - L_i$. It has shown in [14] that during the (real) time interval $[0, \frac{nL_{max}}{r}]$, the CBFS condition is satisfied. In the sequel, we refer to this scenario 2PPS-CBFS (2PPS := 2 Packets Per Second). The left side of Fig. 7 shows an example 2PPS-CBFS scenario.

The following methodology of “reducing to sorting” is used in [14] for proving the lower bounds of $\Omega(\log n)$ per packet, under the aforementioned weaker model. The main idea is to show that to sort n numbers satisfying certain property is “almost equivalent” to scheduling n packets for providing tight delay bound. They first show that to sort n numbers l_i , $i = 1, 2, \dots, n$ needs $\Omega(n \log n)$ under the weaker model, even when a permutation of these numbers are around the neighborhoods of $\frac{L_{max}}{n+1}$, $\frac{2L_{max}}{n+1}$, ..., $\frac{nL_{max}}{n+1}$, respectively. The latter condition is equivalent to saying that $(l_1, l_2, \dots, l_n) \in \mathcal{S}$ and the set \mathcal{S} is as defined in Sec. II. E. Then they show that given a 2PPS-CBFS packet arrival instance, in which the lengths of the first packets are $\{l_i\}_{1 \leq i \leq n}$ and $(l_1, l_2, \dots, l_n) \in \mathcal{S}$, a scheduling algorithm that guarantees $O(1)$ GPS-relative delay will make them almost sorted. Here “almost sorted” means only $o(n \log n)$ additional effort is

needed to get the output of the scheduling algorithm sorted. Therefore, the complexity of the scheduling algorithm to sort these $2n$ packets is $\Omega(n \log n)$, or $\Omega(\log n)$ per packet. Lower bounds concerning $O(n^a)$ GPS-relative delay are proved in a similar way.

Each 2PPS-CBFS scenario also has a corresponding 1PPS-FS scenario, in which the second packet of each session is taken away. Fig. 7 shows at the right side the 1PPS-FS scenario corresponding to the 2PPS-CBFS scenario shown at left side. We call it 1PPS-FS, since the condition “continuously backlogged” is no longer satisfied and each session only has one packet arrival. We will also prove a stronger result based on this scenario.

2) *Our result:* We show that the aforementioned method of “reducing to sorting” will not work under the linear decision tree. We show that under linear decision tree, the complexity of sorting the above numbers $\{l_i\}_{1 \leq i \leq n}$ remain the same (Theorem 5 below). However, the complexity of scheduling these packets to provide $O(1)$ delay bound becomes $\Omega(\log \log n)$ per packet, shown in Theorem 6! Therefore, providing tight delay bounds using scheduling is no longer “almost equivalent” to sorting.

We first show that even if $(l_1, l_2, \dots, l_n) \in \mathcal{S}$ as explained above, to sort the set of numbers $\{l_i\}_{1 \leq i \leq n}$ still needs $\Omega(n \log n)$ complexity.

Theorem 5: The sorting is $\Omega(n \log n)$ even if we know that the numbers are in the set \mathcal{S} under the linear decision tree model.

Proof: [Sketch] Suppose there is a sorting algorithm A for these numbers that is not of complexity $\Omega(n \log n)$. In other words, for any $N > 0$, there exists $n' > N$ such that the worst-case complexity to sort these numbers for all instances of size n' is no more than $\frac{1}{4}n' \log_2 n'$ comparisons. We assume that N is large enough so that the complexity lower bound of \mathcal{S} -membership for instances of size n' is larger than $\frac{1}{2}n' \log_2 n'$ (recall from Lemma 1 that this lower bound is $n \log_2 n - o(n \log_2 n)$) and $2n'$ is no larger than $\frac{1}{4}n' \log_2 n'$, when $n' > N$. We construct an \mathcal{S} -membership algorithm B that has the worst-case complexity no more than $\frac{1}{2}n' \log_2 n'$ comparisons for all instances of size n' as follows. Algorithm B first runs algorithm A . Then we check the \mathcal{S} -membership using $2n'$ comparisons. The to-

tal complexity of the algorithm B is no more than $\frac{1}{4}n' \log_2 n' + 2n' \leq \frac{1}{2}n' \log_2 n'$, a contradiction. ■

Now we are ready to present our $O(\log \log n)$ (per packet) algorithm for scheduling packets to provide $O(1)$ GPS-relative delay bound. Like in [14], we assume that the packet arrival instance is a 2PPS-CBFS scenario. We refer to the first packet of session i as p_i and its size l_i , $i = 1, 2, \dots, n$.

Theorem 6: For a 2PPS-CBFS packet arrival instance where the lengths of the first packets $\{l_i\}_{1 \leq i \leq n}$ satisfy the condition $(l_1, l_2, \dots, l_n) \in \mathcal{S}$, there is an $O(\log \log n)$ per packet scheduling algorithm that provides $O(1)$ GPS-relative delay, under the linear decision tree model.

Proof: We assume that $n + 1 = 2^k$ for the simplicity of discussion. We prove that the scheduling algorithm shown in Fig. 8 provides $O(1)$ GPS-relative delay under the linear decision tree model. Note that the complexity of line 3 is $\Omega(\log \log n)$ since the algorithm can perform binary search among $\log_2 n$ intervals $\left[\frac{2^{j-1}-0.5}{n+1}, \frac{2^j-0.5}{n+1} \right)$, $1 \leq j \leq k$ to find the right one and each iteration processes one packet. Therefore, the total complexity is $\Omega(n \log \log n)$, or $\Omega(\log \log n)$ per packet.

We now show that no packet will be delayed for more than L_{max} (assume the service rate is 1) after scheduling, for all inputs $(l_1, l_2, \dots, l_n) \in \mathcal{S}$. Since $(l_1, l_2, \dots, l_n) \in \mathcal{S}$, by the definition of \mathcal{S} , we know that there exists a permutation π such that $\frac{iL_{max}}{n+1} - \delta < l_{\pi(i)} < \frac{iL_{max}}{n+1} + \delta$. Also clearly $l_{\pi(1)} < l_{\pi(2)} < \dots < l_{\pi(n)}$. Given p_i , suppose $\frac{(2^{j-1}-0.5)L_{max}}{n+1} \leq l_i < \frac{(2^j-0.5)L_{max}}{n+1}$, $1 \leq j \leq k$. Then p_i is in Q_j and will be scheduled earlier than all “first” packets $p_{\pi(m)}$ such that $m \geq 2^j$ and all “second” packets. So the service time of the packet is no more than $l_{\pi(1)} + l_{\pi(2)} + \dots + l_{\pi(2^j-1)}$

$$\leq \sum_{m=1}^{2^j-1} \left(\frac{mL_{max}}{n+1} + \delta \right) \leq \frac{(2^j-1)2^{j-1}L_{max}}{n+1} + (2^j-1)\delta$$

$$\leq \frac{2^{j-1}nL_{max}}{n+1} + n\delta. \text{ However, its GPS virtual finish time is exactly } nl_i \geq \frac{(2^{j-1}-0.5)nL_{max}}{n+1}. \text{ So its GPS-}$$

relative delay is at most $n\delta + \frac{nL_{max}}{2(n+1)} < L_{max}$. ■

We can actually prove a stronger result. We show that even if the second packet does not exist, i.e., for 1PPS-FS scenarios, we still have an $\Omega(\log \log n)$ algorithm. This result is stronger since the GPS-relative delay of the same packet under a 2PPS-CBFS scenario is strictly smaller than that under a

corresponding 1PPS-FS scenario. Such an algorithm is shown in Fig. 9. It is more complicated than the algorithm in Fig. 8. We omit the proof of its correctness since it is similar to the proof of Theorem 6. Its complexity is clearly $\Omega(\log \log n)$ per packet for the reason explained in the proof of Theorem 6.

V. CONCLUSIONS

In this work, we answered the following two sets of important open questions concerning the computational complexity of packet scheduling:

- 1) **On the complexity of tracking the GPS clock:** We first answer the interesting open question as to whether the complexity lower bound of “sorting” the GPS virtual finish times, which is all that are needed in WFQ , can be smaller than computing their exact values. We showed, surprisingly, that the former is at least asymptotically as large as the latter. Then we showed that this lower bound complexity for “sorting” the GPS virtual finish times is $\Omega(\log n)$ per packet, and discovered that an existing algorithm [8] matches this bound. This had an important implication: the WFQ implementation that uses the GPS tracking algorithm in [8] is the optimal algorithm for implementing the WFQ service discipline. We discovered, however, that this GPS tracking algorithm requires “mandatory lazy evaluation”, so that although its theoretical complexity is $\Omega(\log n)$ per packet, its worst-case complexity in practice can be as large as $\Omega(n)$. Finally, we proved the most important result of this paper: the complexity lower bound for tracking the GPS clock perfectly is $\Omega(\log n)$ per packet.
- 2) **On the complexity of providing delay bounds:** Our original goal was to extend the prior work by Xu and Lipton [14] to the stronger model of the linear decision tree. We showed that this extension is important since GPS virtual finish times are all linear functions of the inputs, and therefore the linear decision tree accurately captures the computation power of a router. We obtained an interesting negative result: the methodology (reduce to sorting) used in [14] in proving lower bounds

```

1. FIFO_PacketQueue  $Q_1, Q_2, \dots, Q_k \leftarrow \emptyset$ ; /* Here  $n+1 = 2^k$  */
2. for ( $i = 1$ ;  $i \leq n$ ;  $i++$ ) {
3.   Find  $j, 1 \leq j \leq k$ , such that  $L_{max} * \frac{2^{j-1}-0.5}{n+1} \leq l_i < L_{max} * \frac{2^j-0.5}{n+1}$ ;
4.   Append  $p_i$  to  $Q_j$ ;
5. }
6. for ( $i = 1$ ;  $i \leq k$ ;  $i++$ ) serve packets in  $Q_i$  in the FIFO order;
7. Schedule the remaining ones in the sequential order;

```

Fig. 8. An $O(\log \log n)$ algorithm for providing $O(1)$ GPS-relative delay bound under 2PPS-CBFS scenarios

```

1. FIFO_PacketQueue  $Q_1, Q_2, \dots, Q_k \leftarrow \emptyset$ ; /* Here  $n+1 = 2^k$  */
2. FIFO_PacketQueue  $R_1, R_2, \dots, R_{k-1} \leftarrow \emptyset$ ;
3. for ( $i = 1$ ;  $i \leq n$ ;  $i++$ )
4.   if ( $l_i < L_{max} * \frac{2^{k-1}+0.5}{n+1}$ ) {
5.     Find  $j, 1 \leq j \leq k$ , such that  $L_{max} * \frac{2^{j-1}-0.5}{n+1} \leq l_i < L_{max} * \frac{2^j-0.5}{n+1}$ ;
6.     Append  $p_i$  to  $Q_j$ ;
7.   }
8.   else {
9.     Find  $j, 1 \leq j \leq k-1$ , such that  $L_{max}(1 - \frac{2^j-0.5}{n+1}) \leq l_i < L_{max}(1 - \frac{2^{j-1}-0.5}{n+1})$ ;
10.    Append  $p_i$  to  $R_j$ ;
11.   }
12. for ( $i=1$ ;  $i \leq k$ ;  $i++$ ) serve packets in  $Q_i$  in the FIFO order
13. for ( $i=1$ ;  $i \leq k-1$ ;  $i++$ ) serve packets in  $R_{k-i}$  in the FIFO order

```

Fig. 9. An $O(\log \log n)$ algorithm for providing $O(1)$ GPS-relative delay bound under 1PPS-FS scenarios

will not work for the stronger model. In particular, we show that given a packet arrival scenario used in [14], there is an $O(\log \log n)$ per packet algorithm to guarantee $O(1)$ GPS-relative delay bound under the stronger model. This demonstrates that, under such scenarios, scheduling packets to guarantee tight delay bounds is much easier than sorting the GPS virtual finish times, making the reduction to sorting impossible. This negative result indicates that a new methodology is needed in order to possibly establish this complexity lower bound, which remains an open problem so far.

REFERENCES

- [1] J. Bennett and H. Zhang. wf^2q : worst-case fair weighted fair queuing. In *IEEE INFOCOM'96*, March 1996.
- [2] J. Bennett and H. Zhang. Hierarchical packet fair queuing algorithms. *IEEE/ACM Transactions on Networking*, 5:675–689, 1997.
- [3] H. Chan and X. Guo. *Quality of Service Control in High-Speed Networks*. Wiley, 2001.
- [4] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. *Internetworking: Research and Experience*, pages 3–26, 1990.
- [5] D. Dobkin and R. Lipton. A lower bound of $\frac{1}{2}n^2$ on linear search programs for the knapsack problem. *J. Comput. Syst. Sci.*, 16:413–417, 1978.
- [6] S. Golestani. A self-clocked fair queueing scheme for broadband applications. In *Proc. of Infocom'94*, June 1994.
- [7] P. Goyal, H. Vin, and H. Cheng. Start-time fair queuing: A scheduling algorithm for integrated services packet switching networks. In *CS-TR-96-02, Dpt. of Computer Science, Univ. of Texas, Austin*, 1996.
- [8] A. Greenberg and N. Madras. How fair is fair queuing? *Journal of the ACM*, 39(3):568–598, 1992.
- [9] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single node case. *IEEE/ACM Transaction on Networking*, 1(3):344–357, June 1993.
- [10] M. Sipser. *Introduction to the Theory of Computation*. PWS, 1997.
- [11] D. Stiliadis and A. Varma. Design and analysis of frame-based fair queuing: A new traffic scheduling algorithm for packet switched networks. In *Proc. of ACM Sigmetrics'96*, pages 104–115, May 1996.
- [12] D. Stiliadis and A. Varma. Latency-rate servers: a general model for analysis of traffic scheduling algorithms. In *Proc. of Infocom'96*, March 1996.
- [13] J. Turner. New directions in communications (or which way to the information age?). *IEEE Communications Magazine*, 24:8–15, October 1986.
- [14] J. Xu and R. Lipton. On fundamental tradeoffs between delay bounds and computational complexity in packet scheduling algorithms. In *Proc. of ACM SIGCOMM 2002*, August 2002.
- [15] H. Zhang. Service disciplines for guaranteed performance service in packet switching networks. *Proceedings of the*

