A Proportional Share Resource Allocation Algorithm For Real-Time, Time-Shared Systems

Ion Stoica * Hussein Abdel-Wahab † Kevin Jeffay‡ Sanjoy K. Baruah §
Johannes E. Gehrke ¶ C. Greg Plaxton ||

Abstract

We propose and analyze a proportional share resource allocation algorithm for realizing real-time performance in time-shared operating systems. In a proportional share system, processes are assigned a weight which determines a share (percentage) of the resource they are to receive. The resource is then allocated in discrete-sized time quanta in such a manner that each process makes progress at a precise, uniform rate. Proportional share allocation algorithms are of interest because (1) they provide a natural means of seamlessly integrating real- and non-real-time processing requirements in a general purpose operating system, (2) they are easy to implement (and in particular, easier than more traditional forms of real-time support such as periodic tasks), (3) they provide a simple and effective means of precisely controlling the real-time performance of a process including uniform, predictable degradation in times of system overload, and (4) they provide a natural means of policing processes so that process that use more of a resource than they request have no ill-effect on well-behaved processes.

We present our algorithm and its analysis in the context of an idealized system in which a resource is assumed to be granted in arbitrarily small intervals of time and show that our algorithm guarantees that the difference between the service time that a process should receive in the idealized system and the service time it actually receives in the real system is bounded by the size q of a time quantum. This demonstrates that our algorithm is capable of guaranteeing real-time response times to processes. Moreover, we also show that these bound are the best obtainable in a proportion share system and hence that our algorithm provides optimal performance. Lastly, the algorithm provides support for dynamic operations, such as processes joining or leaving the competition, and for both fractional and non-uniform time quanta. As a proof of concept we have implemented a prototype of a CPU scheduler under FreeBSD operating systems. The experimental results shows that our implementation perform within the theoretical bounds and hence supports real-time execution in a general purpose operating system.

1 Introduction

Currently there is great interest in providing real-time execution and communication support in general purpose operating systems. Indeed, applications such as desktop videoconferencing, distributed shared virtual environments, and collaboration-support systems require real-time computation and communication services to be effective. At present the dominant approach to providing real-time support in a general purpose operating system is to embed a periodic thread or process model into an existing operating system kernel and to use a real-time scheduling algorithm such as rate-monotonic scheduling to schedule the processes. In such as system, aperiodic and non-real-time activities are typically scheduled either as background processes or through the use of a second-level scheduler that is executed quasi-periodically as a server process by the real-time scheduler.

*Department of Computer Science, Old Dominion University, Norfolk, VA 23529-0162 (stoica@cs.odu.edu).
†Department of Computer Science, Old Dominion University, Norfolk, VA 23529-0162 (wahab@cs.odu.edu).
‡Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3175 (jeffay@cs.unc.edu).
§Department of Computer Science, University of Vermont, Burlignton, VT 05405 (sanjoy@cs.uvm.edu).
¶Department of Computer Science, University of Wisconsin-Madisson, Madison, WI 53706-1685 (johannes@cs.wisc.edu).
||Department of Computer Science, University of Texas at Austin, Austin, TX 78712-1188 (plaxton@cs.utexas.edu).

In general, this framework can be quite effective at integrating real-time and non-real-time computing. However, we observe that this approach has yet to be embraced by the larger operating systems community. We believe that this is due in part to the rigid distinctions made between real-time and non-real-time. Real-time activities are programmed differently than non-real-time ones (e.g., as periodic tasks) and real-time activities receive hard and fast guarantees of response time (if admission control is performed). Non-real-time activities are subservient to real-time ones and receive no performance guarantees. While this state of affairs is entirely acceptable for many mixes of real-time and non-real-time activities, for many it is not. For example, consider the problem of supporting real-time videoconferencing on the desktop. This is clearly a real-time application, however, it is not one for which hard and fast guarantees of real-time performance are required. For example, it is easy to imagine situations in which one would like to explicitly degrade the performance of the videoconference (e.g., degrade the rate at which the video is displayed) so that other activities, such as searching a large mail database for a particular message, can complete quicker. In the extreme, one might want to swap the roles of these two processes — to execute the search process in real-time so that given knowledge of the number of files to search, a precise upper bound on the completion time for the search can be predicted, and to execute the video conference in a background, "best-effort" manner. Ideally, a general purpose operating system that supports real-time execution should not a priori restrict the basic tenor of performance guarantees that any process is capable of obtaining.

To address this issue we are investigating an alternate approach to realizing real-time performance in time shared operating systems, namely the use of *proportional share resource allocation* algorithms for processor scheduling. In a proportional share allocation system, every process in the system is guaranteed to make progress at a well-defined, uniform rate. Specifically, each process is assigned a share of the processor — a percentage of the processor's total capacity. If a process's share of the processor is s then in any interval of length t, the process is guaranteed to receive $(s \times t) \pm \epsilon$ units of processor time where $0 \leq \epsilon \leq \delta$, for some constant δ. In a proportional share system, resource allocation is flexible and the share of the processor that a process receives can be changed dynamically. In this manner a process's real-time rate of progress can be explicitly controlled.

Proportional share resource allocation algorithms lie between traditional general purpose and real-time scheduling algorithms. On the one hand, proportional share resource allocation is a variant of the *pure processor sharing* scheduling discipline commonly studied in introductory operating systems courses. In pure processor sharing if there are n processes present in the system, during each time unit each process receives $1/n$ of the processor's capacity. Thus each process appears as it is making uniform progress on a virtual processor that has $1/n$ of the capacity of the physical processor.

On the other hand, traditional real-time scheduling disciplines for periodic tasks can be viewed as coarse approximations of proportional share resource allocation. For example, if a periodic task requires c units of processor time every p time units, then a rate-monotonic scheduler guarantees that for all $k \geq 0$, in each interval $[kp, (k+1)p]$, the periodic task will indeed receive a share of the processor equal to c/p. Specifically, in each of the above intervals, the process will receive $pc/p = c$ units of processor time.

The key distinctions between the proportional share resource allocation policy we develop here and the pure processor sharing discipline is that ours is a discrete approximation to, and a generalization of, the pure processor sharing discipline. Whereas the pure processor sharing discipline is a theoretically ideal algorithm that allocates arbitrarily fine fractional time units to processes, our proportional share algorithm allocates discrete units of the shared resource to processes. Thus in our system a process makes progress at a less than perfect uniform rate. However, our main result is that over any time interval, the deviation between the service time received by a process from a shared resource, and the service time the process should have received if it were making progress at a uniform rate, is bounded. Moreover we show that the bounds we obtain for our scheduling algorithm are optimal in the sense that it is not possible to develop a proportional share algorithm that provides better bounds. This means that one can analyze the real-time performance of *any* process in our system. In addition, by changing scheduling parameters, one can precisely control the real-time performance of a process.

Our proportional share resource allocation policy differs from traditional methods of integrating real- and non-real-time processes in that here all processes, real- and non-real-time, are treated identically. In a proportional share system, real-time (and non-real-time) processes can be implemented very much like traditional processes in a time-shared operating system. Thus in terms of the process model, no "special" support is required to support real-time computing — there is only one type of process. Moreover, as we shall see, proportional share resource allocation algorithms, like many scheduling algorithms used in time-shared systems, allocate resources in discrete units or quanta and hence are easier to realize in time-shared operating systems than traditional real-time policies which are typically event-driven and require the ability to preempt processes at potentially arbitrary points. In addition, because resources are allocated in discrete quanta in a proportional share system, one can better control (and account for) the overhead of the scheduling mechanism as well as tune the system to trade-off fine-grain, real-time control for low scheduling and system overhead. Finally, proportional share algorithms provide a natural means of uniformly degrading system performance in overload situations.

In this paper we present a proportional share processor scheduling algorithm and demonstrate that it can be used to ensure predictable real-time response to all processes. Section 2 discusses related work in scheduling. Section 3 presents our process model and formally introduces the concepts of a share and the requirement for predictable execution with respect to a share. Section 4 presents a deadline-based, virtual-time scheduling algorithm that is used to ensure processes receive their requested share of the processor. Section 5 introduces a key technical problem to be solved in the course of applying our algorithm, namely that of dealing with the dynamic creation and destruction of processes. These events are significant as they indicate the points in time at which the share of the processor to be received by an active process, and hence its real-time rate of progress, potentially changes. Section 6 outlines the proof of correctness of our algorithm and Section 7 presents some experimental results using our proportional share system in the FreeBSD operating system. In particular, we demonstrate how "traditional" real-time processes such as periodic tasks can be realized in a proportional share system.

2 Related Work

Tijdeman was one of the first to formulate and analyze the proportional share allocation problem [14]. The original problem was stated in terms of selecting a union chairman every year, such that the accumulated number of chairmen from each state (of the union) to be proportional to its weight. As shown in [2], Tijdeman's results can be easily applied to solve the proportional share allocation problem. In the general setting, the resource is allocated in fixed time quanta, while the clients' shares may change at the beginning of every time quantum. In this way dynamic operation can be easily accommodated. Tijdeman proved that if the clients' shares are known in advance there exists a schedule with the lag bound less or equal to $1 - 1/(2n-2)$, where n represents the total number of clients. (Note that when $n \to \infty$ the lag bound approaches unity.) Although he gives an optimal algorithm for the *static* case (i.e., when the number of clients does not change over time), he does not give any explicit algorithm for the dynamic case. Furthermore, we note that, even in the general setting, the problem formulation does not accommodate fractional or non-uniform quanta.[1]

Recently, the proportional share allocation problem has received a great deal of attention in the context of various resources, such as communication bandwidth and processor time. Our algorithm is closely related to weighted fair queueing algorithms previously developed for bandwidth allocation in communication networks [4, 5, 9], and general purpose proportional share algorithms, such stride scheduling [16, 17]. Demers, Keshav, and Shenker were the first to apply the notion of fairness to a fluid-flow system that models an idealized communication switch in which sessions are serviced in arbitrarily small increments [4]. Since in practice a packet transmission cannot be preempted, the authors proposed an algorithm, called Packet Fair Queueing (PFQ), in which the packets are serviced in the order in which they would finish in the corresponding fluid-flow system (i.e., in the increasing order of their virtual deadlines). By using the concept of *virtual time*, previously introduced by Zhang [18], Parekh and Gallager have analyzed the PFQ[2] algorithm when the input traffic stream conforms to the *leaky-bucket* constraints [9, 10]. In particular, they shown that no packet is serviced T_{max} latter than it would have been serviced in the fluid-flow system, where T_{max} represents the time to transmit a packet of maximum size. We note that this translates into an $O(1)$ bound for the *positive* lag. However, as shown in [3, 12, 17], the bound for the *negative* lag can be as large as $O(n)$, where n represents the number of active sessions (clients) in the system. Moreover, in PFQ the virtual

[1] The difference between fractional and non-uniform quanta is that while in the first case the fraction from the time quantum (that the client will actually use) is assumed to be known in advance, in the non-uniform quanta case this fraction is not known.

[2] In their work PFQ is known as Packet-by-Packet Generalized Processor Sharing (PGPS for short).

time is updated when a client joins or leaves the competition in the ideal system, and not in the real one. This requires to maintain an additional event queue, which makes the implementation complex and inefficient. As a solution, Golestani has proposed a new algorithm, called Self-Clocked Fair Queueing (SCFQ), in which the virtual time is updated when the client joins/leaves the competition in the real system, and not in the idealized one [5]. Although this scheme can be more efficiently implemented, this does not come for free: the lag bounds increase to within a factor of two of the ones guaranteed by PFQ. Finally, we note that all the dynamic operations, such as a client leaving or joining the competition, can be implemented in $O(\log n)$ time in both PFQ and SCFQ.

Recently, Waldspurger and Weihl have developed a new proportional share allocation algorithm, called stride scheduling [16, 17], which can be viewed as a cross-application of fair queueing to the domain of processor scheduling. Stride scheduling relies on the concept of *global pass* (which is similar to virtual time) to measure the work progress in the system. Each client has an associated *stride* that is inversely proportional to its weight, and a *pass* that measures the progress of that client. The algorithm allocates a time quantum to the client with the lowest pass, which is similar to the PFQ policy. However, by grouping the clients in a binary tree, and recursively applying the basic stride scheduling algorithm at each level, the lag is reduced to $O(\log n)$. Moreover, stride scheduling provides support for both uniform and non-uniform quanta.

In contrast to the above algorithms, by making use of both virtual eligible times and virtual deadlines, EEVDF achieves constant lag bounds, while providing full support for dynamic operations. We note that two similar algorithms were independently developed (in parallel to our original work [12]) by Bennett and Zhang in the context of allocating bandwidth in communication networks [3], and by Baruah, Gehrke and Plaxton in the context of processor scheduling for fixed time quanta [2].[3] In addition to introducing the concept of virtual eligible time (which was also independently introduced in [2] and [3]) our work makes several unique key contributions.

First, by "decoupling" the request size from the size of the time quantum we generalize the previous known theoretical results [9]. Moreover, our analyze can be easily extended to preemptive systems, as well. For example, we can derive lag bounds for a fully preemptive system, by simply taking time quanta to be arbitrarily small. Similarly, by taking the size of a time quantum to be the maximum duration of a critical region, we can derive lag bounds for a preemptive system with critical regions. Finally, by decoupling the request size from the size of a time quantum, a client has the possibility to trade between the allocation accuracy and the scheduling overhead (see Section 6).

Second, we address the problem of a client leaving the competition *before* using the entire service time it has requested. This is an important extension since in the processor case it is hard (sometimes impossible) to predict *exactly* how much service time a client will use for the next request. We note that this problem does not occur and consequently has not been addressed in the context of bandwidth allocation; in this case, the length of a message and therefore its transmission time are assumed to be known upon its arrival. The only previous known algorithms that address this problem are lottery and stride scheduling [16, 17]. However, the lag bounds guaranteed by stride scheduling are as large as $O(n)$, where n represents the number of active clients (being a randomized algorithm, lottery does not guarantee tight bounds). In comparison, EEVDF guarantees optimal lag bounds of one time quantum.

Third, we propose a new approximation scheme for maintaining the virtual time, in which the updating operations are performed when the events (e.g., client leaving, joining) occur in the real system, and not in the ideal one. This simplifies the implementation and eliminates the need of keeping an event queue. It is worth to mention that unlike other previous approximations [5], ours gurantees optimal lag bounds.

Besides the class of fair queuing algorithms, recently, a significant number of other proportional share algorithms have been developed [1, 8, 11, 15]. Although none of them guarantees constant lag bounds in a *dynamic* system, we note that the PD algorithm of Baruah, Gehrke, and Plaxton [1] achieves constant lag bounds in a *static* system.

The idea to apply fair queueing algorithms to processor scheduling was first suggested by Parekh in [10]. Waldspurger and Weihl were the first to actually develop and implement such an algorithm (stride scheduling) for processor scheduling [16, 17].[4] Finally, to our best knowledge we are the first to implement and to test with a proportional share scheduler which guarantees constant lag bounds.

[3] This can be viewed as a direct extension to the dynamic case of the Tijdeman's work.

[4] We note that they have also applied stride scheduling to other shared resources, such as lock access.

3 The Model

We consider an operating system to consist of a set of processes (real-time and non-real-time) that compete for a time shared resource such as a CPU or a communications channel. To avoid confusion with terminology used in the experimental section of the paper we use the term *client* to refer to computational entities (i.e., processes). A client is said to be *active* while it is competing for the resource, and *passive* otherwise. We assume that the resource is allocated in time quanta of size at most q. At the beginning of each time quantum a client is selected to use the resource. Once the client acquires the resource, it may use it either for the entire time quantum, or it may release it before the time quantum expires. Although simple, this model captures the basic mechanisms traditionally used for sharing common resources, such as processor and communication bandwidth. For example, in many preemptive operating systems (e.g., UNIX, Windows-NT), the CPU scheduler allocates the processing time among competing processes in the same fashion: a process uses the CPU until its time quantum expires or another process with a higher priority becomes active, or it may voluntarily release the CPU while it is waiting for an event to occur (e.g., an I/O operation to complete). As another example, consider a communication switch that multiplexes a set of incoming sessions on a packet-by-packet basis. Since usually the transmission of a packet cannot be preempted, we take a time quantum to be the time required to send a packet on the output link. Thus, in this case, the size q of a time quantum represents the time required to send a packet of maximum length.

Further, we associate to each client a *weight* that determines the relative *share* of the resource that it should receive. Let w_i denote the weight associated to client i, and let $\mathcal{A}(t)$ be the set of all clients active at time t. We define the (instantaneous) share $f_i(t)$ of an active client i at time t as

$$f_i(t) = \frac{w_i}{\sum_{j \in \mathcal{A}(t)} w_j}. \qquad (1)$$

If the client's share remains constant during a time interval $[t, t+\Delta t]$, then it is entitled to use the resource for $f_i(t)\Delta t$ time units. In general, when the client share varies over time, the service time that client i should receive in a *perfect fair* system, while being active during a time interval $[t_0, t_1]$, is

$$S_i(t_0, t_1) = \int_{t_0}^{t_1} f_i(\tau) d\tau \qquad (2)$$

time units. The above equation corresponds to an ideal fluid-flow system in which the resource can be granted in arbitrarily small intervals of time[5]. Unfortunately, in many practical situations this is not possible. One of the reasons is the overhead introduced by the scheduling algorithm itself and the overhead in switching from one client to another: taking time quanta of the same order of magnitude as these overheads could drastically reduce the resource utilization. Another reason is that some operations cannot be interrupted, i.e., once started they must complete in the same time quantum. For example, once a communication switch begins to send a packet of one session, it cannot serve any other session until the entire packet is sent. As another example, a process cannot be preempted while it is in a critical section. Thus, in the first example we can chose the size of a quantum q as being the time required to send a packet of maximum length, while in the second example we can chose q as being the maximum duration of a critical section.

Due to quantization, in a system in which the resource is allocated in discrete time quanta it is not possible for a client to always receive *exactly* the service time it is entitled to. The difference between the service time that a client should receive at a time t, and the service time it *actually* receives is called service time *lag* (or simply lag). Let t_0^i be the time at which client i becomes active, and let $s_i(t_0^i, t)$ be the service time the client receives in the interval $[t_0^i, t]$ (here, we assume that client i is active in the entire interval $[t_0^i, t]$). Then the lag of client i at time t is

$$lag_i(t) = S_i(t_0^i, t) - s_i(t_0^i, t). \qquad (3)$$

Since the lag quantifies the allocation accuracy, we use it as the main parameter in characterizing our proportional share algorithm. In particular, what we show below is that our proportional share algorithm (presented next) (1) provides bounded lag for all clients, and that (2) this bound is optimal in the sense that it is not possible to develop an algorithm that affords better bounds. Together, these properties indicate that our algorithm will provide real-time response guarantees to clients and that with respect to the class of proportional share algorithms, these guarantees are the best possible.

4 The EEVDF Algorithm

In order to obtain access to the resource, a client must issue a *request* in which it specifies the duration of the service time it needs. Once a client's request is fulfilled, it may either issue a new request or

[5] A similar model was used by Demers *et al* [4] in studying fair-queuing algorithms in communication networks.

become passive. For uniformity, throughout this paper we assume that the client is the sole initiator of the requests[6]. For flexibility we allow the requests to have any duration. Note that a client may request the same amount of service time by generating either fewer longer requests, or many shorter ones. For example, a client may ask for one minute computation time either by issuing 60 requests with a duration of one second each, or by issuing 600 requests with a duration of 100 ms each. As we will show in Section 6, shorter requests guarantee better allocation accuracy, while longer requests decrease the system overhead. In this way a client could trade between the allocation accuracy and scheduling overhead.

We formulate our scheduling algorithm in terms of the behavior of an ideal, fluid-flow system that executes clients in a *virtual-time* time domain [18, 9].[7] Abstractly, the virtual fluid-flow system executes each client for w_i real-time time units during each virtual-time time unit. More concretely, virtual-time is defined to be the following function of real-time

$$V(t) = \int_0^t \frac{1}{\sum_{j \in \mathcal{A}(\tau)} w_j} d\tau. \tag{4}$$

Note that virtual-time increases at a rate inversely proportional to the sum of the weights of all active clients. That is, when the competition increases virtual-time slows down, while when the competition decreases it accelerates. Intuitively, the flow of virtual-time changes to "accommodate" all active clients in one virtual-time time unit. That is, the size of a virtual-time unit is modified such that in the corresponding fluid-flow system each active client i receives w_i real-time units during one virtual-time time unit. For example, consider two clients with weights $w_1 = 2$ and $w_2 = 3$. Then the rate at which virtual-time increases is $\frac{1}{w_1 + w_2} = 0.2$, and therefore a virtual-time time unit equals five real-time units. Thus, in each virtual-time time unit the two clients should receive $w_1 = 2$, and $w_2 = 3$ time units.

Ideally we would like for our proportional share algorithm to approach the behavior of the virtual fluid-flow system. Thus, since in the fluid-flow system, at all points in time a client is best characterized by the service time it has received up to the current time, to compare our approach with the ideal, we must be able to compute the service time that a client should receive in the fluid-flow system. By combining Eq. (1) and (2) we can express the service time that an active client i should receive in the interval $[t_1, t_2)$ as

$$S_i(t_1, t_2) = w_i \int_{t_1}^{t_2} \frac{1}{\sum_{j \in \mathcal{A}(\tau)} w_j} d\tau. \tag{5}$$

Once the integral in the above equation is computed, we can easily determine the service time that *any* client i should receive during the interval $[t_1, t_2)$, by simply multiplying the client's weight by the integral's value. Next, from Eq. (5) and (4) it follows that

$$S_i(t_1, t_2) = (V(t_2) - V(t_1)) w_i. \tag{6}$$

To better interpret the above equation consider a much simpler model in which the number of active clients is constant and the sum of their weights is one ($\sum_{i \in \mathcal{A}} w_i = 1$), i.e., the share of a client i is $f_i = w_i$. Then, in this model, the service time that client i should receive during an interval $[t_1, t_2)$ is simply $S_i(t_1, t_2) = (t_2 - t_1) w_i$. Next, notice that by replacing the real times t_1 and t_2 with the corresponding virtual-times $V(t_1)$ and $V(t_2)$ we arrive at Eq. (6). Thus, Eq. (6) can be viewed as a generalization for computing the service time $S_i(t_1, t_2)$ in a *dynamic* system — one in which clients are dynamically (joining and leaving the competition (i.e., being created and destroyed).

Our scheduling algorithms uses measurements made in the virtual-time domain to make scheduling decisions. To each client's request we associate an *eligible time* e and a *deadline* d which represent the starting and finishing time respectively for the request's service in the corresponding fluid-flow system. Let t_0^i be the time at which client i becomes active, and let t be the time at which it initiates a new request. Then, a request becomes eligible at a time e when the service time that the client should receive in the corresponding fluid-flow system $S_i(t^0, e)$ *equals* the service time that the client has already received in the real system $s_i(t^0, t)$, i.e., $S_i(t_0^i, e) = s_i(t_0^i, t)$. Note that if at time t client i has received more service time than it was supposed to receive (i.e., $lag_i(t) < 0$), then it will be the case that $e > t$ and hence the client should wait until time e before the new request becomes eligible. In this way a client that has received more service time than its share is "slowed down", while giving the other active clients the opportunity to "catch up". On the other hand, if at time t client i has received less service time than it was supposed to receive (i.e., its lag is positive), then it will be the case that $e < t$,

[6] In practice this is not necessarily true. For example, in the processor case the scheduler itself could be the one to issue the requests on behalf of the client. In this case, the request duration (length) is either specified by the client, or otherwise the scheduler assumes a "default" duration. This allows us to treat all continuous media, interactive, and batch activities in a consistent way.

[7] We note that Waldspurger and Weihl [16, 17] use a similar concept (denoted *global pass*) in their stride scheduling.

and therefore the new request is immediately eligible at time t. By using Eq. (6) we can express the virtual eligible time $V(e)$ as

$$V(e) = V(t_0^i) + \frac{s_i(t_0^i, t)}{w_i}. \quad (7)$$

Similarly, the *deadline* of the request is chosen such that the service time that the client should receive between the eligible time e and the deadline d *equals* the service time of the new request, i.e., $S_i(e, d) = r$, where r represents the length of the new request. Further, by using again Eq. (6), we derive the virtual deadline $V(d)$ as

$$V(d) = V(e) + \frac{r}{w_i}. \quad (8)$$

Notice that although Eq. (7) and (8) give us the virtual eligible time $V(e)$ and the virtual deadline $V(d)$, they do not *necessarily* give us the values of the real times e and d ! To see why, consider the case in which e is larger than the current time t. Then e cannot be computed exactly from Eq. (4) and (7), since we do *not* know how the slope of the virtual-time mapping will vary in the future (it changes dynamically while clients join and leave the competition). Therefore we will formulate our algorithm in terms of *virtual* eligible times and deadlines and not of the real times. With this, the Earliest Eligible Virtual Deadline First (EEVDF) algorithm can be simply stated as follows:

EEVDF Algorithm. *A new quantum is allocated to the client which has the eligible request with the earliest virtual deadline.*

Since EEVDF is formulated in terms of virtual-times, in the remaining of this paper we use ve and vd to denote the virtual eligible time and virtual deadline respectively, whenever the corresponding real eligible time and the deadline are not given. Let $r^{(k)}$ denote the length of the k^{th} request made by client i, and let $ve^{(k)}$ and $vd^{(k)}$ denote the virtual eligible time and the virtual deadline associated to this request. If the client uses each time the *entire* service time it has requested, then by using Eq. (7) and (8) we obtain the following recurrence which computes both the virtual eligible time and the virtual deadline of each request:

$$\begin{aligned} ve^{(1)} &= V(t_0^i), && (9) \\ vd^{(k)} &= ve^{(k)} + \frac{r^{(k)}}{w_i}, && (10) \\ ve^{(k+1)} &= vd^{(k)}. && (11) \end{aligned}$$

Next, we consider the more general case in which the client does not use the entire service time it has requested. Since a client never receives more service time than requested, we need to consider only the case when the client uses the resource for less time than requested. Let $u^{(k)}$ denote the service time that client i actually receives during its k-th request. Then the only change in Eq. (9)–(11), will be in computing the eligible time of a new request. Specifically, Eq. (11) is replaced by

$$ve^{(k+1)} = ve^{(k)} + \frac{u^{(k)}}{w_i}. \quad (12)$$

Example. To fix the ideas, let us take a simple example (see Figure 1). Consider two clients with weights $w_1 = w_2 = 2$ that issue requests with lengths $r_1 = 2$, and $r_2 = 1$, respectively. We assume that the time quantum is of unit size ($q = 1$) and that client 1 is the first one which enters competition at real time $t_0 = 0$. Thus, according to Eq. (9) and (10) the virtual eligible time for the first request of client 1 is $ve = 0$, while its virtual deadline is $vd = 1$. Being the single client that has an outstanding eligible request, client 1 receives the first quantum. At real time $t = 1$, client 2 enters the competition. Since during the interval $[0, 1)$ the only active client in the system is client 1, from Eq. (4), the value of virtual-time at real-time 1 is $V(1) = \int_0^1 \frac{1}{w_1} d\tau = 0.5$. Thus, virtual-time increases at half the rate of real-time. In this way, in an ideal system, during every virtual-time time unit, client 1 receives $w_1 = 2$ real time units. Next, after the second client enters the competition, the rate of virtual-time slows down further to $\frac{1}{w_1+w_2} = 0.25$. Hence, in the ideal system during one virtual-time time unit, each client will receive 2 real time units (since $w_1 = w_2 = 2$). Next, assume that client 2 issues its first request just before the second quantum is allocated. Then at real time $t = 1$ there are two pending requests: one of client 1 with the virtual deadline 1 (which waits for another time quantum to fulfill its request), and one of client 2 which has the same virtual deadline, i.e., 1. In this situation we arbitrarily break the tie in favor of client 2, which therefore receives the second quantum. Since this quantum fulfills the current request of client 2, client 2 issues immediately, at real time 3 (virtual-time 1), a new request. From Eq. (11) and (10) the virtual eligible time and the virtual deadline of the new request are 1 and 1.5, respectively. Thus, at real time $t = 2$ (virtual-time 0.75) the single eligible request is the one of client 1, which therefore receives the next quantum. Further, at real time $t = 3$ (virtual-time 1) there are again two eligible requests: the one of client 2 that has just become eligible, and the new request issued by client 1. Since the virtual

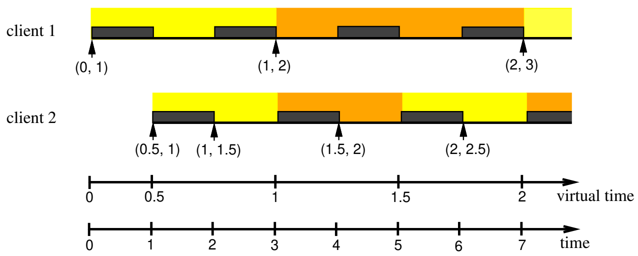

An example of EEVDF scheduling involving two clients with equal weights $w_1 = w_2 = 2$. All the requests generated by client 1 have length 2, and all of the requests generated by client 2 have length 1. Client 1 becomes active at time 0 (virtual-time 0), while client 2 becomes active at time 1 (virtual-time 0.5). The vertical arrows represent the times when the requests are initiated (the pair associated to each arrow represents the virtual eligible time and the virtual deadline of the corresponding request). The shaded regions in the background show the the durations of servicing successive requests (of the same client) in the fluid-flow system (note that the starting and the finishing times of servicing a request in the ideal system corresponds to the virtual eligible time and the virtual deadline of the corresponding request).

deadline of the second client's request (1.5) is earlier than the one of the first client (2), the fourth quantum is allocated to client 2. Further, Figure 1 shows how the next four quanta are allocated.

Note the uniform progress of the two clients in Figure 1. Although the uniformity is perfect in this contrived example, we show in Section 6 that in fact the deviation of a client's progress from the perfectly uniform rate (i.e., its rate of progress in the ideal fluid-flow system) is bounded and that these bounds are the best possible. This shows that for a given quanta q, the EEVDF algorithm provides the best possible guarantees of real-time progress.

5 Fairness in Dynamic Systems

In this section we address the issue of *fairness* in dynamic systems. Throughout this paper, we assume that a dynamic system provides support for client joining and leaving the competition[8]. To understand the main issues of implementing dynamic operations, first recall that the client's lag represents (see Eq. (3)) the difference between the service time that the client *should* receive and the service time it *has* actually received. An important property of the EEVDF algorithm is that at *any* time the sum of the lags of all active clients is zero (see Lemma 2). Thus, if a client leaves the competition with a negative lag (i.e., after receiving more service time than it was supposed to), the remaining clients should have received *less* service time than they were entitled to. In short, a gain for one client translates into a loss for the other active clients. Similarly, when a client with positive lag leaves, this translates into a gain for the remaining clients. The main question here is how to distribute this loss/gain among the remaining clients. In [12] we answered this question by distributing it in *proportion* to the clients' weights. In the remaining of this section we show that the same answer is obtained by approaching the problem from a different angle.

The basic observation is that this problem does not occur as long as a client with zero lag leaves the competition, because there is nothing to distribute. Since in the corresponding fluid-flow system the lag of any client is *always* zero, a simple solution would be to consider the time when the client leaves to be the time when it leaves in the corresponding fluid-flow system, and not in the real system. Unfortunately, this solution has two major drawbacks. First, in many situations, such as scheduling incoming packets in a high speed networking switch, maintaining the events in the fluid-flow system is too expensive in practice [5]. Second and more important, this solution assumes implicitly that the service time that a client will use is known in advance. While this is generaly true in the case of the communication switch, where the length of a message (and consequently its service time) is assumed to be known when the packet arrives, in the processor case this is not always possible. To see why this is a potential problem, consider the previous example (see Figure 1) in which the first client leaves the competition after using only 1.1 time-units of the second request, i.e., at time 6.1 in the real system and the corresponding virtual time 1.725. However, according to Eq. (12), in the ideal system the client should complete its service and therefore leave the competition at the virtual time 1.525 ($= ve^{(2)} + u^{(2)}/w_1$), which in our example corresponds to the real time 5.2. Unfor-

[8]Note that having this two operation, changing the client weight can be easily implemented as a leave followed by a rejoin operation [12].

tunately, since at this point we do not know for how long client 1 will continue to use the resource (we know only that it has made a *two* time units request from which it has received one time unit, and it still has to receive another one) we cannot update the virtual time correctly !

Next, we present our solution to this problem for a dynamic system in which the following two (reasonable) restrictions hold: (1) all the clients that join the competition are assumed to have zero lag, and (2) rejoins the competition. a client has to leave the competition as soon as it is finished using the resource (i.e., when a client terminates it is not allowed to remain in the system). We consider two cases whether the client's lag is negative, or positive. From Eq. (3), (4), (6) it follows that the client's lag increases as long as the client receive service time, and decreases otherwise. Thus, when a client with negative lag wants to leave, we can simply delay that client (without allocating any service time to it) until its lag becomes zero. This can be simply accomplished by generating a dummy request of zero length. However, note that since a request cannot be processed before it becomes eligible, and since the virtual eligible time of the dummy request is equal to its deadline (see Eq. (8)), this request cannot be processed earlier than its deadline. In this way, we have reduced the first case to the second one, in which the client leaving the competition has a positive lag. Next, we consider only the latter case.

Our first solution is based on the same simple idea: the client is delayed until its lag becomes zero. For clarity, consider the example in Figure 2(a), where three clients become simultaneously active. Next, suppose that at time t_1, client 1 decides to leave the competition while having a positive lag. Then the client will be simply delayed, while continuing to receive service time, until its lag becomes zero, i.e., until time t_1'. If we assume that the slope of virtual-time with respect to real-time does not change between t_1 and t_1', then from Eq. (5) and (6) we obtain $S_1(t_1, t_1') = (V(t_1') - V(t_1))w_1 = w_1(t_1' - t_1)/(w_1 + w_2 + w_3)$. Further, by using Eq. (3) and (5), and the fact that $s_1(t_1, t_1') = t_1' - t_1$ we can compute the virtual-time at t_1' as

$$V(t_1') = V(t_1) + \frac{lag_1(t_1)}{w_2 + w_3} \qquad (13)$$

The main drawback of this approach is that client 1 continues to receive service time between t_1 and t_1', although it does not need it (since it has already finished using the resource) ! Thus, this service time will be wasted, which is unacceptable. Our solution is to simply let any client with positive lag to leave immediately, while correctly updating the value of virtual-time to account for the change (see Figure 2(b)). In this way the virtual-times corresponding to the times when a client decides to leave and when it actually leaves are the *same* in both systems. More precisely consider a client k leaving the competition at time t_k with a positive lag (i.e., $lag_k(t_k) > 0$). Then, by generalizing Eq. (13), the value of virtual-time is updated as follows

$$V(t_k) = V(t_k) + \frac{lag_k(t_k)}{\sum_{j \in \mathcal{A}(t_k) \setminus \{k\}} w_j}, \qquad (14)$$

where $\mathcal{A}(t_k)$ represents the set of all active clients just *before* client k leaves. For example, in Figure 2(b) $\mathcal{A}(t_1) = \{1, 2, 3\}$. Thus, $\mathcal{A}(t_k) \setminus \{k\}$ represents the set of all active clients just *after* client k leaves the competition. Further note that according to Eq. (3) the lag of any remaining client $i \in \mathcal{A}(t_k) \setminus \{k\}$ changes to

$$lag_i(t_k) = lag_i(t_k) + w_i \frac{lag_k(t_k)}{\sum_{j \in \mathcal{A}(t_k) \setminus \{k\}} w_j}. \qquad (15)$$

Thus the lag of client i is *proportionally* distributed among the *remaining* clients, which is consistent with our interpretation of fairness in dynamic systems, i.e., a gain/loss is proportionally distributed among the remaining clients.

Since virtual-time is updated only when the events actually occur in the real system (as opposed to when they occur in the ideal one), the EEVDF algorithm can be easily and efficiently implemented. Even in a system in which the service times are known in advance (e.g., packet transmission times in communication switches), and therefore it is theoretically possible to update virtual-time as in the ideal system, in practice this is hard to achieve. Mainly, this is because we need to implement an event queue which has to trade between the timer granularity and the overhead. As we have shown in [12] all the basic operation required to implement the EEVDF algorithm, i.e., inserting and deleting a request, and finding the eligible request with the earliest deadline can be implemented in $O(\log n)$, where n represents the number of active clients.

We note that in the worst case it may be possible that all the dummy requests occur at the same time. In this situation, the scheduler should perform $O(n)$ deletions before the next "real" request is serviced. Although, this is a potential problem in the case of a communication switch, where the selection of the next packet to be serviced is assumed to be done *during* the transmission of the current packet, it does not

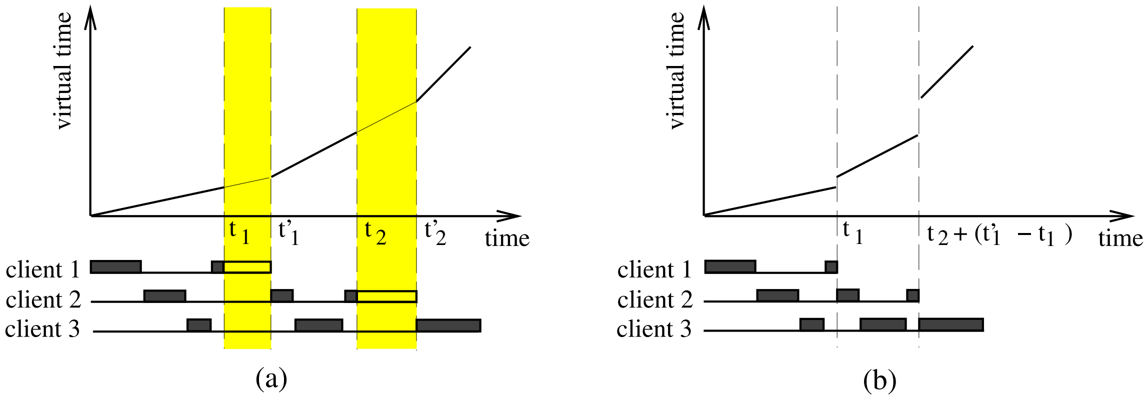

Three clients become active at the same time, after which client 1 and client 2, both with positive lags, leave the competition. In (a) clients are allowed to leave only after their lags become zero; in (b) clients are allowed to leave immediately. The shaded regions in (a) represents the time intervals during which the system allocates service time to the clients with a positive lag that want to leave (thus delaying their departure from the system until their lags become zero). In both cases the virtual-time just before a client wants to leave and just after it has actually left are equal.

significantly increase the complexity in the processor case. This is mainly because a processor, besides servicing the active clients (processes), also executes the scheduling algorithm, as well as other related operating system functions (e.g., starting a new process, or terminating an existing one). Consequently, in a complete model we need to account for these overheads anyway. A simple solution would be to charge each process for the related overheads. For example, the time to select the next process to receive a time quantum should be charged to that client. In this way, from the processor perspective, a dummy request is *no* longer a 0-duration request, since it should account at least for the scheduling overhead (and eventually for the process termination). In the current model we ignore these overheads, which, as the experimental results suggest (see Section 7), is an acceptable approximation for many practical situations. However, we plan to address this aspect in the future.

6 Fairness Analysis

The proportional share scheduling algorithm we have proposed executes clients at a precise rate. One can determine if a client has a desired real-time response time property by simply computing the amount of service time it is to receive during the interval(s) of time of interest using either Eq. (5) or (6). However, because service time is allocated in discrete quanta, this computation is off by the client's lag (the difference between the service time that the client should have received in the ideal system, and the service time that it actually receives in the real system). Thus, in order to use our proportional share algorithm for real-time computing, we must demonstrate that the lag incurred by any client is bounded at all times. This is done next.

The problem is stated as that of demonstrating that the EEVDF algorithm is *fair* in the sense that all clients make progress according to their weights. By demonstrating that the lag is of client is bounded at all times we conclude that our algorithm is fair. Here we sketch the argument that lags are bounded. The complete proof of each result in given in the Appendix.

Theorem 1 shows that any request is fulfilled no latter than q time units after its deadline in the corresponding fluid-flow system, where q represents the maximum size of a time quantum. Theorem 2 gives tight bounds for the lag of any client in a system in which all the clients that join and leave the competition have zero lags. Similarly, Theorem 3 gives tight bounds for the client's lag in a system in which a client with positive lag may leave at any time (recall from the previous section that the case of a client with negative lag leaving the competition can be easily reduced to the case when a client with positive lag leaves the competition). Finally, as a corollary we show that in a dynamic system in which no client request is larger than the maximum size q of a time quantum the lag of any client is bounded by q. Moreover, this result is optimal with respect to *any* proportional share algorithm. We begin by defining formally the systems we are analyzing.

Definition 1 *A* **steady** *system (***S-system** *for short) is a system in which the lag of any client that joins, or leaves the competition is zero.*

The next definition is a formal characterization of the system described in Section 5 (see Figure 2(b)).

Definition 2 *A* **pseudo-steady** *system (***PS-system** *for short) is a system in which the lag of any client that joins is zero, and the lag of any client that leaves is positive. Moreover, when a client with positive lag leaves, the value of virtual-time is updated according to Eq. (14).*

The following theorem gives the upper bound for the maximum delay of fulfilling a request in an S-system. We note that this result generalizes a previous result of Parekh and Gallager [9] which holds for the particular case in which a request is no larger than a time quantum.

Theorem 1 *Let d be the deadline of the current request issued by client k in an S-system with quantum q, and let f be the actual time when this request is fulfilled. Then*

1) the request is fulfilled no later than $d + q$, i.e., $f < d + q$, and

2) if $f > d$, for any time $t \in [d, f)$, $lag_k(t) < q$.

The next theorem gives tight bounds for a client lag in an S-system.

Theorem 2 *Let r be the size of the current request issued by client k in an S-system with quantum q. Then the lag of the client k at any time t while the request is pending is bounded as follows*

$$-r < lag_k(t) < max(r, q),$$

Moreover, these bounds are asymptotically tight.

Notice that the bounds given by Theorem 2 apply independently to each client and depend only on the length of their requests. While shorter requests offer a better allocation accuracy, the longer ones reduce the system overhead since for the same total service time fewer requests need to be generated. It is therefore possible to trade between the accuracy and the system overhead, depending on the client requirements. For example, for an intensive computation task it would be acceptable to take the length of the request to be in the order of seconds. On the other hand, in the case of a multimedia application we need to take the length of a request no greater than several tens of milliseconds, due to the delay constraints. Theorem 2 shows that EEVDF can accommodate clients with different requirements, while guaranteeing tight bounds for the lag of each client, which are independent on the other clients. As the next theorem shows this is not true for PS-systems. In this case the lag of a client can be as large as the maximum request issued by *any* client in the system.

Theorem 3 *Let r be the size of the current request issued by client k in an PS-system with quantum q. Then the lag of the client k at any time t while the request is pending is bounded as follows*

$$-r < lag_k(t) < max(R_{max}, q),$$

where R_{max} represents the maximum duration of any request issued by any client in the system. Moreover, these bounds are asymptotically tight.

The following corollary follows directly from Theorems 2 and 3.

Corollary *If no request of client k is larger than a time quantum, then at any time t its lag is bounded as follows:*

$$-q < lag_k(t) < q.$$

Finally, we note that according to the following simple lemma (the proof can be found in [12]) the bounds given in the above corollary are optimal, i.e., they hold for *any* proportional share algorithm.

Lemma *Given any system with time quanta of size q and any proportional share algorithm, the lag of any client is asymptotically bounded by $-q$ and q.*

7 Experimental Results

As a proof of concept we have implemented a CPU scheduler prototype based on our EEVDF algorithm under FreeBSD v 2.0.5. All the experiments were run on a PC compatible with a 75 MhZ Pentium and 16 MB of RAM. The scheduler time slice (quantum), and the duration of any client's request were set to 10 ms.

Excepting the CPU scheduler, we did not alter the FreeBSD kernel.[9] Our scheduler coexists with the

[9] Indeed, the fact that we could perform these experiments on top of a largely unmodified general purpose operating system indicates the good fit between proportional share resource allocation scheme we advocate and general purpose operating system design.

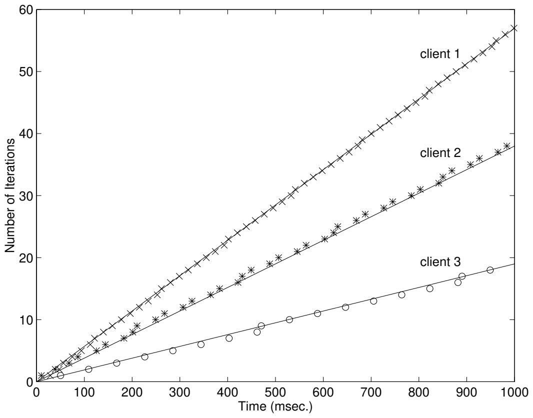

(a)

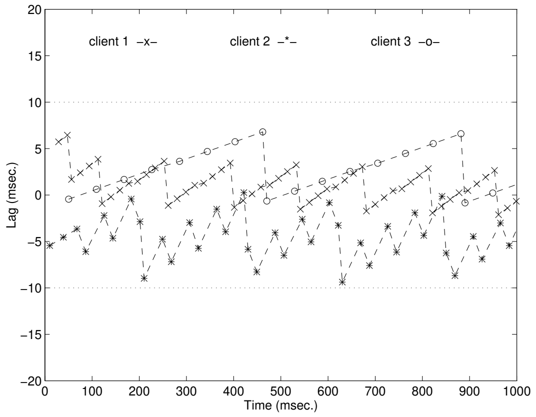

(b)

(a) plots the number of iterations over one second period for three clients with weights 3, 2, and 1, respectively; (b) plots the clients' lags over the same interval.

original FreeBSD scheduler [6]. All the processes that request proportional share or reservation services are assigned a reserved (user-level) priority, and are handled by our scheduler. All the other processes are scheduled by the regular FreeBSD scheduler. In this way, the kernel processes are scheduled over any process in the proportional share or the reservation class. Since FreeBSD lacks real-time support, such as preemptive kernel or page pinning, in our experiments we tried to avoid as much as possible the interaction with the kernel. For example, we make sure that all the measurements are performed after the entire program is loaded into memory[10]. Also, with the exception of `gettimeofday` function used for time measurements, we do not use any other system function while running the experiments (all the data are recorded in memory and saved at the end of each experiment). Addressing these issues in a first-class manner would only serve to improve our (already good) results.

To measure the allocation accuracy we have written a simple iterative applications which performs some arithmetic computations. Each iteration takes close to 9 ms. In each experiment we run several copies of the program, by assigning to each copy (client) a different weight.

In the first experiment we run three clients (processes), with weights 3, 2, and 1, respectively. All clients are synchronized via shared-memory to start the actual computation at the same time. Figure 3(a) shows the number of iterations executed by each client during the first second. The solid lines depict the ideal number of iterations for each client. As it can be seen, the number of iterations actually performed by each client is very closed to the ideal one. In particular, note that the client with a weight of 1 executes 1/2 the number of iterations as the client with a weight of 2 and at 1/3 the rate as the client with weight 3. Figure 3(b) depicts the lag of each client over the same interval. Note that the lags are always between -10 and 10 ms, which is consistent with the bounds given by the corollary in Section 6. Thus, each client is executing at a rate which is precise enough to afford one the ability to predict its performance in real-time (modulo 10 ms.) over any interval.

For example, assume client 3 in the first experiment is an audio player that receives and plays 1000 audio samples per minute from a remote site. Moreover, assume the computational cost of playing an audio sample is 10 ms. In a traditional real-time system this client would be modeled as a periodic task with a period p of 60 ms. and a cost c of 10 ms. However, in a proportional share system this client is a generic process that repeatedly makes requests for execution to play the next audio sample; each request having a duration of 10 ms. Since client 3 needs to execute for 10 ms. every 60 ms. it needs to receive a share of the processor equal to $10/60 = 1/6$ of the processor's total capacity. Thus, if clients 1 and 2 are the only other processes in the system, with weights 3 and 2 respectively, then to ensure client 3 executes properly we

[10]We achieve this by executing the entire code and touching all the data once *before* the actual measurements are performed (obviously, here we assume that the applications are small enough to fit all in the memory at the same time).

assign it a weight w such that its share of the processor is $f = \frac{w}{3+2+w} = \frac{1}{6}$. In this case, a weight of $w = 1$ will ensure client 3 receives the appropriate share of the processor and thus executes with the appropriate period, since $\frac{c}{f} = \frac{10}{1/6} = 60$ ms. That is, when client 3 receives share f of the processor, each request for execution will be completed $\frac{c}{f}$ time units after it is made. Moreover, at all times client 3 is guaranteed to have received within $q = 10$ ms of the service time it has requested. Thus, since our quanta in this example is the same as the duration of a client 3 request for execution, at all times t client 3 is guaranteed to have processed $\frac{t}{60} \pm 1$ audio samples. That is, at all times, client 3's real-time performance as measured in terms of the number of audio samples played, is within one sample of its ideal real-time execution. This difference can be changed as desired by adjusting the quanta with which the shared resource is allocated.

This example illustrates the power of a proportional share resource allocation system. In a proportional share system, one can either adjust the weight of a client as needed when new clients enter and leave the competition to achieve constant real-time performance, or leave a client's weight constant and allow a client's performance to proportionally improve and/or degrade as clients come and go. (For a complete solution see [13].) In the former case, one can ensure constant real-time performance so long as the set of active clients requiring constant performance do not cumulatively request shares of the processor greater than 1.0 (i.e., so long as these clients do not over utilize the processor). In the latter case, graceful (uniform, predictable) degradation/improvement is possible across all system loads. Moreover, unlike periodic task scheduling systems, since allocation is done in terms of discrete quanta, if a client executes for "too long," it only slows down itself. No other clients are effected. Thus if a process "misbehaves" it does not effect the rate of progress of any other process. Hence in addition to its other benefits, proportional resource allocation systems provide a natural means of "policing" tasks.

In the second experiment we consider again three clients with weights 3, 2, and 1, respectively, but in a more "dynamic" scenario. While clients 2 and 3 begin execution at the same time, client 1 is delayed for 300 ms. Each client performs 50 iterations after which it leaves the competition. As shown in Figure 4(a) there are four distinct regions. In the first region (i.e., between 0 and 300 ms) there are only two active clients: 2 and 3. Therefore client 2 (having weight 2) receives 66% percent of the CPU, while client 3 (having weight 1) receives 33% from the CPU. Consequently, after 300 ms client 2 completes 22 iterations, while client 3 completes only 11 iterations. After 300 ms, client 1 joins the competition and therefore in the second region (between 300 and 998 ms) all three clients are active. Further, at time $t = 998$ ms client 2 finishes all its iterations and leaves the competition. Thus, in the next region only clients 1 and 3 remains active. Finally, at time $t = 1128$ ms, client 1 finishes, and client 3 remain the only one active (in region four). Figure 4(b) depicts the clients lags, which are again between the theoretical bounds, i.e, -10 and 10 ms. However, notice that the sum of the lags is not always zero, as Lemma 2 implies (see Appendix). This is mainly because the elapsed times measured at the kernel level, and which are actually used in scheduling, are not equal to the times measured at the process level. One reason is that at the kernel level the elapsed times include also the context switching and other various scheduling related overheads. A second reason is the inaccuracy in measurements due using a system call.

8 Conclusions

We have described a new proportional share resource allocation scheduler that provides a flexible control, and provides strong timeliness guarantees for the service time received by a client. In this way we provide a unified approach for scheduling "firm" real-time, interactive, and batch applications. We achieve this by uniformly converting the application requirements regardless of their type in a sequence of requests for the resource. Our algorithm guarantees that the difference between the service time that a client should receive in the idealized system and the service time it actually receives in the real system is bounded by one time quantum and that this bound is optimal. At our best knowledge, this is the first algorithm to achieve these bounds in a dynamic system that provides support for both fractional and non-uniform quanta. As a proof of concept we have also implemented a prototype of a CPU scheduler under the FreeBSD operating system. Our experimental results shows that our implementation performs within the theoretical bounds.

References

[1] S. K. Baruah, J. E. Gehrke and C. G. Plaxton, "Fast Scheduling of Periodic Tasks on Multiple Resources", *Proc. of the 9th International Parallel Processing Symposium*, April 1995, pp. 280–288.

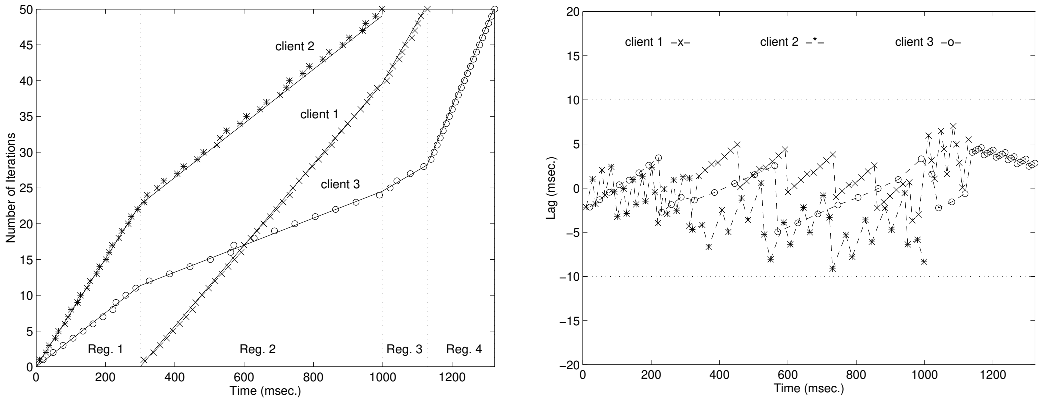

(a) (b)

(a) plots the number of iterations for three clients with weights 3, 2, and 1, respectively. Each client performs 50 iterations. Clients 2 and 3 begins at the same time; client 3 is delayed for 300 ms. (b) plots the clients' lags.

[2] S. K. Baruah, J. E. Gehrke and C. G. Plaxton, "Fair On-Line Scheduling of a Dynamic Set of Tasks on a Single Resource", *Technical Report* UTCS-TR-96-03, Department of Computer Science, University of Texas at Austin. February 1996.

[3] J. C. R. Bennett and H. Zhang, "WF^{2}Q : Worst-case Fair Queueing", INFOCOM'96, San-Francisco, March 1996.

[4] A. Demers, S. Keshav, and S. Shenkar, "Analysis and Simulation of a Fair Queueing Algorithm", *Journal of Internetworking Research & Experience*, October 1990, pp. 3–12.

[5] S. J. Golestani, "A Self-Clocked Fair Queueing Scheme for Broadband Applications", *Proc. of IEEE INFOCOM'94*, April 1994, pp. 636–646.

[6] S. J. Leffler, M. K. McKusick, M. J. Karels and J. S. Quarterman. "The Design and Implementation of the 4.3BSD UNIX Operating System," Addison-Wesley, 1989.

[7] C. L. Liu and J. W. Layland, "Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment", *Journal of the ACM*, Vol. 20, No. 1, January 1973, pp. 46–61.

[8] U. Maheshwari, "Charged-based Proportional Scheduling", Technical Memorandum MIT/LCS/TM-529, Laboratory for CS, MIT, July 1995.

[9] A. K. Parekh and R. G. Gallager, "A Generalized Processor Sharing Approach To Flow Control in Integrated Services Networks-The Single Node Case", *ACM/IEEE Transactions on Netwrking*, Vol. 1, No. 3, 1992, pp. 344–357.

[10] A. K. Parekh, "A Generalized Processor Sharing Approach To Flow Control in Integrated Services Networks", *Ph.D Thesis*, Department of EE and CS, MIT, 1992.

[11] I. Stoica, H. Abdel-Wahab, "A new approach to implement proportional share resource allocation", *Technical Report* TR-95-05, CS Dpt., ODU, April 1995, pp. 17 (also available at `http://www.cs.odu.edu/~stoica/pubs.html`)

[12] I. Stoica, H. Abdel-Wahab, "Earliest Eligible Virtual Deadline First: A Flexible and Accurate Mechanism for Proportional Share Resource Allocation", *Technical Report* TR-95-22, CS Dpt., ODU, November 1995, pp. 37 (also available at `http://www.cs.odu.edu/~stoica/pubs.html`)

[13] I. Stoica, H. Abdel-Wahab, K. Jeffay, "On the Duality between Resource Reservation and Proportional Share Resource Allocation", *Technical Report* TR-96-19, CS Dpt., ODU, May 1996, pp. 14 (also available at `http://www.cs.odu.edu/~stoica/pubs.html`)

[14] R. Tijdeman, "The Chairmain Assignment Problem", *Discrete Mathematics*, vol. 32, 1980, pp. 323–330.

[15] C. A. Waldspurger and W. E. Weihl. "Lottery Scheduling: Flexible Proportional-Share Resource Management," *Proc. of the First Symposium on Operating System Design and Implementation*, November 1994, pp. 1–12.

[16] C. A. Waldspurger and W. E. Weihl. "Stride Scheduling: Deterministic Proportional Share Resource Menagement," Technical Memorandum, MIT/LCS/TM-528, Laboratory for CS, MIT, July 1995.

[17] C. A. Waldspurger. "Lottery and Stride Scheduling: Flexible Proportional-Share Resource Management," *PhD Thesis*, Technical Report, MIT/LCS/TR-667, Laboratory for CS, MIT, September 1995.

[18] L. Zhang, "VirtualClock: A New Traffic Control Algorithm for Packet-Switched Networks", *ACM Transactions on Computer Systems*, vol. 9, no. 2, May 1991, pp. 101–124.

Appendix

In this appendix we prove the main results of the paper presented in Section 6. The following lemma gives the eligibility condition for a client request.

Lemma 1 *If an active client k has a positive lag at time t (i.e., $lag_k(t) \geq 0$), then its current pending request is eligible.*

Proof. Let r be the length of the request issued by client k that is pending at time t (recall that an active client has always a pending request), and let ve denote the virtual eligible time of this request. For the sake of contradiction, assume the request is not eligible at time t, i.e., $ve > V(t)$.
Let $t'(< t)$ be the time when the request was issued. Then from Eq. (7)

$$s_k(t_0^k, t') = w_k(ve - V(t_0^k)). \tag{16}$$

Since between t' and t the request was not eligible, it follows that the client has not received any service time during $[t', t)$, and thus $s_k(t_0^k, t') = s_k(t_0^k, t)$. By using Eq. (3), (6), and (16) we obtain

$$\begin{aligned} lag_k(t) &= (V(t) - V(t_0^k))w_k - s_k(t_0^k, t) = \\ &\quad (V(t) - V(t_0^k))w_k - (ve - V(t_0^k))w_k = \\ &\quad (V(t) - ve)w_k < 0, \end{aligned} \tag{17}$$

which, by contradiction, proves the lemma. □

From Lemma 1 and from the fact that any zero sum has at least a positive term, we have

Corollary 1 *Let $\mathcal{A}(t)$ be the set of all active clients at time t, such that*

$$\sum_{i \in \mathcal{A}(t)} lag_i(t) = 0. \tag{18}$$

Then there is at least one eligible request at time t.

The next lemma shows that at any time t the sum of the lags over all active clients is zero. An immediate consequence of this result and the above corollary is that at any time t at which there is at least one active client, there is also at least one pending eligible request in the system. Thus, in the sense of the definition given in [9], the EEVDF algorithm is *work-conserving*, i.e., as long as there are active clients the resource is busy.

Lemma 2 *At any moment of time t, the sum of the lags of all active clients is zero, i.e,*

$$\sum_{i \in \mathcal{A}(t)} lag_i(t) = 0. \tag{19}$$

Proof. Here we give only the outline of the proof; for details the reader is referred to [12]. The proof goes by induction. First, we assume that at time $t = 0$ there is no active client and therefore Eq. (19) is trivially true. Next, we show that Eq. (19) remains true after each of the following events occur: (i) a client joins the competition, (ii) a client leaves the competition, (iii) a client changes its weight. Finally, we show that (iv) during any interval $[t_1, t_2)$ in which none of the above events occurs if Eq. (19) holds at time t_1, then it also holds at time t_2. □

Lemma 3 *Let e be the eligible time of some request issued by client k. Then*

$$lag_k(e) \geq 0. \tag{20}$$

Proof. Let t be the time at which client k issues the request. First, assume the request is not eligible at time t, i.e., $t < e$. Since client k does not receive any service time between t and e we have $s_k(t_0^k, t) = s_k(t_0^k, e)$. On the other hand, if the request is eligible by time t, i.e., $t > e$, we have $s_k(t_0^k, t) \geq s_k(t_0^k, e)$. Finally, by using Eq. (3), (6) and (7) we obtain

$$\begin{aligned} lag_k(e) &= (V(e) - V(t_0^k))w_k - s_k(t_0^k, e) \geq \\ &\quad (V(e) - V(t_0^k))w_k - s_k(t_0^k, t) = 0. \end{aligned} \tag{21}$$

which completes the proof. □

Before giving the next theorem we note, from Eq. (3), that for any client active during an interval $[t_1, t_2]$ we have:

$$lag_k(t_2) = lag_k(t_1) + S_k(t_1, t_2) - s_k(t_1, t_2). \tag{22}$$

Theorem 1 *Let d be the deadline of the current request issued by client k in an S-system with quantum q, and let f be the actual time when this request is fulfilled. Then*
1) the request is fulfilled no later than $d + q$, i.e., $f < d + q$, and
2) if $f > d$, for any time $t \in [d, f)$, $lag_k(t) < q$.

Proof. Let e be the eligible time of the request (with deadline d) issued by client k. Assume the request is not fulfilled by time $T \geq d$ (otherwise, the theorem is trivially true). Then, let t_1 be the latest time no greater than T at which a client l with a deadline greater than d *starts* to receive a time quantum.
Case 1 (t_1 exists). As shown in Figure 5, consider the partition of all active clients at time t into two sets: $C(t)$ containing all the clients active at time t that have a request with a deadline *larger* than d (note that

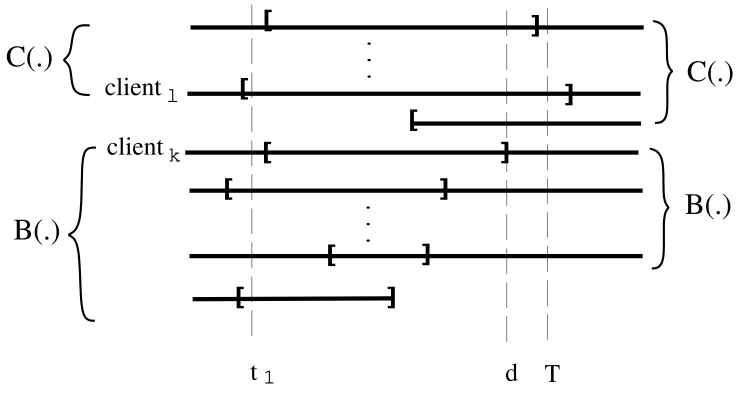

The current pending request of client k has the eligible time e and deadline d. Set $C(t)$ contains all clients active at time t that have deadlines larger than d, while set $B(t)$ contains all the other clients, including client k, that are active at time t. t_1 is the latest time $\leq d$, when a client in $C(t_1)$ receives service time.

$l \in C(t_1)$), and $B(t)$ containing all the other clients (with deadlines no larger than d), i.e., $B(t) = \mathcal{A}(t) \setminus C(t)$ (note that $k \in B(T)$).

Next we show that all the clients in $B(t_1)$ have negative lags. By contradiction assume there is a client $i \in B(t_1)$, such that $lag_i(t_1) \geq 0$. Then from Lemma 1 the pending request of client i is eligible at time t_1. But then client l cannot be scheduled at time t_1, since its request has a larger deadline than the deadline of the request of client i, which contradicts our assumption. Thus we have $\sum_{i \in B(t_1)} lag_i(t_1) < 0$, and since the sum of the lags of all active clients is zero (Lemma 2)

$$\sum_{i \in C(t_1)} lag_i(t_1) > 0. \tag{23}$$

Further, consider the following facts:
1. Excepting l, no other client with a deadline greater than d receives any service time between t_1 and T.
2. Since the maximum size of a time quantum is q, the service time received by client l during $[t_1, T]$ is at most q.
3. All the clients that leave or join the competition have zero lags.
4. $C(T)$ consists of all clients in $C(t_1)$ that are still active at time T, and all the clients which join the competition between $[t_1, T]$ by issuing a request with a deadline larger than d.
¿From (1) and (2) it follows that the *entire* service time received by all clients in $C(T)$ during $[t_1, T]$ is at most q. Further, from (3) and (4) above, from Eq. (23), and the fact that the lag of a client increases as long as it does not receive service time, it can be shown that $\sum_{i \in C(T)} lag_i(T) > -q$. Further, by using Lemma 2, we obtain

$$\sum_{i \in B(T)} lag_i(T) < q. \tag{24}$$

Since all the deadlines of the requests issued by clients in $B(T)$ are smaller than d (and consequently smaller than T), it follows that no request issued by any client in T is fulfilled before its deadline. Next, recall that if a request is not fulfilled by its deadline, then this means that it has received *less* service time than it was supposed to, and therefore its lag is positive. But then, at time T, all clients in $B(T)$ have positive lags, and therefore from Ineq. (24) we have

$$lag_k(T) < q, \tag{25}$$

which proves the second part of the theorem.

Next we prove the first part. From the deadline definition, recall that if the request of a client i is *not* fulfilled by its deadline d', then $lag_i(d')$ represents the *remaining* service time to fulfill that request. Further since the virtual time increases monotonically, while the remaining service time can only decrease, and since the deadline of any client in $B(T)$ is smaller than T, it follows that at time T the lag of any client in $B(T)$ is *larger* than the remaining service time required to fulfill the client's request.

Next, consider two sub-cases: $t_2 \leq d$, and $t_2 > d$. In the first sub-case, from Eq. (24), we have $\sum_{i \in B(d)} lag_i(d) < q$. Since no client with deadline larger than d can receive any time quantum before all the requests of clients in $B(d)$ are satisfied, and since at most q time units are necessary to satisfy all these requests, it follows that by time $d + q$, all these requests, including the one of client k, are satisfied. In the second sub-case note that since the pending request of any client i in $B(t_2)$ has an eligible time e' larger than t_1 (otherwise, client l cannot be scheduled at time t_1) and a deadline d' smaller than d, from the monocity of the virtual time and from Eq (8) the size r' of the request issued by client i is $r' = (V(d') - V(e'))w_i \leq (V(d) - V(t_1))w_i$. From here, it is easy to show that the sum of all the outstanding requests issued by the clients in $B(t_2)$ is no larger than $d - t_1$. Consequently, the requests of all clients in $B(t_2)$, including client k, are fulfilled no later than $d - t_1$ after time t_2. Since $t_2 - t_1 < q$ the proof follows.

Case 2 (t_1 does not exist). Then we take $t_1 = e$. Assume $C(e)$ is not empty (the case when $C(e)$ is empty reduces trivially to this case). Since no client in $C(e)$ receives any service time before e, we have $\sum_{i \in C} lag_i(e) \geq 0$, $(i \in C(e))$, and further

$$\sum_{i \in B(e)} lag_i(e) \leq 0. \qquad (26)$$

¿From here the proof proceeds similar to the previous case, the difference being that in this case we can actually show that the request issued by client k is *always* fulfilled before its deadline d. □

Theorem 2 *Let r be the size of the current request issued by client k in an S-system with quantum q. Then the lag of the client k at any time t while the request is pending is bounded as follows*

$$-r < lag_k(t) < max(r, q), \qquad (27)$$

Moreover, these bounds are asymptotically tight.

Proof. Let e and d be the eligible time and the deadline of an arbitrary request with duration r issued by client k. Since the client's lag increases as long as the client is not serviced, and decreases otherwise (see Eq. (3), (5) and (4)), and since a request is not serviced before it is eligible, it is easy to see that the lag is minimized when the client receives the entire service time as soon as the request becomes eligible. In other words, the minimum lag occurs at time $e + r$, if the request is fulfilled by that time. Then, for any t while the request is pending, by using Eq. (22) we have

$$\begin{aligned} lag_k(t) \;\; \geq \;\; & lag_k(e + r) = \qquad (28) \\ & lag_k(e) + S_k(e, e + r) - s_k(e, r) \geq \\ & -s_k(e, r) > -r. \end{aligned}$$

Since any client with positive lag is eligible (see Lemma 1), to prove the right-hand side of Ineq. (27) we consider only times larger or equal to e. Next, for any such a time $t > e$, while the request is pending, we consider two cases: $t < d$ and $t \geq d$. In the first case, from Eq. (22), and since $S_k(t_0^k, t) < S_k(t_0^k, d)$ and $s_k(t_0^k, t) \geq s_k(t_0^k, e)$, we have

$$lag_k(t) = S_k(t_0^k, t) - s_k(t_0^k, t) < S_k(t_0^k, d) - s_k(t_0^k, e) \leq r. \qquad (29)$$

In the second case, from Theorem 1 we have directly $lag_k(t) < q$. Finally, by combining this inequality, with Ineq. (28) and (29), we obtain (27).

To show that the bound $lag_k(t) > -r_1$ is asymptotically tight, consider the following example. Let w_1, w_2 be the weights of two active clients, such that $w_1 \ll w_2$. Next, suppose that both clients become active at time t_0 and their first requests have the lengths r_1 and r_2, respectively. We assume that r_1 and r_2 are chosen such that the virtual deadline of the first client's request is smaller than the virtual deadline of the second client's request, i.e., $t_0 + \frac{r_1}{w_1} < t_0 + \frac{r_2}{w_2}$. Then client 1 receives the entire service time before client 2, and thus from Eq. (3) we have $lag_1(r_1) = S_1(t_0, t_0 + r_1) - r_1$. Next, by using Eq. (5) we obtain $S_1(t_0, t_0 + r_1) = \frac{w_1}{w_1 + w_2}$, which approaches zero when $\frac{w_1}{w_2} \to \infty$, and consequently $lag_1(r_1)$ approaches $-r_1$. The proof for $lag_k(t) < max(r_1, q)$ is similar; therefore we omit it here. □

Theorem 3 *Let r be the size of the current request issued by client k in an PS-system with quantum q. Then the lag of the client k at any time t while the request is pending is bounded as follows*

$$-r < lag_k(t) < max(R_{max}, q),$$

where R_{max} represents the maximum duration of any request issued by any client in the system. Moreover, these bounds are asymptotically tight.

Proof. The basic idea is to reduce any PS-system A to an "equivalent" S-system B. Intuitively, system B is derived from A by simply "stretching" the time axis whenever a client with a strict positive lag leaves the competition in system A (for an example see Figures 2(b) and 2(a)). More formally, system B is defined as follows.

1. For any time t in A there is a time $t' = f(t)$ in B, such that:

$$f(t) = t + \sum_{i \in L(t)} \Delta_i, \qquad (30)$$

where $L(t)$ denotes all the clients with positive lags that leave the competition before t in system A, and Δ_i represents the service time that client $i \in L(t)$ receives in B before leaving the competition (such that its lag to become zero).

2. If a client with zero lag joins/leaves the competition at time t in A, the same client joins/leaves at time $f(t)$ in B.

3. If a client k with a strict positive lag leaves at time t in A, then it leaves at time $f(t) + \Delta_k$ in B, i.e., when its lag becomes zero.

Note that in this reduction we implicitly use the restriction (2) from Section 6, i.e., the restriction that a client leaves the competition only when it holds the shared resource. Next, we show that the lag of any *active* client in A at time t is equal to the lag of the same client in B at time $f(t)$. This results easily by observing that: (1) the virtual time at time t in system A is equal to the virtual time at time $f(t)$ in system B, and (2) the service time that a client *active* at time t receives before t in A is equal to the service time it receives before $f(t)$ in B. (However, note that *every* client k that leaves with a positive lag in A, receives an extra Δ_k service time in B.) But then from Eq. (3) the two lags are equal.

Thus we have reduced the problem of computing the lags' bounds for system A to computing these bounds for system B. Next, consider a request with the virtual eligible time e and the virtual deadline d issued by some client k in system B. Then, we show that if the request is not fulfilled by time T $(\geq d)$, then we have

$$lag_k(t) \leq R_{max}. \tag{31}$$

The proof is very similar to the proof for the *second* part of Theorem 1 (in the followings we use the same notations as in Theorem 1). The only difference is that in case 1 we have also to consider the sub-case where the latest client l, with deadline larger than d that receives service time starting with t_1, is a client with a *strict positive* lag that wants to leave the competition. First, notice that since time t_1 cannot be larger than d (otherwise, client l cannot be scheduled at time t_1 since the request of client k has an earlier deadline), it follows that t_1 is also smaller than the deadline of client l. From here we have $lag_l(t_1) < r \leq R_{max}$. Since the lag of client l cannot decrease by more than $lag_k(t_1)$ by time T (because once it becomes zero client l leaves the competition), and since no other client in $C(T)$ receives any service time between t_1 and T, by using Eq. (23) we can show that $\sum_{i \in C(T)} lag_i(T) > lag_l(t_1) > -R_{max}$. Further, from Lemma 2, we have $\sum_{i \in B(T)} lag_i(T) < R_{max}$. Finally, since all the clients in $B(T)$ have positive lags at time T we obtain Eq. (31). Further the proof is similar to the one of Theorem 3; therefore we omit it here. $\square$