Analysis of a Window-Constrained Scheduler for Real-Time and Best-Effort Packet Streams

Richard West and Christian Poellabauer

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332

Abstract

This paper describes how Dynamic Window-Constrained Scheduling (DWCS) can guarantee real-time service to packets from multiple streams with different performance objectives. We show that: (1) DWCS can guarantee that no more than x packets miss their deadlines for every y consecutive packets requiring service, (2) using DWCS, the delay of service to real-time packet streams is bounded even when the scheduler is overloaded, (3) DWCS can ensure the delay bound of any given stream is independent of other streams, and (4) a fast response time for best-effort packet streams, in the presence of real-time packet streams, is possible. As long as the *minimum* aggregate bandwidth requirement of all real-time packet streams does not exceed the available bandwidth, DWCS can guarantee that each such stream does not miss more than x deadlines for every y requests. Furthermore, if a feasible schedule exists, each stream is guaranteed a minimum fraction of available bandwidth over a finite window of time. Consequently, DWCS can provide bounded delay of service to each real-time stream in a manner which is independent of other streams, while also meeting per-stream, explicit delay and loss constraints.

1. Introduction

Low latency, high bandwidth integrated services networks have introduced opportunities for new applications such as video conferencing, tele-medicine, virtual environments[4, 15], groupware[11], distributed interactive simulations (DIS)[20] and other real-time multimedia applications. Many of these applications require strict performance (or *quality of service*) requirements on the information transferred across a network. Typically, these performance objectives are expressed as some function of throughput, delay, jitter and loss-rate[8]. With many multimedia applications, such as video-on-demand or streamed audio, it is important that information is received and processed at an almost constant rate, such as 30 frames per second for video information. However, some packets comprising a video frame or audio sample can be lost, resulting in little or no noticeable degradation in the quality of playback at the receiver. Similarly, a data source can lose a certain fraction of information during its transfer across a network as long as the receiver processes the received data to compensate for the loss. Consequently, loss-rate is an important performance measure for this category of applications. We define the term *loss-rate*[19] as the fraction of packets in a stream [1] either received later than allowed or not received at all at the destination.

One of the problems with using loss-rate as a performance metric is that it does not describe when losses are allowed to occur. For most loss-tolerant applications, there is usually a restriction on the number of *consecutive* packet losses that are acceptable. For example, losing a series of consecutive packets from an audio stream might result in the loss of a complete section of audio, rather than merely a reduction in the signal-to-noise ratio. A suitable performance measure in this case is a *windowed loss-rate*, i.e. loss-rate constrained over a finite range, or *window*, of consecutive packets. More precisely, an application might tolerate x packet losses for every y arrivals at the various service points across a network. Any service discipline attempting to meet these requirements must ensure that the number of violations to the loss-tolerance specification is minimized (if not zero) across the whole stream.

In contrast to loss-constrained applications, computer data transferred between hosts using a file-transfer protocol cannot tolerate any loss at all. In this case, a more appropriate performance measure is mean delay, to ensure that the delay incurred by packets from this class of applications is minimized.

In summary, integrated services networks must be able to support diverse performance objectives. Therefore, a suitable service discipline at the network access points (and, possibly, switches) must be able to schedule the transmission of packets from various streams so that the objectives of as many of the most important packets as possible are met.

This paper describes the real-time properties of Dynamic Window-Constrained Scheduling (DWCS)[23, 24], an algorithm which is suitable for packet scheduling in real-time communications. DWCS is designed to explicitly service packet streams in accordance with their loss and delay constraints, using just two attributes per packet stream. Furthermore, DWCS has the desirable property of supporting 'fair' allocation of bandwidth to packet streams, in proportion to their loss-constraints and per-packet deadlines[23]. In fact, DWCS can behave as a static-priority (SP), earliest-deadline-first (EDF), and fair scheduling algorithm[7, 25, 9, 2, 10, 18, 22]. Moreover, DWCS is intended to support multimedia traffic streams in the same manner as the SMART scheduler[17], but DWCS is less complex and requires maintenance of less state information than SMART.

[1] Corresponding to each stream is the notion of a flow, session, or logical connection.

Although DWCS has a lot in common with fair scheduling algorithms, it is more closely related to algorithms which attempt to guarantee service to packet streams, whereby at least m out of k packet deadlines are met for each and every stream. Hamdaoui and Ramanathan[12] were the first to introduce the notion of (m, k)-firm deadlines, which is similar to the concept of 'Skip-Over' by Koren and Shasha[14]. However, in some cases, skip over algorithms unnecessarily skip service to one or more activities (such as periodic tasks or packet streams), even if it is possible to meet the deadlines of those activities.

The (m, k)-firm deadline algorithm of Hamdaoui and Ramanathan guarantees that a statistical number of deadlines will be met, by using a 'distance-based' priority scheme to increase the priority of an activity in danger of missing more than m deadlines over a window of k requests for service. By contrast, Bernat and Burns[3] schedule activities with (m, k)-*hard* deadlines, but their approach requires such hard temporal constraints to be guaranteed by off-line feasibility tests. Moreover, Bernat and Burns work focuses less on the issue of providing a solution to on-line scheduling of activities with (m, k)-hard deadlines, but more on the support for fast response time to best-effort activities, in the presence of activities with hard deadline constraints.

Pinwheel scheduling[13, 5, 1] is also similar to DWCS. In essence, the generalized pinwheel scheduling problem is equivalent to determining a schedule for a set of n activities $\{a_i \mid 1 \leq i \leq n\}$, each requiring at least m_i deadlines are met in *any* window of k_i deadlines, given that the time between consecutive deadlines is a multiple of some fixed-size time slot, and resources are allocated at the granularity of one time slot. DWCS can be thought of as a special case of pinwheel scheduling, whereby DWCS guarantees a minimum of m_i deadlines are met every *fixed* (i.e., non-overlapping) window of k_i deadlines, for a given activity a_i. In fact, DWCS is capable of producing a feasible schedule, *independent of an activity's window size* k_i, when 100% of available resources (such as bandwidth) are utilized. By comparison, Baruah and Lin[1] have developed a pinwheel scheduling algorithm, that is capable of producing a feasible schedule when the utilization of resources approaches 100%, given that $k \to \infty$.

In contrast to the related work described above, the significant contributions of this work are: (1) the description and analysis of an *on-line* version of DWCS that can guarantee (m, k)-hard deadlines (or, equivalently, no more than x missed packet deadlines for every fixed window of y consecutive packets in a given stream), (2) an approach to ensure fast response time to best-effort packet streams in the presence of real-time packet streams, and (3) a proof that DWCS ensures the delay of service to packets in any given stream is bounded, even in overload situations. In fact, DWCS can ensure the delay bound of any given stream is independent of other streams.

Using DWCS in an end-host QoS architecture, or in network switches, can guarantee service to packet streams with loss and delay constraints. In fact, assuming all packets in each and every stream have the same *worst-case* service time requirements, and each stream tolerates no more than x missed packet deadlines every y requests, then if the *minimum* aggregate bandwidth requirement of all real-time packet streams does not exceed the available bandwidth, DWCS can guarantee to meet the service constraints of all such packet streams. Observe that the service time of a packet is a function of its length (in bits) and service rate (in bits per second), due to server capacity, or link bandwidth (whichever is limiting). If we assume the scheduler has the capacity to process packets fast enough to saturate a network link, and link bandwidth is constant, then all packets will have the same service time if they have the same length. However, if packets vary in length, or if the server capacity fluctuates (either due to variations in link bandwidth, or variations in the service rate due to scheduling latencies associated with supporting different numbers of streams), then packet service times can be variable. In such circumstances, if it is possible to impose an upper bound on the worst-case service time of each and every packet, then DWCS can still guarantee that no more than x packet deadlines are missed every y requests. Note that, for these service guarantees to be made with DWCS, resources are allocated at the granularity of one *time slot* (see Figure 1), where the size of a time slot is typically determined by the (worst-case) service time of the largest packet in any stream requiring service. Therefore, it is assumed that when scheduling packets from a given stream, at least one packet in a stream is serviced in a time slot, and no other packet (or packets) from any other stream can be serviced until the start of the next time slot. Unless stated otherwise, we assume throughout this paper that at most one packet from any given stream is serviced in a single time slot but, in general, it is possible for multiple packets from the same stream to be aggregated together and serviced in a single time slot, as if they were one large packet.

The remainder of this paper is organized as follows: Section 2 precisely defines the scheduling problem, and then describes a version of Dynamic Window-Constrained Scheduling that guarantees real-time service constraints. Section 3 analyzes the performance of DWCS, including the bounds on service delay for competing packet streams, and constraints under which real-time service guarantees can be made. Section 4 describes an approach to guarantee fast response time to best-effort packet streams in the presence of packet streams with real-time constraints. Finally, conclusions and future work are described in Section 5.

2. Dynamic Window-Constrained Scheduling

We begin this section by defining the problem of guaranteeing a feasible schedule for real-time packet streams, which can tolerate at most x missed deadlines every fixed window of y requests. We then describe how the DWCS algorithm works, so that it can produce a feasible schedule on-line.

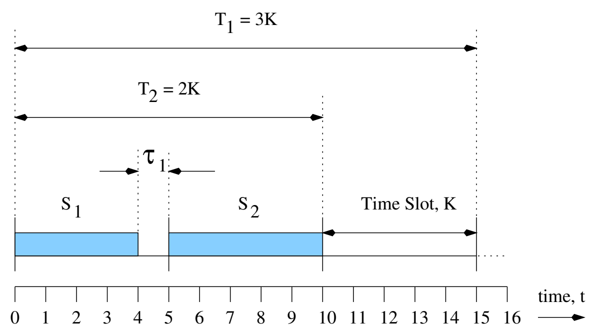

Figure 1. Example of two packets from different streams, S_1 and S_2 being serviced in their respective time slots. Each time slot is of constant size K. Observe that the packet in S_1 requires $K - \tau_1$ service time, thereby wasting τ_1 time units before the packet in S_2 is serviced. In this example, S_1 has a request period of 3 time slots, while S_2 has a request period of 2 time slots.

2.1. Problem Definition

In order to define the real-time scheduling problem addressed as part of this paper, we introduce the following definitions:

Bandwidth Utilization. This is a measure of the fraction (or percentage) of available bandwidth used by packet streams to meet their service constraints. A series of packet streams is said to *fully utilize*[16] available bandwidth, B, if all packet streams using B satisfy their service constraints, and any increase in the use of B violates the service constraints of one or more packet streams.

Dynamic Window-Constrained Scheduling (DWCS). DWCS is an algorithm for scheduling packet streams, each having the following pair of service attributes, which are used to define each stream's delay and loss constraints:

- *Request Period* – A request period, T_i, for a packet stream, S_i, is the interval between the deadlines of consecutive pairs of packets in S_i. Observe that the end of a request period, T_i, determines a *deadline* by which a packet in stream S_i must be serviced.
- *Window-Constraint* – this is specified as a value x_i/y_i, where the window-numerator, x_i, is the number of packets that can be lost or transmitted late for every fixed *window*, y_i (the window-denominator), of consecutive packet arrivals in the same stream, S_i. Hence, for every y_i packet arrivals in stream S_i, a minimum of $y_i - x_i$ packets must be scheduled for service by their deadlines. At any time, all packets in the same stream, S_i, have the same window-constraint, W_i, while each successive packet in a stream, S_i, has a deadline that is offset by a fixed amount, T_i, from its predecessor. After servicing a packet from S_i, the scheduler adjusts the window-constraint of S_i

and all other streams whose head packets have just missed their deadlines due to servicing S_i. Consequently, a stream S_i's *original* window-constraint, W_i, can differ from its *current* window-constraint, W_i'. Observe that a stream's window-constraint can also be thought of as a *loss-tolerance*.

Feasibility. A schedule, comprising a sequence of packet streams, is feasible if no original window-constraint of any packet stream is ever violated. DWCS attempts to schedule all packet streams to meet as many window-constraints as possible.

Problem Statement. The problem is to produce a feasible schedule using an on-line algorithm. The algorithm should attempt to maximize network bandwidth. In fact, we show in Section 3 that, under certain conditions, Dynamic Window-Constrained Scheduling can guarantee a feasible schedule as long as the *minimum* aggregate bandwidth utilization of a set of packet streams does not exceed 100% of available bandwidth. Before we analyze DWCS, we first describe the algorithm in more detail.

2.2. The DWCS Algorithm

This section describes a revised version of Dynamic Window-Constrained Scheduling (DWCS) from that described in [23, 24], so that hard real-time guarantees can be made on packet streams tolerating x missed deadlines every y requests. This version does not replace the original algorithm but shows how, by working in a *non-work-conserving* mode, deterministic service guarantees are possible for more real-time streams. The original algorithm is a work-conserving algorithm and only guarantees a statistical number of real-time service constraints.

DWCS orders packets for service based on the values of their *current* window-constraints and deadlines (where each deadline is derived from the current time and the request period). Precedence is given to packets in streams according to the rules shown in Table 1. Whenever a packet in S_i misses its deadline, the window-constraint for all subsequent packets in S_i is adjusted to reflect the increased importance of servicing S_i. This approach avoids starving the service granted to a given packet stream, and attempts to increase the importance of servicing any stream likely to violate its original window-constraint. Conversely, any packet in a stream serviced before its deadline causes the window-constraint of any subsequent packets in the same stream (yet to be serviced) to be increased, thereby reducing their priority.

The window-constraint of a packet stream changes over time, depending on whether or not another (earlier) packet from the same stream has been serviced by its deadline. If a packet cannot be serviced by its deadline, it is either transmitted late or it is dropped and the next packet in the stream is assigned a deadline corresponding to the latest time it must complete service.

Table 1 shows the rules for ordering pairs of packets. It differs slightly from the precedence rules used

Pairwise Packet Ordering
Earliest deadline first (EDF)
Equal deadlines, order lowest window-constraint first
Equal deadlines and zero window-constraints, order highest window-denominator first
Equal deadlines and equal non-zero window-constraints, order lowest window-numerator first
All other cases: first-come-first-serve

Table 1. Precedence amongst pairs of packets in different streams

in the original design of DWCS[23]. It should be noted that DWCS is more than merely an earliest deadline first (or earliest due date) algorithm. Observe that earliest deadline first scheduling (EDF) considers that each packet's importance (or priority) increases as the urgency of completing that packet's service increases. By contrast, static priority algorithms all consider that one packet is more important to service than another packet, based solely on each packet's time-invariant priority. DWCS combines both the properties of static priority and earliest deadline first scheduling by considering each packet's individual importance when the urgency of servicing two or more packets is the same. That is, if two packets have the same deadline, DWCS services the packet which is more important. In practice it makes sense to set packet deadlines in different streams to be some multiple of a, possibly worst-case, packet service time. This increases the likelihood of multiple head packets of different streams having the same deadlines. In fact, using a slotted time system, as described earlier, deadlines can be aligned on time slot boundaries. Observe that packets are ordinarily serviced in earliest deadline first order. However, if at least two packet streams have head packets with equal deadlines, the packet from stream S_i with the lowest *current* window-constraint W_i' is serviced first. If $W_i' = W_j' > 0$, and $d_i = d_j$ for S_i and S_j, respectively, S_i and S_j are ordered such that a packet from the stream with the lowest window-numerator is serviced first. By ordering based on the lowest window-numerator, precedence is given to the packet with *tighter* window-constraints, since fewer consecutive late or lost packets from the same stream can be tolerated. Likewise, if two packet streams have zero window-constraints and equal deadlines, the packet in the stream with the highest window-denominator is serviced first. All other situations are serviced in a first-come-first-serve manner.

We now describe how a stream's window-constraints are adjusted. As part of this approach, a *tag* is associated with each stream S_i, to denote whether or not S_i has violated its window-constraint W_i at the current point in time. In what follows, let S_i's *original* window-constraint be $W_i = x_i/y_i$, where x_i is the original window-numerator and y_i is the original denominator. Likewise, let $W_i' = x_i'/y_i'$ denote the *current* window-constraint. Before a packet in S_i is serviced, $W_i' = W_i$. Upon servicing a packet in S_i

before its deadline, W'_i is adjusted for subsequent packets in S_i, as follows:

(A) Window-constraint adjustment for a packet in S_i serviced before its deadline:

if $(y'_i > x'_i)$ then $y'_i = y'_i - 1$;
else if $(y'_i = x'_i)$ and $(x'_i > 0)$ then $x'_i = x'_i - 1; y'_i = y'_i - 1$;
if $(x'_i = y'_i = 0)$ or (S_i is tagged) then $x'_i = x_i; y'_i = y_i$;
if (S_i is tagged) then reset tag;

At this point in time, the window-constraint, W_j, of any other packet stream, $S_j \mid j \neq i$, comprising one or more late packets, is adjusted as follows:

(B) Window-constraint adjustment when a packet in $S_j \mid j \neq i$ misses its deadline:

if $(x'_j > 0)$ then
  $x'_j = x'_j - 1; y'_j = y'_j - 1$;
  if $(x'_j = y'_j = 0)$ then $x'_j = x_j; y'_j = y_j$;
else if $(x'_j = 0)$ and $(y_j > 0)$ then
  $y'_j = y'_j + \epsilon$;
  Tag S_j with a violation;

Observe that with DWCS, window-constraints do not change for packet streams whose packets do not have deadlines. Packet streams comprising packets without deadlines are *non-time-constrained*, and their window-constraints act as static priorities. Consequently, the pseudo-code for DWCS is shown in Figure 2.

Usually, a packet stream is eligible for service if a packet in that stream has not yet been serviced in the current request-period, which is the time between the deadline of the previous packet and the deadline of the current packet in the same stream. That is, no more than one packet in a given stream can be serviced in a given request period, and exactly one packet must be serviced by the end of its request period to prevent a deadline being missed. However, DWCS allows packet streams to be marked as eligible for scheduling multiple times in the same request period. This ensures DWCS is *work-conserving*, in that the jth packet, $p_{i,j}$ in a stream, S_i, can be serviced before the deadline of a prior packet, $p_{i,j-1}$ in the same stream, if $p_{i,j-1}$ has been serviced before the end of its request period. This implies $p_{i,j-1}$ is serviced before its deadline. For the purposes of *real-time*, as opposed to best-effort streams, this paper assumes DWCS works as a non-work-conserving scheduler. However, all best-effort streams can be serviced whenever there is available time to service such streams.

In the absence of a feasibility test, it is possible that window-constraint violations can occur. A violation actually occurs when $W'_i = x'_i/y'_i \mid x'_i = 0$ and another packet in S_i then misses its deadline. Before S_i is serviced, x'_i remains zero, while y'_i is increased by a constant, ϵ, every time a packet in S_i misses a deadline. The exception to this rule is when $y_i = 0$ (and, more specifically, $W_i = 0/0$). This special case allows DWCS to *always* service packet streams in EDF order, if such a service policy is desired.

```
Let Si = Stream i
    di = deadline of the next packet in Si
    Ti = request period of Si
    Wi'= current window-constraint of Si

while (TRUE) {
  for (each packet in all streams eligible for service at the current time)
    find the next packet in stream, Si, with the highest priority,
    according to the rules in Table 1;
  service next packet in Si;
  adjust Wi' according to rules in (A);
  /* Adjust deadline of next packet in Si. */
  di = di + Ti;
  for (each packet in Sj, other than Si, missing its deadline) {
    while (deadline missed) {
      adjust Wj' according to rules in (B);
      if (current packet can be dropped) {
        drop current packet in Sj;
        /* Adjust deadline of next packet in Sj, by adding
           Tj to the previous packet's deadline. */
        dj = dj + Tj;
      }
      else {
        /* Adjust deadline of current packet in Sj
           by adding Tj to the current deadline. */
        dj = dj + Tj;
      }
    }
  }
}
```

Figure 2. The DWCS algorithm.

If S_i violates its original window-constraint, it is tagged for when it is next serviced. Tagging ensures that a stream is never starved of service even in overload. Theorem 2 shows the delay bound for a stream which is tagged with window-constraint violations. Consequently, S_i is assured of service, since it will eventually take precedence over all packet streams with a zero-valued current window-constraint. Consider the case when S_i and S_j both have current window-constraints, W_i' and W_j', respectively, such that $W_i' = 0/y_i'$ and $W_j' = 0/y_j'$. Even if both deadlines, d_i and d_j, are equal, precedence is given to the packet stream with the highest window-denominator. Suppose that S_j is serviced before S_i, because $y_j' > y_i'$. At some later point in time, S_i will have the highest window-denominator, since its denominator is increased by ϵ every request period, T_i, that a packet in S_i is delayed, while S_j's window-constraint is reset once it is serviced. For simplicity, we assume every stream has the same value of ϵ but, in practice, it may be beneficial for each stream to have its own value, ϵ_i, to increase its need for service at a rate independent of other streams, even when window-constraint violations occur. Unless stated otherwise,

$\epsilon = 1$ is used throughout the rest of this paper.

To complete this section, Figure 3 shows an example schedule using both DWCS and EDF, for three streams, S_1, S_2, and S_3. For simplicity, assume that every time a packet in one stream is serviced, another packet in the same stream requires service. It is left to the reader to verify the scheduling order for DWCS. Observe that, using DWCS, all window-constraints are met over non-overlapping windows of y_i deadlines (for each stream, S_i), and no time slots are unused in this example. Moreover, the three streams are serviced in proportion to their original window-constraints and request periods. Consequently, S_1 is serviced twice as much as S_2 and S_3 over the interval $t = [0, 16]$. By contrast, EDF arbitrarily schedules packets with equal deadlines, irrespective of which packet is from the more critical stream in terms of its window-constraint. In this example, EDF selects packets with equal deadlines in strict alternation but the window-constraints of the streams are not guaranteed.

Note that EDF scheduling is optimal in the sense that if it is possible to produce a schedule in which all deadlines are met, such a schedule can be produced using EDF. Consequently, if C_i is the service time for a packet in stream S_i, then if $\sum_{i=1}^{n} \frac{C_i}{T_i} \leq 1.0$ all deadlines will be met using EDF[16]. However, in this example, $\sum_{i=1}^{n} \frac{C_i}{T_i} = 3.0$ so not all deadlines can be met. Since, $\sum_{i=1}^{n} \frac{(1-W_i)C_i}{T_i} = 1.0$, it is possible to strategically miss deadlines for certain packets and thereby guarantee the window-constraints of each stream. By considering window-constraints when deadlines are tied, DWCS is able to make guarantees that EDF cannot.

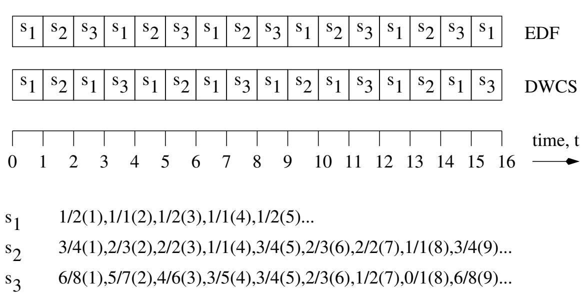

Figure 3. Example showing the scheduling of 3 streams, S_1, S_2, and S_3, using EDF and DWCS. All packets in each stream have unit service times and request periods. The window-constraints for each stream are shown as fractions, x/y, while packet deadlines are shown in brackets.

2.3. DWCS Complexity

DWCS's time complexity is divided into two parts: (1) the cost of *selecting* the next packet according to the precedence rules in Table 1, and (2) the cost of *adjusting* stream window-constraints and packet deadlines *after* servicing a packet. Using heap data structures for prioritizing packets, the cost of selecting the next packet for service is $O(log(n))$, where n is the number of streams awaiting service. However, after servicing a packet, it may be necessary to adjust the deadlines of the head packets, and window-constraints, of all n queued streams. This is the case when all $n-1$ streams (other than the one just serviced) have packets that miss their deadlines. This can lead to a worst-case complexity for DWCS of $O(n)$. However, the average case performance is typically a logarithmic function of the number of streams (depending on the data structures used for maintaining scheduler state), because not all streams always need to have their window-constraints adjusted after a packet in any given stream is serviced. When only a *constant* number of packets in different streams miss their deadlines after servicing some other packet, a heap data structure can be used to determine those packet deadlines and stream window-constraints that need to be adjusted. It follows that a constant number of updates to service constraints using heaps, as described in an earlier paper[24], requires $O(log(n))$ operations. Observe that there is an $O(1)$ cost per stream to update the corresponding service constraints, *after servicing a packet.*

In reality, the costs associated with DWCS compare favorably to those of fair queueing algorithms. Observe that with Weighted Fair Queueing (WFQ)[7] (as with other fair queueing variants), the time complexity consists of: (1) the cost of calculating a per packet virtual time, $v(t)$, *upon packet arrival* at the input to the scheduler, which is then used to derive an ordering *tag* (typically a packet start or finish tag), and (2) the cost of determining the next packet for service based on each packet's tag. The cost of part (2) is the same as the cost of selecting the next packet for service in DWCS, and can be implemented in $O(log(n))$ time using a heap. The calculation of the virtual time, $v(t)$, in part (1), is $O(n)$ in WFQ, since it is a function of all backlogged sessions (i.e., streams) at time t. Algorithms such as Start-time Fair Queueing (SFQ)[10, 18], Self-Clocked Fair Queueing (SCFQ)[9], or Frame-Based Fair Queueing (FFQ) and Starting-Potential Fair Queueing (SPFQ)[21], have an $O(1)$ complexity for calculating virtual timestamps (and ordering tags) *per packet.* However, if a burst of packets arrive from n different streams, then a timestamp and, hence, tag has to be calculated for all n packet arrivals. When the scheduler is heavily loaded, for example, during bursty periods, it is possible for packets to arrive from n different streams after scheduling one packet. Thus, the cost of scheduling with an efficient fair queueing algorithm is similar to that of DWCS in many situations, especially when the scheduler is heavily loaded.

The per-stream state requirements of DWCS include the head packet's deadline (computed from a stream's request period and the current time), a stream's window-constraint, and a single-bit violation tag. Due to the time and space requirements of DWCS, we feel it is possible to implement the algorithm at network access points and, possibly, within switches too. An earlier paper[24] shows how DWCS can be approximated, to further reduce its scheduling latency, thereby improving service scalability at the cost of potentially violating some service constraints. Moreover, it may be appropriate to combine multiple streams into one session, with DWCS used to service the aggregate session. Such an approach would reduce the scheduling state requirements and increase scalability. In fact, this is the approach taken in our simulated experiments in Section 3.4.

Having described DWCS, we analyze the algorithm's performance in the following section.

3. Analysis of DWCS

In this section we show the following important characteristics of the DWCS algorithm (as defined in this paper):

- If a feasible schedule is known to exist, DWCS ensures that the maximum delay of service to a real-time packet stream is bounded. The exact value of this maximum delay is characterized below.
- If window-constraint violations occur (because the scheduler is overloaded), the maximum queueing delay of a packet stream (and, hence, packet) is still bounded. Again, the exact value of this maximum delay is characterized below.
- A simple on-line feasibility test for DWCS exists, assuming each stream is serviced at the granularity of a fixed-sized time slot, and all request periods are multiples of such a time slot (see Figure 1). A time slot can be thought of as the time to service one or more packets from any given stream, and no two streams can be serviced in the same time slot. For simplicity, we assume that at most one packet from any given stream is serviced in a single time slot. Consequently, if the *minimum* aggregate bandwidth requirement of all real-time packet streams does not exceed the total available bandwidth, then a feasible schedule is guaranteed using DWCS.
- For networks with fixed-length packets, a time slot is at the granularity of the service time of one packet. However, for variable rate servers, or in networks where packets have variable lengths, the service times can vary for different packets. In such circumstances, if it is possible to impose an upper bound on the worst-case service time of each and every packet, then DWCS can still guarantee that no more than x packet deadlines are missed every y requests. In this case, service is granted to packet streams at the granularity of a time slot, which represents the worst-case service time of any packet. Alternatively, if it is possible to fragment variable-length packets and later reassemble them at the destination, per-stream service requirements can be translated and applied to fixed-length packets with constant service times, representing a time slot in a DWCS-based system.

3.1. Delay Characteristics

Theorem 1 *If a feasible schedule exists, the maximum delay of service to a stream, $S_i \mid 1 \leq i \leq n$, is at most $(x_i + 1)T_i - C_i$, where C_i is the service time for one packet in S_i*[2].

Proof Every time a packet in S_i misses its deadline, x'_i is decreased by 1 until x'_i reaches 0. A packet misses its deadline if it is delayed by T_i time units without service. Observe that, at all times, $x'_i \leq x_i$. Therefore, service to S_i can be delayed by at most $x_i T_i$ until $W'_i = 0$. If S_i is delayed more than another $T_i - C_i$ time units, a window-constraint violation will occur, since service of the next packet in S_i will not complete by the end of its request period, T_i. Hence, S_i must be delayed at most $(x_i + 1)T_i - C_i$ if a feasible schedule exists. □

We now characterize the delay bound for a packet stream when window-constraint violations occur, assuming all request periods are greater than or equal to each and every packet's service time. That is, $T_i \geq C_i, x_i \geq 0, y_i > 0, \forall i \mid 1 \leq i \leq n$.

Theorem 2 *If window-constraint violations occur, the maximum delay of service to S_i is no more than $T_i(x_i + y_{max} + n - 1) + C_{max}$, where $y_{max} = max[y_1, \cdots, y_n]$ and C_{max} is the maximum packet service time amongst all queued packets.*

Proof The details of this proof are shown in the Appendix.

If $T_i \to \infty$, then S_i experiences unbounded delay in the worst-case. This is the same problem with static-priority scheduling, since a higher priority packet stream will always be serviced before a lower priority packet stream. Observe that in calculating the worst-case delay experienced by S_i, it is assumed that $dy'_i/dt = \epsilon/T_i \mid \epsilon = 1$ (see Figure 5). If $\epsilon > 1$ or there is a unique value, $\epsilon_i > 1$ for each stream S_i, then the worst-case delay experienced by S_i is $\frac{T_i(x_i + y_{max} + n - 1)}{\epsilon_i} + C_{max}$. If $\epsilon_i = (x_i + y_{max} + n - 1)$ then the worst-case delay of S_i is $T_i + C_{max}$, which is independent of the number of streams. Consequently, the worst-case delay of service to each stream can be made to be independent of all other streams, even in overload situations.

3.2. Bandwidth Utilization

As stated earlier, $W_i = x_i/y_i$ for stream S_i. Therefore, a minimum of $y_i - x_i$ packets in S_i must be serviced 'on time' every window of y_i consecutive packets, for S_i to satisfy its window-constraints. Since one packet is required to be serviced every request period, T_i, to avoid any packets in S_i being

[2] For simplicity, we assume all packets in the same stream have the same service time. However, unless stated otherwise, this constraint is not binding and the properties of DWCS should still hold.

late, a minimum of $y_i - x_i$ packets must be serviced every y_iT_i time units. Therefore, if each packet takes C_i time units to be serviced, then y_i packets in S_i require at least $(y_i - x_i)C_i$ units of service time every y_iT_i time units. For a packet stream, S_i, with request period, T_i, the *minimum* utilization factor is $U_i = \frac{(y_i - x_i)C_i}{y_iT_i}$, which is the minimum required fraction of available service capacity and, hence, bandwidth by consecutive packets in S_i. Hence, the utilization factor for n packet streams is at least $U = \sum_{i=1}^{n} \frac{(1-W_i)C_i}{T_i}$. Furthermore, the *least upper bound* on the utilization factor is the minimum of the utilization factors for all packet streams that fully utilize all available bandwidth[16]. If U exceeds the least upper bound on bandwidth utilization, a feasible schedule is not guaranteed. In fact, it is necessary that $U \leq 1.0$ is true for a feasible schedule, using any scheduling policy.

We now characterize the least upper bound on bandwidth utilization, assuming that at most one packet from any given stream is serviced in a single, fixed-sized time slot of size K, and all request periods are multiples of such a time slot. That is, $C_i \leq K, T_i = q_iK, x_i \geq 0, y_i > 0, \forall i \mid 1 \leq i \leq n$, K is a constant, and q_i is a positive integer.

Theorem 3 *Using DWCS, the least upper bound on the utilization factor is* 1.0*, if all streams comprise packets with the same service times, and all request periods are multiples of the packet service times. That is, DWCS is optimal in the sense that a feasible schedule exists if* $\sum_{i=1}^{n} \frac{(1-W_i)C_i}{T_i} \leq 1.0$*, given* $C_i = K$ *and* $T_i = q_iK$ *for* $q_i \in Z^+$*, where* Z^+ *is the set of positive integers.*

Proof The details of this proof are also shown in the Appendix.

Observe that in this proof, all packet service times are assumed equal, and all request periods are assumed to be multiples of the constant packet service time. That is, each packet is serviced for exactly K time units, which is the size of one time slot. This is a necessary condition, because the optimality of earliest-deadline first scheduling (which forms the basis of DWCS as presented in this paper) assumes that arbitrary preemption of service is possible. However, for packet scheduling, packets are indivisible entities and, hence, cannot be preempted.

3.3. Supporting Packets with Variable Service Times

For variable rate servers, or in networks where packets have variable lengths, the service times can vary for different packets. In such circumstances, if it is possible to impose an upper bound on the *worst-case* service time of each and every packet, then DWCS can still guarantee that no more than x packet deadlines are missed every y requests. This implies that the scheduling granularity, K (i.e., one time slot), should be set to the worst-case service time of any packet scheduled for transmission. For situations where a packet's service time, C_i, is less than K (see Figure 1), then a feasible schedule is still possible using DWCS, but the least upper bound on the utilization factor is less than 1.0. That is, if $\tau_i = K - C_i$, then, the least upper bound on the utilization factor is $1.0 - \sum_{i=1}^{n} \frac{(1-W_i)\tau_i}{T_i}$.

Alternatively, if it is possible to fragment variable-length packets and later reassemble them at the destination, per-stream service requirements can be translated and applied to fixed-length packet fragments with constant service times. This is similar to CPU scheduling, in which variable-length threads (or processes) can be preempted at fixed intervals (e.g., every $10mS$ timeslice). Moreover, ATM networks have fixed-length (53 byte) cells and the SAR component of the ATM Adaptation Layer segments application-level packets into cells, which are later reassembled. Consequently, the scheduling granularity, K, can be set to a time which is less than the worst-case service time of a packet.

For fragmented packets, the per-stream service constraints are translated as follows. Let C_i be the *worst-case* service time of a packet in stream S_i before fragmentation, and let $c_i = K$ be the constant service time of each and every fragment. Likewise, let W_i and T_i be the window-constraint and request period, respectively, for a stream before fragmentation, while w_i and t_i are the translated window-constraint and request period, respectively, after fragmentation. Then:

$c_i = K, t_i = \lfloor \frac{T_i}{C_i} \rfloor K$ and $w_i = a_i/b_i$, where a_i and b_i are the smallest values satisfying $a_i/b_i = \frac{T_i c_i - C_i(1-W_i)t_i}{T_i c_i}$.

Example. Consider three streams, S_1, S_2 and S_3 with the following constraints: $(C_1 = 3, W_1 = 2/3, T_1 = 5)$, $(C_2 = 4, W_2 = 23/35, T_2 = 6)$ and $(C_3 = 5, W_3 = 1/5, T_3 = 7)$. The total utilization factor is 1.0 in this example, but due to the non-preemptive nature of the variable-length packets, a feasible schedule cannot be constructed. However, if the packets are fragmented and the per-stream service constraints are translated to be $(c_1 = 1, w_1 = 4/5, t_1 = 1)$, $(c_2 = 1, w_2 = 27/35, t_2 = 1)$ and $(c_3 = 1, w_3 = 3/7, t_3 = 1)$, then a feasible schedule exists. In the latter case, all fragments are serviced so that their corresponding stream's window-constraints are met. These translated window-constraints are equivalent to the original window-constraints, thereby guaranteeing each stream its exact share of bandwidth. Observe that $c_i = t_i = 1$ is the normalized time to service one fragment of a packet. This fragment could be a single cell in an ATM network but, more realistically, it makes sense for one fragment to map to multiple ATM cells, thereby reducing the scheduling overheads per fragment. Similarly, a fragment might correspond to a maximum transmission unit in an Ethernet-based network.

3.4. Simulated Results

To show that it is possible to feasibly schedule a set of packet streams when the demand for bandwidth is no more than 100% of available bandwidth, we simulated the number of missed deadlines and window-constraint violations for a number of streams, comprising fixed (unit) length packets, with different request periods and original window-constraints. The following three scenarios were considered:

- *Scenario 1*: There were 8 scheduling classes for packet streams. The original window-constraints for each class of packet streams were $1/10, 1/20, 1/30, 1/40, 1/50, 1/60, 1/70$, and $1/80$, and the request period for each packet in every stream was 480 time units. Each packet in every stream required at most one time unit of service in its request period, or else that packet missed its deadline.
- *Scenario 2*: This was the same as *Scenario 1* except that the request periods for packets in streams belonging to the first 4 classes (with window-constraints $1/10$ to $1/40$) were 240 time units, and the request periods for packets in streams belonging to the remaining 4 classes were 320 time units.
- *Scenario 3*: This was the same as *Scenario 1* except that the request periods for packets in streams belonging to each pair of classes (starting from the class with a window-constraint of $1/10$) were 400, 480, 560, and 640 time units, respectively.

Tables 2(a), 2(b) and 2(c) show the results of scenarios 1, 2, and 3, respectively. n is the total number of packet streams, D is the number of missed deadlines, V is the number of window-constraint violations, and U is the total utilization factor. The number of packet streams in each case was uniformly distributed between each scheduling class, and a total of a million packets across all streams were serviced. It should be clear from the tables that, although some packets miss their deadlines when the worst-case total utilization factor of all packet streams is less than 1.0, there are no window-constraint violations for any streams until the utilization factor exceeds 1.0.

(a) Scenario 1			
n	D	V	U
240	0	0	0.4830
320	0	0	0.6440
400	0	0	0.8050
480	0	0	0.9660
488	16664	0	0.9821
496	33328	0	0.9982
504	49992	14344	1.0143
512	66656	30295	1.0304
520	83320	44753	1.0465

(b) Scenario 2			
n	D	V	U
80	0	0	0.2810
160	0	0	0.5620
240	0	0	0.8430
256	0	0	0.8921
272	0	0	0.9554
280	20820	0	0.9835
288	49968	11215	1.0116
304	108264	54444	1.0678
320	166560	88544	1.1240

(c) Scenario 3			
n	D	V	U
480	0	0	0.9156
496	0	0	0.9461
504	0	0	0.9613
512	15152	0	0.9766
520	30990	0	0.9919
528	46828	7038	1.0071
544	78528	31873	1.0376
560	110240	53455	1.0681
640	268800	148143	1.2207

Table 2. Number of missed deadlines, D, and window-constraint violations, V, for increasing numbers of packet streams, n, and increasing utilization factors, U. In all scenarios, the number of window-constraint violations remains zero until the total utilization factor, U, exceeds 1.0.

Summary. DWCS ensures the delay of service to real-time packet streams is bounded even in the absence of a feasibility test, whereby the scheduler may be overloaded and window-constraint violations can occur. Consequently, DWCS guarantees that a packet stream will never suffer starvation. Furthermore, the least upper bound on bandwidth utilization using DWCS can be as high as 100%.

Having described some important characteristics of DWCS for real-time packet streams, we now describe how best-effort packet streams can be serviced with low average delay, while still guaranteeing service to real-time packet streams.

4. Heterogeneous Packet Streams

In many situations, it is desirable, or even necessary, to service a mixture of both real-time and best-effort packet streams. Many researchers have proposed that best-effort, or non-time-constrained packet streams are only scheduled when all real-time packet streams have been serviced. Other researchers, in real-time systems research[6, 3], have attempted to reduce the mean delay of non-time-constrained activities (such as threads or packets) by giving them precedence over real-time activities until it is essential to service the real-time activities.

$U_{min,WC}$	U_{BE}	$U_{min,WC} + U_{BE}$
0.0000	1.00	1.0
0.1526	0.84	0.9926
0.3052	0.68	0.9852
0.4578	0.52	0.9778
0.6104	0.37	0.9804
0.7630	0.21	0.9730
0.9156	0.05	0.9656
0.9461	0.02	0.9661
0.9613	0.0001	0.9614
0.9766	0.0	0.9766
0.9919	0.0	0.9919
1.0071	0.0	1.0071

Table 3. $U_{min,WC}$ is the minimum utilization factor of all the window-constrained (WC) packet streams when there are no window-constraint violations. $U_{min,WC}$ is *calculated* from *Scenario 3* in Section 3.2. In this scenario, U_{BE} is the *measured* utilization factor of best-effort packet streams when they are serviced. A best-effort packet stream is serviced when one packet from each and every real-time, window-constrained packet stream has been serviced in its current request period.

For combined best-effort and window-constrained, real-time packet streams, our approach is to service best-effort packet streams only when one packet from each and every real-time, window-constrained packet stream has been serviced in its current request period. Note that we only allow multiple packets in real-time streams to be serviced in the same request period if such packet streams have been *marked* as eligible for scheduling multiple times in one request period. For all normal cases, only one packet in a real-time stream can be serviced in any given request period. Furthermore, the next packet in a real-time stream has a deadline that is offset from the previous packet's deadline by the stream's request period.

Table 3 shows the measured utilization factor, U_{BE}, of best-effort packet streams in the presence of real-time packet streams having the same service constraints as in *Scenario 3*, in Section 3.2. The number of real-time packet streams is increased, thereby increasing their utilization factor. $U_{min,WC}$ shows the minimum utilization factor of all the real-time, window-constrained (WC) packet streams when there are no window-constraint violations. If $U_{min,WC}$ were the measured utilization factor of real-time packet streams, then the best-effort packet streams could use the remaining fraction of bandwidth to minimize their delay. One way to minimize the delay of best-effort packet streams is to calculate a *pseudo* request period, T_{BE}, and window-constraint, W_{BE}, so that $1 - \sum_{i=1}^{n} \frac{(1-W_i)C_i}{T_i} = \frac{(1-W_{BE})C_{BE}}{T_{BE}}$, when there are n real-time, window-constrained packet streams. However, with this approach, there can be cases where real-time packet streams miss deadlines due to best-effort packet streams being serviced. In some cases, this may be acceptable, since each real-time stream only violates a tolerable number of packet deadlines, and does not violate its window-constraint. In other cases, we want to ensure real-time packet streams *never* miss deadlines when best-effort packet streams are serviced. Hence, our alternative approach is to service best-effort packet streams only when a packet from each and every window-constrained packet stream has been serviced in each real-time stream's current request period. This guarantees packets in real-time streams do not miss any deadlines due to servicing best-effort packet streams. From Table 3, the sum, $U_{min,WC} + U_{BE}$, is still close to 1.0. Since the *actual* utilization factor of both real-time and best-effort packet streams is 1.0, in this scenario, the real-time packet streams are actually experiencing a utilization factor above their minimum required utilization factor, $U_{min,WC}$. Consequently, there are some cases when a real-time, window-constrained stream is being serviced when it need not be serviced. This means the delay of best-effort packet streams is greater than the minimum possible delay. However, since $U_{min,WC} + U_{BE}$ is close to 1.0, the best-effort packet streams are experiencing close to the minimum possible delay.

Figure 4(a) shows the number of best-effort packet streams serviced, as a function of all packets serviced from both best-effort and real-time streams. Each set of real-time packet streams has a different utilization factor (hence, the six different lines in the graph). In all cases, the service constraints of real-time packet streams were the same as *Scenario 3*, in Section 3.2. The utilization factor of these real-time packet streams was increased, by increasing the number of packet streams in each of 8 different scheduling classes, from 10 to 60 packet streams per class. From the figure, it can be seen that there is a constant rate of service to best-effort packet streams at each of the different loads from real-time packet streams. This is useful, in that best-effort packet streams will not experience large variations in delay

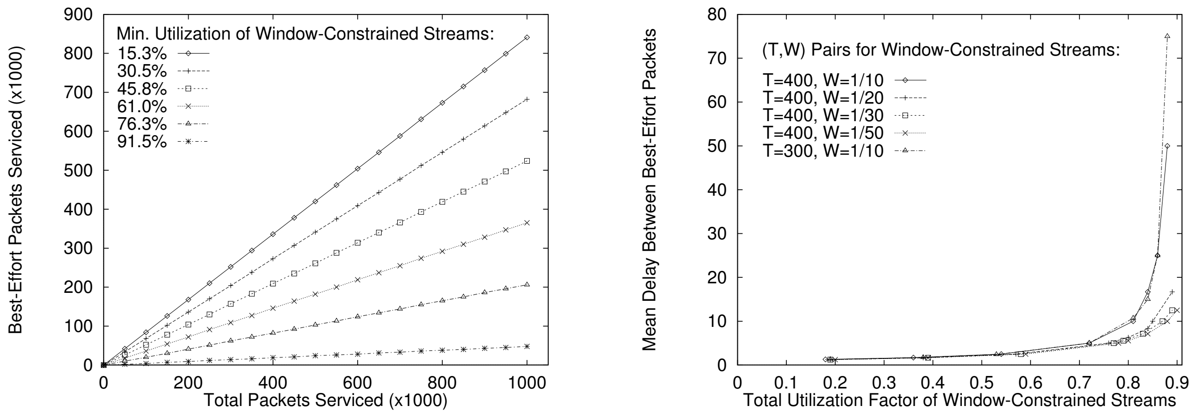

Figure 4. (a) The number of best-effort packets serviced, as a function of all packets serviced from both best-effort and real-time streams. The results are shown when real-time streams have different *minimum* **utilization factors; (b) the mean delay (as a function of packet service times) between servicing consecutive packets in best-effort streams under different loads from real-time packet streams.**

(and, hence, jitter) in the presence of real-time packet streams.

Figure 4(b) shows the mean delay between servicing consecutive packets in best-effort packet streams under different loads from real-time packet streams. The mean delay is a multiple of packet service times, which are assumed constant. Observe that, when real-time packet streams impose the same load but have different request periods, T, and window-constraints, W, then best-effort packet streams experience different delays. In fact, for loads above about 55%, there are different delays for best-effort packet streams that are dependent upon the service constraints of real-time packet streams. It can be seen that, for a given real-time load, increasing the request period, T, of real-time packet streams, can reduce the delay of best-effort packet streams. This is because only one packet in a real-time stream is serviced in its request period, T. By increasing T, for any given real-time load, there is a larger window of time to service best-effort packet streams, since real-time packet streams are not eligible for service again until their next request periods. Likewise, reducing the magnitude of the window-constraints (in this example, from 1/10 to 1/50), also reduces the delay incurred by packets in best effort streams. For any given real-time load, there will be fewer real-time packet streams if the utilization factor of each stream is higher. By reducing the window-constraints from 1/10 to 1/50, we are increasing the minimum required utilization factor of real-time packet streams. Consequently, there are fewer real-time packet streams that must be serviced in any request period, T. This means that more best-effort packet streams can be serviced in a given window of time. As a result, for higher loads, the mean delay of packets in best-effort streams is reduced, when real-time window-constraints are reduced.

Summary. For combined best-effort and window-constrained, real-time packet streams, our approach is to service best-effort packet streams only when one packet from each and every real-time, window-constrained packet stream has been serviced in its current request period. This is conservative, in that it ensures no packet in any window-constrained stream misses its deadline as a result of servicing best-effort packet streams. However, there may be cases when we want to allow some packets to miss their deadlines in real-time streams, as long as the window-constraints of these packet streams are not violated. In the latter case, we can calculate a *pseudo* request period, T_{BE}, and window-constraint, W_{BE} for best-effort packet streams. Furthermore, these best-effort packet streams can be prioritized to ensure precedence is given to the highest priority, non-time-constrained stream when it is safe to service such a packet stream. Finally, best-effort packet streams can be serviced at a constant rate. This minimizes the variations in delay of service to consecutive packets in best-effort streams, in the presence of real-time streams.

5. Conclusions and Future Work

This paper describes a modified version of Dynamic Window-Constrained Scheduling (DWCS)[23, 24]. DWCS was originally designed as a packet scheduler to provide (m, k)-*firm* deadline guarantees[12] and fair queueing[7, 25, 9, 2, 10, 18, 22] properties, for loss and delay constrained traffic streams such as multimedia audio and video streams. In this paper, we have shown: (1) a version of DWCS that can guarantee (m, k)-*hard* deadlines (or, equivalently, no more than x missed packet deadlines for every y consecutive packets in a given stream), (2) using DWCS, the delay of service to real-time packet streams is bounded even when the scheduler is overloaded, (3) DWCS can ensure the delay bound of any given stream is independent of other streams, and (4) a fast response time for best-effort packet streams, in the presence of real-time packet streams, is possible.

DWCS can be thought of as a special case of pinwheel scheduling[13], in which a minimum of m deadlines are guaranteed to be met, every *fixed* (as opposed to *sliding*) window of k deadlines. However, DWCS has a least upper bound on resource (such as bandwidth) utilization[16] of 100%, independent of the window size, k. By comparison, Baruah and Lin[1] have developed an algorithm for pinwheel systems with a least upper bound on utilization that *asymptotically approaches* 100%, only when $k \to \infty$. Like Baruah and Lin's algorithm, DWCS is computationally efficient, with a cost linear in the number of packet streams, to determine the next packet for service.

Finally, our future work includes implementing a DWCS packet scheduler on Intel i960-based network interface boards and in the Linux kernel, to complement our DWCS CPU scheduler in Linux (which can be downloaded from: *http://www.cc.gatech.edu/~west/dwcs.html*). This work will form the basis of our ongoing adaptive systems research, involving the coordination of communication and computation resources.

A Appendix

A.1. Proof of Theorem 2

The worst-case delay experienced by S_i can be broken down into three parts: (1) the time for the next packet in S_i to have the earliest deadline amongst all packets queued for service, (2) the time taken for W'_i to become the minimum amongst all current window-constraints, $W'_k \mid 1 \leq k \leq n$, when the head packets in all n streams have the same (earliest) deadline, and (3) the time for y'_i to be larger than any other current denominator, $y'_j \mid j{\neq}i, 1 \leq j \leq n$, amongst each packet stream, S_j, with the minimum current window-constraint and earliest packet deadline. At this point, S_i may be delayed a further C_{max} due to another packet currently in service.

Part (1): The next packet in S_i is never more than T_i away from its deadline. Consequently, S_i will have a packet with the earliest deadline after a delay of at most T_i.

Part (2): $W'_i = 0$ is the minimum possible current window-constraint. From Theorem 1, $W'_i = 0$ after a delay of at most $x_i T_i$.

Parts (1) and (2) contribute a maximum delay of:

$$(x_i + 1)T_i \tag{1}$$

Part (3): Assuming all packet streams have the minimum current window-constraint and comprise a head packet with the earliest deadline, the next stream chosen for service is the one with the highest current window-denominator. Moreover, the worst-case scenario is when all other packet streams have the same or higher current window-denominators than S_i and every time another stream, S_j is serviced, deadline $d_j{\leq}d_i$. To show that $d_j{\leq}d_i$ holds, all deadlines must be at the same time, t, when some stream S_j is serviced in preference to S_i. After servicing a packet in S_j for C_j time units, all packet deadlines d_k that are earlier than $t + C_j$ are incremented by a multiple of the corresponding request periods, $T_k \mid 1 \leq k \leq n$, depending on how many request periods have elapsed while servicing S_j. The worst-case is that $T_j{\leq}T_i, \forall j \neq i$. Furthermore, every time a stream, S_j, other than S_i is serviced, $W'_j = 0$. This is true regardless of whether or not S_j is tagged with a violation, if $W_j = 0$, which is the case when $x_j = 0$.

Hence, the worst-case delay incurred by S_i when $W'_i = 0$ is $T_i + \delta_i$, where δ_i is the maximum time for y'_i to become larger than any other current denominator, $y'_j \mid j{\neq}i, 1 \leq j \leq n$, amongst all packet streams with the minimum current window-constraint and earliest packet deadline. Now, let state ϕ be when each stream, S_k, has $W'_k = 0$ for the first time. Moreover, $W'_k = 0/y'_{k_\phi}$, and $y'_{k_\phi} > 0$ is the current window-denominator for S_k when in state ϕ.

Suppose $T_j{\leq}T_i, \forall j \neq i$ and T_j is finite. For n packet streams, the worst-case δ_i is when $T_j = K$ and $T_i >> K$, for some constant, K, equal to the largest packet service time, C_{max}. Without loss of

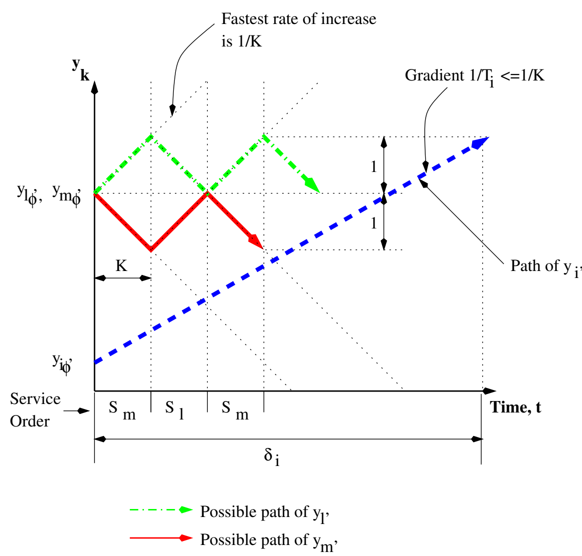

Figure 5. The change in current window-denominators, y'_i, y'_l and y'_m for three packet streams, S_i, S_l and S_m, respectively, when all request periods, except possibly T_i, are finite. The initial state, ϕ, is when all current window-constraints first equal 0, and the current window-denominators are all greater than 0.

generality, it can be assumed in what follows that all packet service times equal C_{max}. Now, it should be clear that, if T_i tends to infinity, then the rate of increase of y'_i approaches 0. Moreover, if each and every packet stream, $S_j \mid j \neq i$, has a request period, $T_j = K$, then S_i will experience its worst delay before $y'_i \geq y'_j$. This is because y'_j rises at a rate of $1/K$ for each stream S_j experiencing a delay of K time units without service, while y'_i increases at a rate of $1/T_i$, which is less than or equal to $1/K$.

Figure 5 shows the worst-case situation for three packet streams, S_i, S_l, and S_m, which causes S_i the largest delay, δ_i, before y'_i is the largest current window-denominator. From the figure, $y'_{l_\phi} = y'_{m_\phi}$, and y'_i increases at a rate $dy'_i/dt = \epsilon/T_i \mid \epsilon = 1$, until S_i is serviced. When S_m is serviced, y'_m decreases at a rate of $1/K$, while y'_l increases at a rate of $1/K$. Conversely, when S_l is serviced, y'_l decreases at a rate of $1/K$, while y'_m increases at a rate of $1/K$. Only when $y'_m = 0$ is W'_m reset. Likewise, only when $y'_l = 0$ is W'_l reset. Consequently, $y'_i \geq max[y'_l, y'_m]$ is true when $y'_i = y'_{l_\phi} + 1 = y'_{m_\phi} + 1$.

Suppose now, another stream, S_o (with $y'_{o_\phi} = y'_{l_\phi} = y'_{m_\phi}$ and $T_o = K$), is serviced before either S_l or S_m when in state ϕ. Then, $y'_l = y'_m = y'_{l_\phi} + 1 = y'_{m_\phi} + 1$ after K time units. If S_l is now serviced, then $y'_m = y'_{m_\phi} + 2$ after a further K time units. In this case, $y'_i \geq max[y'_l, y'_m, y'_o]$ is true when $y'_i = y'_{l_\phi} + 2 = y'_{m_\phi} + 2 = y'_{o_\phi} + 2$. By induction, for each of the $n-1$ packet streams, $S_j \mid j \neq i, 1 \leq j \leq n$, other than S_i, each with $T_j = K$ and $y'_{j_\phi} \geq y'_{i_\phi}$, $y'_i \geq max[y'_1, y'_{i-1}, \cdots, y'_{i+1}, y'_n]$ is true when $y'_i = y'_{1_\phi} + (n-2) = \cdots = y'_{n_\phi} + (n-2)$. Therefore, since $dy'_i/dt = 1/T_i$, it follows that $\delta_i \leq T_i(y'_{j_\phi} - y'_{i_\phi} + (n-2))$.

Now observe that $y'_{j_\phi} \leq y_j$ for each and every stream, $S_j \mid j \neq i$, since state ϕ is the first time W'_j is 0. Furthermore, we have the constraints that $y_j = max[y_1, y_{i-1}, y_{i+1}, y_n]$, $y_i \leq y_j$, and $y'_{i_\phi} \geq 1$. Therefore,

$$\delta_i \leq T_i(y_j + (n-2)) \tag{2}$$

If $T_j > T_i, \forall j \neq i$ and both T_j and T_i are finite, then y'_i and y'_j converge more quickly than in the case above, when $T_j \leq T_i$. Therefore, if window-constraint violations occur, the maximum delay of service to S_i (from Equations 1 and 2) is no more than

$$(x_i + 1)T_i + T_i(y_{max} + n - 2) + C_{max} = T_i(x_i + y_{max} + n - 1) + C_{max},$$

where $y_j = y_{max}$ in Equation 2, and C_{max} is the worst-case additional delay due to another packet in service when a packet in S_i reaches the highest priority. □

A.2. Proof of Theorem 3

There are three possible cases to consider.

Case 1: All packet streams have current window-constraints equal to 0. That is, $W'_l = 0, \forall l \mid 1 \leq l \leq n$. In this situation, it is necessary to meet all deadlines, and DWCS services the packets of each stream in earliest deadline first order. This is known to be optimal in the sense that if all packets in each and every stream can be serviced by their deadlines, then these packet streams can be serviced by their deadlines using earliest deadline first scheduling[16].

Case 2: All packet streams have current window-constraints greater than 0. That is, $W'_l > 0, \forall l \mid 1 \leq l \leq n$. Since DWCS services packets in each stream in earliest deadline first order, we only need to consider window-constraints when two or more packet streams have head packets with the same earliest deadline. Consider packet streams S_i and S_j and the set of packet streams, $\overline{S_k}$, which includes all packet streams except S_i and S_j. Let W'_k and T_k be the current window-constraint and the request period, respectively, for any stream in $\overline{S_k}$. Observe that $W'_i = x'_i/y'_i$, $W'_j = x'_j/y'_j$ and $W'_k = x'_k/y'_k$. If all current window-constraints are *normalized* with the same denominator, $W'_i = \frac{x'_i y'_1 \cdots y'_{i-1} y'_{i+1} \cdots y'_n}{\prod_{l=1}^{n} y'_l}$, $W'_j = \frac{x'_j y'_1 \cdots y'_{j-1} y'_{j+1} \cdots y'_n}{\prod_{l=1}^{n} y'_l}$ and $W'_k = \frac{x'_k y'_1 \cdots y'_{k-1} y'_{k+1} \cdots y'_n}{\prod_{l=1}^{n} y'_l}$.

Let $W'_i \leq W'_j \leq W'_k$ be true, so that

$$x'_i y'_1 \cdots y'_{i-1} y'_{i+1} \cdots y'_n \leq x'_j y'_1 \cdots y'_{j-1} y'_{j+1} \cdots y'_n \leq x'_k y'_1 \cdots y'_{k-1} y'_{k+1} \cdots y'_n.$$

From Theorem 1, the *normalized* maximum delay, $\Delta_{i_{norm}}$, of S_i, before it violates W_i' is

$$(x_i'y_1'\cdots y_{i-1}'y_{i+1}'\cdots y_n'+1)T_i.$$

Suppose now that, for any pair of packet streams, S_m and S_p, in $\overline{S_k}$, that if $W_m'\leq W_p'$ is true, then $\Delta_{m_{norm}}\leq\Delta_{p_{norm}}$. Consequently, if $T_i\leq T_j\leq T_k$ is true, then $\Delta_{i_{norm}}\leq\Delta_{j_{norm}}\leq\Delta_{k_{norm}}$. Therefore, servicing S_i before S_j (before any other stream in the set $\overline{S_k}$) when $W_i'\leq W_j'\leq W_k'$ is the same as servicing packets in earliest deadline first order, which (as stated earlier) is known to be optimal. This is exactly what happens with DWCS, since it services the stream with the lowest current window-constraint, out of those packet streams with the earliest packet deadlines. Consequently, when there are two or more packet streams tied for the earliest packet deadline, scheduling the packet streams in increasing order of current window-constraints is analogous to earliest deadline first scheduling of packets.

We shall now verify that the method of dynamically adjusting window-constraints does indeed ensure each stream meets its minimum utilization requirements. Suppose that, in the worst-case, S_j is not serviced during the next $\delta_j = x_j'y_1'\cdots y_{j-1}'y_{j+1}'\cdots y_n'T_j$ time units. After δ_j time units, $W_j' = 0$. However, packets in S_j are still serviced in (optimal) earliest deadline first order with all other packets in streams that have a current window-constraint of 0. Consequently, S_j is guaranteed K units of service time in the next T_j time units. Therefore, S_j is serviced once for K time units over the normalized interval $(x_j'y_1'\cdots y_{j-1}'y_{j+1}'\cdots y_n'+1)T_j$. Furthermore, while $W_j' = 0$, S_j must be serviced for K time units every T_j, which is guaranteed with earliest deadline first scheduling. In the worst-case, the fraction of bandwidth required by S_j must never exceed its minimum utilization factor. Furthermore, the fraction of bandwidth required by S_j must not drop below S_j's minimum utilization factor. Hence, in the worst-case, the utilization factor of S_j is $U_j = \frac{(y_j-x_j)K}{y_jT_j}$. To show this, first observe how W_j' changes with time. In what follows, let

$$\prod_{l=1}^{n} y_l = y_{norm}, \text{ and } x_jy_1\cdots y_{j-1}y_{j+1}\cdots y_n = x_{j_{norm}}.$$

If $W_j' = W_j$ initially, then the change of W_j' with time (where time is shown above the arrows) is

$$\frac{x_{j_{norm}}}{y_{norm}} \xrightarrow{\delta_j} \frac{0}{y_{norm}-x_{j_{norm}}} \xrightarrow{T_j} \frac{0}{y_{norm}-x_{j_{norm}}-1} \longrightarrow \overset{(y_{norm}-x_{j_{norm}}-1)T_j}{\cdots} \longrightarrow W_j.$$

When $W_j' = 0$, S_j is serviced for K time units every T_j, over the interval $(y_{norm}-x_{j_{norm}})T_j$. During this time, the denominator of W_j' is decremented by 1 every interval T_j, until it is zero, at which point W_j' is reset to its original value, W_j. Hence, over the interval $(x_{j_{norm}}+1)T_j + (y_{norm}-x_{j_{norm}}-1)T_j$, S_j is serviced for $K + (y_{norm} - x_{j_{norm}} - 1)K$ time units. Therefore, the utilization factor

$$U_j = \frac{K + (y_{norm} - x_{j_{norm}} - 1)K}{(x_{j_{norm}} + 1)T_j + (y_{norm} - x_{j_{norm}} - 1)T_j} = \frac{(y_{norm} - x_{j_{norm}})K}{y_{norm}T_j} = \frac{(y_j - x_j)K}{y_jT_j}.$$

Now observe that, $x'_j \leq x_j$ and $y'_j \leq y_j$ is always true, due to the rules for window-constraint adjustments when no violations are allowed. Furthermore, $x'_j = x_j - \sigma$ and $y'_j = y_j - \sigma$, where $\sigma \in \{0, 1, 2, \cdots\}$. Therefore, since, $y_j \geq x_j$, the worst-case minimum utilization factor for any stream, S_j is $U_j = \frac{(y_j - x_j)K}{y_jT_j}$. This guarantee ensures the least upper bound on the utilization factors of all packet streams never exceeds 1.0.

Case 3: a packet streams have current window-constraints greater than 0 and $n - a$ packet streams have current window-constraints equal to 0. In this case, DWCS services all packet streams in earliest packet deadline first order, except when two or more packet streams have the earliest packet deadline. If there is a tie for the earliest packet deadline, DWCS services one of the $n - a$ packet streams with a current window-constraint of 0. While each of the a packet streams is not serviced in its request period, its window-constraint is decreased. In the worst-case, each of the a packet streams reaches a state where the current window-constraint is 0. In such a case, these packet streams are guaranteed to be serviced in their next request periods, otherwise a window-constraint violation will occur. In fact, if all n packet streams reach a state where they have window-constraints equal to 0, DWCS schedules packet streams in a pure earliest packet deadline first order, as in **Case 1**.

Having satisfied all three cases, Theorem 3 is proved. □

References

[1] S. K. Baruah and S.-S. Lin. Pfair scheduling of generalized pinwheel task systems. *IEEE Transactions on Computers*, 47(7), July 1998.

[2] J. C. Bennett and H. Zhang. WF^2Q: Worst-case fair weighted fair queueing. In *IEEE INFOCOMM'96*, pages 120–128. IEEE, March 1996.

[3] G. Bernat and A.Burns. Combining (n/m)-hard deadlines and dual priority scheduling. In *Proceedings of the 18th IEEE Real-Time Systems Symposium*, pages 46–57, San Francisco, December 1997. IEEE.

[4] C. Carlsson and O. Hagsand. DIVE-a platform for multi-user virtual environments. *Computers and Graphics*, 17(6):663–669, November-December 1993.

[5] M. Chan and F. Chin. Schedulers for the pinwheel problem based on double-integer reduction. *IEEE Transactions on Computers*, 41(6):755–768, June 1992.

[6] R. Davis and A.Wellings. Dual priority scheduling. In *Proceedings of the 16th IEEE Real-Time Systems Symposium*, pages 100–109. IEEE, 1995.

[7] A. Demers, S. Keshav, and S. Schenker. Analysis and simulation of a fair-queueing algorithm. *Journal of Internetworking Research and Experience*, pages 3–26, October 1990.

[8] D. Ferrari. Client requirements for real-time communication services. *IEEE Communications Magazine*, 28(11):76–90, November 1990.

[9] S. Golestani. A self-clocked fair queueing scheme for broadband applications. In *INFOCOMM'94*, pages 636–646. IEEE, April 1994.

[10] P. Goyal, H. M. Vin, and H. Cheng. Start-time fair queueing: A scheduling algorithm for integrated services packet switching networks. In *IEEE SIGCOMM'96*. IEEE, 1996.

[11] S. Greenberg and D. Marwood. Real-time groupware as a distributed system: Concurrency control and its effect on the interface. In *Proceedings of the ACM Conference on Cooperative Support for Cooperative Work, ACM press*, pages 207–217. ACM, 1994.

[12] M. Hamdaoui and P. Ramanathan. A dynamic priority assignment technique for streams with (m,k)-firm deadlines. *IEEE Transactions on Computers*, April 1995.

[13] R. Holte, A. Mok, L. Rosier, I. Tulchinsky, and D. Varvel. The pinwheel: A a real-time scheduling problem. In *Proceedings of the 22nd Hawaii International Conference of System Science*, pages 693–702, Jan 1989.

[14] G. Koren and D. Shasha. Skip-over: Algorithms and complexity for overloaded systems that allow skips. In *Proceedings of the 16th IEEE Real-Time Systems Symposium*, pages 110–117. IEEE, December 1995.

[15] R. Kravets, K. Calvert, P. Krishnan, and K. Schwan. Adaptive variation of reliability. In *HPN-97*. IEEE, April 1997.

[16] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard real-time environment. *Journal of the ACM*, 20(1):46–61, January 1973.

[17] J. Nieh and M. S. Lam. The design, implementatin and evaluation of SMART: A scheduler for multimedia applications. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*. ACM, October 1997.

[18] X. G. Pawan Goyal and H. M. Vin. A hierarchical CPU scheduler for multimedia operating systems. In *2nd Symposium on Operating Systems Design and Implementation*, pages 107–121. USENIX, 1996.

[19] J. M. Peha and F. A. Tobagi. A cost-based scheduling algorithm to support integrated services. In *IEEE INFOCOMM'91*, pages 741–753. IEEE, 1991.

[20] S. Singhal. *Effective Remote Modeling in Large-Scale Distributed Simulation and Visualization Environments*. PhD thesis, Stanford University, August 1996.

[21] D. Stiliadis and A. Varma. Efficient fair queueing algorithms for packet-switched networks. *IEEE/ACM Transactions on Networking*, 6(2), 1998.

[22] I. Stoica, H. Abdel-Wahab, K. Jeffay, S. K. Baruah, J. E. Gehrke, and C. G. Plaxton. A proportional share resource allocation algorithm for real-time, time-shared systems. In *Real-Time Systems Symposium*. IEEE, December 1996.

[23] R. West and K. Schwan. Dynamic window-constrained scheduling for multimedia applications. In *6th International Conference on Multimedia Computing and Systems, ICMCS'99*. IEEE, June 1999. Also available as a Technical Report: GIT-CC-98-18, Georgia Institute of Technology.

[24] R. West, K. Schwan, and C. Poellabauer. Scalable scheduling support for loss and delay constrained media streams. In *Proceedings of the 5th IEEE Real-Time Technology and Applications Symposium*. IEEE, June 1999.

[25] H. Zhang and S. Keshav. Comparison of rate-based service disciplines. In *Proceedings of ACM SIGCOMM*, pages 113–121. ACM, August 1991.