

A Virtual Deadline Scheduler for Window-Constrained Service Guarantees

Yuting Zhang, Richard West and Xin Qi

Computer Science Department
Boston University
Boston, MA 02215
{danazh,richwest,xqi}@cs.bu.edu

Abstract

This paper presents a new approach to window-constrained scheduling, suitable for multimedia and weakly-hard real-time systems. We originally developed an algorithm, called Dynamic Window-Constrained Scheduling (DWCS), that attempts to guarantee no more than x out of y deadlines are missed for real-time jobs such as periodic CPU tasks, or delay-constrained packet streams. While DWCS is capable of generating a feasible window-constrained schedule that utilizes 100% of resources, it requires all jobs to have the same request periods (or intervals between successive service requests). We describe a new algorithm called Virtual Deadline Scheduling (VDS), that provides window-constrained service guarantees to jobs with potentially different request periods, while still maximizing resource utilization.

VDS attempts to service m out of k job instances by their virtual deadlines, that may be some finite time after the corresponding real-time deadlines. Notwithstanding, VDS is capable of outperforming DWCS and similar algorithms, when servicing jobs with potentially different request periods. Additionally, VDS is able to limit the extent to which a fraction of all job instances are serviced late. Results from simulations show that VDS can provide better window-constrained service guarantees than other related algorithms, while still having as good or better delay bounds for all scheduled jobs. Finally, an implementation of VDS in the Linux kernel compares favorably against DWCS for a range of scheduling loads.

1. Introduction

The ubiquity of the Internet has led to widespread delivery of content to the desktop. Much of this content is now stream-based, such as video and audio, having quality of service (QoS) constraints in terms of throughput, delay, jit-

ter and loss. More recently, developments have focused on large-scale distributed sensor networks and applications, to support the delivery of QoS-constrained data streams from sensors to specific hosts [11], hand-held PDAs and even actuators. Many stream-based applications can tolerate late or lost data delivery as long as a minimum fraction is guaranteed to reach the destination in a timely fashion. However, there are constraints on which pieces of the data can be late or lost. For example, the loss of too many consecutive packets in a video stream sent over a network, might result in significant picture breakup rather than a tolerable reduction in signal-to-noise ratio.

While stream-based applications are often tolerant of late or lost information, other real-time applications (e.g., in embedded systems) are sometimes capable of functioning at acceptable levels even when a number of tasks are executed late or not at all. For example, a CPU-bound task that must sample and process sensor data may skip some samples as long as the minimum sampling rate is above a certain threshold.

To deal with the above classes of applications, we have developed a number of ‘window-constrained’ scheduling algorithms. Window-constrained scheduling is a form of weakly-hard [3, 4] service, in which a minimum number of consecutive job instances (e.g., periodic tasks or consecutive packets in a real-time stream) must be processed by their deadlines in every finite window. Similar to the assumption in the Lui and Layland model [10], every job J_i is periodic, with a constant request period, T_i , for each job instance $J_{i,j}$.

One such algorithm we developed in prior work is Dynamic Window-Constrained Scheduling (DWCS) [16, 14, 15]. DWCS attempts to guarantee no more than x out of a fixed *window* of y deadlines are missed for consecutive job instances. DWCS is capable of guaranteeing a feasible schedule for each job, J_i , such that no more than x_i out of y_i instances of J_i are serviced late, or skipped, as long as the total utilization of all required job instances does not exceed 100%. However, DWCS is only capable of guaran-

teering a feasible schedule when all jobs have the same request period. Although this seems restrictive, a similar constraint applies to pinwheel schedulers [7, 5, 1], and it can be shown by careful manipulation of x_i and y_i that minimum fractions of service are guaranteed for each J_i in finite and tunable windows of time.

Mok and Wang extended our original work by showing that the *general* window-constrained problem is NP-hard for arbitrary service time and request periods [12]. While they also developed a solution to the window-constrained scheduling problem for unit service time and arbitrary request periods, it is only capable of guaranteeing a feasible schedule when resources are utilized up to 50%. This has prompted us to devise a new algorithm, called Virtual Deadline Scheduling (VDS), that guarantees resource shares to a specific fraction of all job instances, even when resources are 100% utilized and request periods *differ* between jobs. That is, two jobs J_i and J_j , may have different request periods, T_i and T_j ¹.

In order to generate a feasible schedule for the window-constrained problem, both the request deadlines and window-constraints of jobs must be considered. Instead of considering these two factors separately as in DWCS, VDS combines them together to determine a *virtual deadline* that is used to order job instances. Virtual deadlines are set at specific points within a window of time, to ensure each job is given a proportionate share of service. Unlike other approaches that attempt to provide proportional sharing of resources, VDS dynamically adjusts virtual deadlines as the urgency of servicing a job changes.

From experimental results, VDS is able to outperform other algorithms that attempt to satisfy the window-constrained scheduling problem. However, VDS is specifically designed to satisfy a relaxed form of the window-constrained scheduling problem, in which m out of k job instances must be serviced by their virtual (as opposed to real) deadlines. In effect, this guarantees a fraction of resource usage to each job over a finite interval of time, while bounding the delay of each job instance. Although a job instance may miss its real deadline, VDS is still able to ensure a minimum of m job instances are serviced in a specific window of time. This is suitable for applications that can tolerate some degree of delay up to some maximum amount. In fact, VDS imposes the same delay bounds on jobs serviced according to both the relaxed and original window-constrained scheduling models. These properties of VDS make it suitable for a number of multimedia applications and those supported by weakly-hard real-time systems.

The contributions of this paper can now be summarized as follows:

- We present a relaxed (m, k) window-constrained model, that is appropriate for many classes of applications, such as multimedia streaming and real-time data sampling.
- We present a new algorithm, called virtual deadline scheduling (VDS), that combines a job's *period-based deadline* and *window-constraint* to determine the scheduling order. We show how VDS can make full use of resources, while still managing to service n jobs, such that each job J_i is serviced at least m_i times every non-overlapping window of $k_i T_i$ real-time. Moreover, we ensure that each job instance is serviced in keeping with its delay constraints, defined by a virtual deadline. For jobs with different request periods, VDS outperforms DWCS and similar algorithms for the original window-constrained scheduling problem, that requires m out of k deadlines to be met. Additionally, VDS can also guarantee service in the relaxed model up to 100% utilization, for jobs with different request periods.
- We compare the performance of VDS to algorithms such as DWCS, Eligibility-based Window-Deadline-First (EWDF), and EDF based on Pfair scope [12], using a series of simulations. In these simulations, a set of random jobs are serviced according to various window-constraints and arbitrary request periods.

The rest of the paper is organized as follows. In the next section, we define the window-constrained scheduling problem, in both its original and relaxed forms. The VDS algorithm and an analysis of its characteristics are then described in Section 3. In Section 4, we simulate the performance of VDS, and compare it with other window-constrained scheduling algorithms. Additionally, we show the performance of VDS for real-time workloads when operating as a CPU scheduler in the Linux kernel. This is followed by a description of related work in Section 5. Finally, conclusions and future work are described in Section 6.

2. Window-Constrained Scheduling

We originally developed the DWCS algorithm to address the window-constrained scheduling problem. Given a set of n periodic jobs, J_1, \dots, J_n , a valid window-constrained schedule requires at least m_i out of k_i instances of a job J_i to be serviced by their deadlines. Deadlines of consecutive job instances are assumed to be separated by request periods of size T_i , for each job J_i , as in Rate Monotonic scheduling [10]. One can think of a job instance's request period as the interval between when it is ready and when it must complete service for a specific amount of time. Moreover, the ready time of one job instance is also the deadline of previous job instance. Therefore, the request period T_i is also the interval between deadlines of successive instances of J_i . Thus, if the j th instance of J_i is denoted by $J_{i,j}$, then

¹Note that we assume the request period, T_i , for job J_i is a constant but this is more for analysis reasons than any implicit restriction on the algorithm.

the deadline of $J_{i,j}$ is $d_{i,j} = d_{i,j-1} + T_i$.

We assume that every instance of J_i has the same service time requirement, C_i ². This implies that a window-constrained schedule must (a) ensure at least m_i instances of J_i are serviced by their respective deadlines, and (b) the minimum service share for J_i is $m_i C_i$ time units every window of $k_i T_i$ time.

In prior work, the assumption was that each window of $k_i T_i$ time, or k_i deadlines spaced apart by T_i time units, was non-overlapping with a previous or successive window. This differs from the pinwheel scheduling model that considers windows to be sliding. While we can transform window-constraints (m, k) to their equivalent values $(m, 2k - m)$ for sliding windows, we will assume windows are non-overlapping throughout the rest of this paper.

Based on the above, a window-constrained job, J_i , is defined by a 4-tuple (C_i, T_i, m_i, k_i) . A minimum of m_i out of k_i consecutive job instances must be serviced for C_i time in every window time of $k_i T_i$ for each job J_i with request period T_i . This implies the minimum utilization factor of each job J_i is $U_i = \frac{m_i C_i}{k_i T_i}$. Additionally, the minimum required utilization for a set of n periodic jobs is $U_{min} = \sum_{i=1}^n \frac{m_i C_i}{k_i T_i}$. When the system is overloaded, the total resource utilization $U = \sum_{i=1}^n \frac{C_i}{T_i} > 1.0$, and it is therefore impossible to service every instance of all n jobs. However, if the minimum required utilization $U_{min} \leq 1.0$, a feasible window-constrained schedule *may* exist.

It can be shown that a feasible window-constrained schedule *must* exist if each and every job J_i meets m_i deadlines every $k_i T_i$ window of time during the hyper-period of size $lcm(k_i T_i)$. However, the general window-constrained problem with *arbitrary* service times and request periods has been shown to be NP-hard [12]. With arbitrary service times, it is not possible to guarantee a feasible window-constrained schedule for all job sets even if the minimum required utilization $U_{min} \leq 1.0$. Figure 1 shows an example job set for which a feasible window-constrained schedule cannot be produced. It should be clear that J_1 and J_3 cannot both satisfy their window-constraints. However, if the service time of each and every job instance is constant, and all request periods are a fixed multiple of this constant, then a feasible window-constrained schedule exists even when $U_{min} \leq 1.0$ [16].

Relaxing the window-constrained scheduling problem:

If we consider a schedule that starts at time, $t = 0$, then J_i requires service for at least $m_i C_i$ units of time by $t = k_i T_i$. However, as stated earlier, each job instance must be serviced for C_i time units in the current request period. This prevents J_i from receiving a continuous burst of service of $m_i C_i$ units from $t = k_i T_i - m_i C_i$ to $t = k_i T_i$. In effect,

² C_i can be thought of as the worst-case execution time of any instance of J_i .




Figure 1. Example of an infeasible window-constrained schedule when service times are not all the same.

a window-constrained schedule prevents large bursts of service to one job at the cost of others. However, a relaxed version of the problem, in which job instances may be serviced after their deadlines as long as a job receives at least $m_i C_i$ units of service every interval $k_i T_i$ may be acceptable for some real-time applications. This is true for many multimedia applications, and those which can tolerate a bounded delay, as long as they receive a minimum fraction of service in fixed time intervals. For example, packets carrying multimedia data streams can experience finite buffering delays before transmission, or processing at a receiver.

This has prompted us to relax the original window-constrained problem, to allow job instances to be serviced after their deadlines as long as we guarantee a minimum fraction of service to a job. As will be seen later, Virtual Deadline Scheduling (VDS) attempts to guarantee a feasible schedule according to these relaxed constraints. However, VDS still prevents a job being serviced entirely at the end of a window of size $k_i T_i$ time units, by spreading out where the m_i instances of a job must be serviced in that interval. In effect, VDS adopts a form of “proportional fair” scheduling of at least m_i instances of each job, J_i , every interval $k_i T_i$.

For clarification, Figure 2 shows the difference between the original and relaxed window-constrained scheduling problems. Case (a) describes the original window-constrained problem, in which at most one instance of a job, J_i , is serviced every request period. A feasible schedule results in service for a J_i in at least m_i out of k_i periods, every adjacent window of $k_i T_i$ time slots. Case (b) shows the relaxed window-constrained scheduling problem. Up to α instances of a given job can be serviced in a single period of size, T_i , if $\alpha - 1$ instances have missed their real-time deadlines in the current window of size $k_i T_i$. In case (b) of Figure 2, up to 2 instances of J_i can be serviced in period $T_{i,5}$, according to the relaxed window-constrained problem. However, case (c) shows that with the relaxed window-constrained scheduling model, only one job instance can be serviced in period $T_{i,4}$, because no deadlines have been missed in the current window.

In previous work, we show how the DWCS algorithm can meet window-constraints for n jobs when the mini-

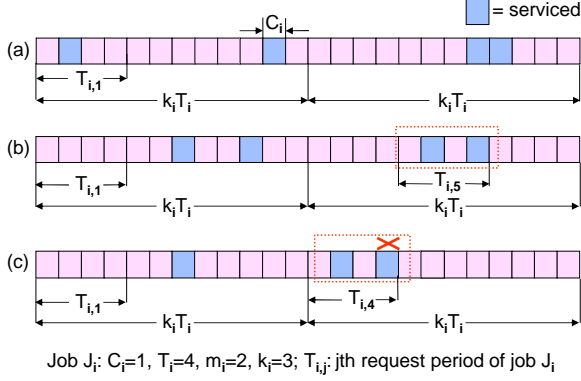


Figure 2. Original versus relaxed versions of the window-constrained scheduling problem.

minimum required utilization factor, $U_{min} = \sum_{i=1}^n \frac{m_i C_i}{k_i T_i} \leq 1.0$, if all service times are a constant, and request periods are a fixed multiple of this constant. That is, DWCS is capable of producing a feasible window-constrained schedule when resources are 100% utilized, if scheduling is performed at discrete time intervals, Δ , when $C_i = \Delta$ and $T_i = q\Delta$, for all i , such that $1 \leq i \leq n$ and q is a positive integer [16]. However, when jobs have different request periods, DWCS may not generate a feasible schedule even if $U_{min} < 1.0$. This has motivated us to develop the VDS algorithm, to provide service guarantees to jobs with potentially different request periods, while maximizing resource utilization.

3. Virtual Deadline Scheduling

Virtual deadline scheduling (VDS) is designed to provide service guarantees according to the relaxed form of the window-constrained scheduling problem. However, it is able to outperform algorithms such as DWCS for the original window-constrained problem, when jobs have different request periods. VDS derives “virtual deadlines” for each job instance from the corresponding window-constraint and request period, and the job instance with the earliest such deadline is scheduled first. In effect, a virtual deadline is used to loosely enforce proportional fairness on the service granted to a job in a specific window of time. This means the amount of service currently granted to a job in a specific window of real-time should be proportional to the minimum fraction of service required in the entire window.

3.1. Virtual Deadlines

Earliest-deadline-first (EDF) scheduling has the property that it can guarantee all deadlines will be met if resource usage does not exceed 100%. However, when the total utilization exceeds 100%, it is impossible for *any* schedule to

meet every deadline. With window-constrained scheduling, strategic deadlines may be missed, so that a minimum of m_i out of k_i deadlines are met every non-overlapping window of $k_i T_i$ real-time. As a result, it may be possible for the total utilization for n jobs, $\sum_{i=1}^n \frac{C_i}{T_i}$, to exceed 1.0, while the minimum utilization to guarantee a feasible schedule, $U_{min} = \sum_{i=1}^n \frac{m_i C_i}{k_i T_i}$ is less than or equal to 1.0. To produce such a feasible schedule, job instances must be prioritized using a combination of deadlines (based on their corresponding request periods) and window-constraints.

As stated above, VDS gives precedence to the job instance with earliest virtual deadline. A job’s virtual deadline with respect to real-time, t , is shown in Equation 1. T'_i is the remaining time in the current request period for J_i . $(t + T'_i - T_i)$ is the start time of current request period, which we denote as $t_r(t)$ for brevity. Similarly, (m'_i, k'_i) represent the *current* window-constraint at time, t . This implies that window-constraints change dynamically, depending on whether or not a job instance is serviced by its deadline.

The exact rules that control the dynamic adjustment of window-constraints will be described later. At this point, it is worth outlining the intuition behind a job’s virtual deadline. If at time t , J_i ’s current window-constraint is (m'_i, k'_i) , then $m_i - m'_i$ out of $k_i - k'_i$ job instances have been serviced. There are still m'_i job instances that need to be serviced in the remaining time in the current window, which is $k'_i T_i$. If one instance of J_i is serviced every interval $\frac{k'_i T_i}{m'_i}$, then m'_i job instances will be serviced in the current remaining window-time, $k'_i T_i$. A side-effect of this is that J_i is assured proportional fairness guarantees with respect to other window-constrained jobs. Additionally, the delay bound is minimized, by preventing at least m_i instances of J_i being serviced in a single burst at the end of a given real-time window.

$$\begin{aligned} Vd_i(t) &= \frac{k'_i T_i}{m'_i} + t_r(t) \\ t_r(t) &= (t + T'_i - T_i) \end{aligned} \quad (1)$$

Figure 3 gives an example of the virtual deadline calculation. We can see that, if a job’s current window-constraint does not change within a request period, its virtual deadline will not change either. This example corresponds to the relaxed window-constrained model, where more than one job instance can be served in one request period.

3.2. The VDS Algorithm

Although VDS gives precedence to the job with the earliest virtual deadline, it will only do so if that job is eligible for service. There are several cases that preclude a job from being scheduled, even when it has the lowest virtual deadline:




Figure 3. Example showing how to calculate virtual deadlines.

1) A job instance cannot be serviced before the start of its request period, even if it arrives early for service. It follows that if all currently available instances of a job have been serviced, the job is ineligible until a new arrival is ready.

2) If J_i has been serviced at least m_i times in its current window, it is given lower priority than a job yet to meet its window-constraint. Only if all jobs have achieved their minimum level of service can they again be considered in their current windows.

When a job is serviced its current window-constraint is adjusted. Job J_i has an original window-constraint of (m_i, k_i) that is set to a current value of (m'_i, k'_i) , to reflect how many more instances require service in the remainder of the active window. Figure 4 shows how current window-constraints are updated.

Here, the assumption is that scheduling decisions and service constraint adjustments are made once every time-slot, Δ . Unless stated otherwise, we assume throughout the rest of the paper that Δ represents a unit time-slot. Additionally, we assume the service time, C_i , of each and every job, J_i is the same as Δ . In Figure 4, C'_i and T'_i represent the remaining service time, and time remaining in the current request period of J_i , respectively. Every time job J_i is serviced, its remaining service time, C'_i , is decremented by Δ . At the start of a new request period J_i 's remaining service time, C'_i , is reset to its original value, C_i . If C'_i decreases to 0, J_i is ineligible for service until the start of the next request period.

At every scheduling point, J_i 's remaining time in its request period, T'_i , is decreased by Δ . If T'_i reaches 0, it is reset to the original value, T_i . Since the current time, t , increases by Δ every time the scheduler is activated, we need only update the value of t_r in Equation 1 once every request period, to determine J_i 's new virtual deadline, Vd_i .

The last few lines of the pseudo-code in Figure 4 show how constraint adjustments differ between the relaxed and original models. In the relaxed model, if there are outstanding job instances in the previous request period of the current window, C'_i is reset. In the original model, C'_i is reset

```

if (( $C'_i == 0$ ) || ( $m'_i \leq 0$ ))
    job  $J_i$  is ineligible for service ;

Serve eligible job  $J_i$  with lowest virtual deadline & update  $m'_i, C'_i$ :
 $C'_i = C'_i - \Delta$ ;
if ( $C'_i == 0$ )
     $m'_i = m'_i - 1$  ;

For every job  $J_j$ , check violations and update constraints:
if (( $Vd_j \leq \Delta + t$ ) && ( $j \neq i$ ))
    Tag  $J_j$  with a violation;
 $T'_j = T'_j - \Delta$ ;
if ( $T'_j == 0$ )
     $k'_j = k'_j - 1$ ;  $C'_j = C_j$ ;  $T'_j = T_j$ ;
if ( $k'_j == 0$ ) {
     $m'_j = m_j$ ;  $k'_j = k_j$ ;
    Discard the remaining job instances in the previous window
}
if ( $m'_i > 0$ )
    Update  $Vd_j$  according to Equation 1

// Only for the relaxed model
if ((( $k_j - k'_j$ )  $\geq$  ( $m_j - m'_j$ )) && ( $C'_j == 0$ ))
     $C'_j = C_j$ ;

```

Figure 4. Updating service constraints using VDS.

only at the beginning of each request period. This makes the relaxed window-constrained model more flexible, increasing the potential for more job instances to be serviced over time.

As stated above, the current window-constraint (m'_i, k'_i) is dynamically adjusted according to the service received by J_i . When an instance of J_i is serviced, m'_i is decreased by 1, because fewer instances need to be serviced in current window. If m'_i reaches 0 in the current window, J_i has met its window-constraint and becomes ineligible for service until the start of the next window, unless all other jobs have met their current window-constraints. The value of k'_i is decreased by 1 every request period, T_i , until it reaches 0, which indicates the end of the current window. At this point, J_i 's current window-constraint (m'_i, k'_i) is reset to its original value, (m_i, k_i) . A window-constraint violation is observed if any job instance misses its virtual deadline.

3.3. VDS versus Other Algorithms

The Earliest-deadline-first (EDF) algorithm produces a schedule that meets all deadlines, if such a schedule is known to theoretically exist. For the window-constrained scheduling problem, if each job J_i requires that $m_i = k_i$, then every real-time deadline must be met. In this case, the virtual deadlines of job instances serviced by VDS are the same as their corresponding real-time deadlines. In effect, VDS and EDF behave the same when $m_i = k_i$ for each J_i .

This implies that VDS shares the same optimal characteristics of EDF, when it is possible to meet all deadlines. Now, when $m_i = 1$ for each and every J_i , virtual deadlines using VDS are at the end of the current request window of size $k_i T_i$. Here, VDS behaves the same as an EDF scheduler for jobs with request periods of length $k_i T_i$. Furthermore, when k_i is a multiple of m_i for each and every J_i , the corresponding window-constraint can be reduced to $(1, \frac{k_i}{m_i})$. Once again, this is equivalent to servicing jobs using EDF with deadlines at the ends of periods of length $\frac{k_i T_i}{m_i}$.

DWCS was our first algorithm designed explicitly to support jobs with window-constraints. In ordering jobs for service, DWCS compares deadlines and window-constraints *separately*. In one version of the algorithm [16, 14], DWCS first compares the deadlines of jobs, giving precedence to the one with the earliest such deadline. If two or more jobs have the earliest deadline, their *current* window-constraints are then compared. In this case, the job, J_i , with the highest ratio, $\frac{m'_i}{k'_i}$, is given precedence. It can be shown that if all jobs have the same request periods, DWCS can generate a feasible window-constrained schedule, even when $U_{min} = 1.0$. This implies that a feasible window-constrained schedule is possible even when all resources (e.g., CPU cycles) are utilized.

Comparing VDS to DWCS, if all request periods are equal, then each job's virtual deadline only depends on its *current* window-constraint. Moreover, if all jobs have the same request periods then their current instances have the same real-time deadlines. In this case, DWCS will give precedence to the job with the highest value of $\frac{m'_i}{k'_i}$. Likewise, VDS will select the job with highest ratio $\frac{m'_i}{k'_i}$, since (from Equation 1) it has the earliest virtual deadline. Consequently, VDS is also able to produce a feasible window-constrained schedule that utilizes 100% of resources when all job request periods are equal.

Now, when jobs have *different* request periods and window-constraints, DWCS may fail to produce a valid schedule. As an example, consider the job set in Figure 5, with a minimum required utilization, $U_{min} = \sum_{i=1}^n \frac{m_i C_i}{k_i T_i} = \frac{8}{9} < 1.0$. The figure shows a number of schedules during the hyper period $(0, 9]$, for four different algorithms. As can be seen, J_3 cannot be scheduled in the first window using either EDF or DWCS, so it violates its window-constraint. Observe that EDF and DWCS both choose J_1 first, because it has the earliest deadline, rather than J_2 or J_3 that have "tighter" window-constraints. In contrast, VDS produces a schedule that satisfies the service constraints of all jobs. This is because VDS *combines* deadlines and window-constraints to derive a virtual deadline and, hence, priority for ordering jobs.

We stated earlier that VDS is designed to explicitly service jobs according to the relaxed form of window-

constrained scheduling. That is, it guarantees each job, J_i , receives at least $m_i C_i$ service time every window $k_i T_i$, rather than meeting at least m_i out of k_i deadlines. However, Figure 5 shows that it can outperform algorithms such as DWCS, according to the original window-constrained scheduling problem, when job request periods differ. Now, with the relaxed model, jobs may be serviced late as long as a minimum number of instances are scheduled in a given window. Thus, we can consider every job instance scheduled in the same window as having a common deadline at the end of that window. By setting deadlines at the ends of windows, an alternative to VDS is to use a deadline-driven scheduler that we call "Eligibility-based Window-Deadline-First" (EWDF). It behaves similar to EDF but gives precedence to the job with the earliest window deadline that is eligible for service. Section 3.2 describes the two conditions preventing a job from being eligible for service. With EWDF, k_i instances of J_i all have deadlines at the end of the current window of size $k_i T_i$, rather than each instance having a separate deadline at the end of its request period. As can be seen from Figure 5, EWDF is able to service all three jobs according to their window-constraints.

In general, EWDF is able to guarantee $m_i C_i$ units of service every $k_i T_i$ for each job J_i , if $U_{min} \leq 1.0$. However, it may delay the service of a job until the end of a window, $k_i T_i$. In the worst case, all m_i instances of J_i may be serviced in a single burst during the last $m_i \Delta$ time units in the current window. Hence, the worst-case delay of a job instance with EWDF is $(k_i T_i - m_i C_i + T_i - C_i)$. This compares to the maximum delay with VDS of $(k_i - m_i + 1)T_i - C_i$, as shown in the next section.




Figure 5. A comparison of scheduling algorithms.

Figure 6 shows an example of the differences in delays experienced by jobs using the VDS and EWDF algorithms. Using EWDF, all three job instances for J_1 are serviced in the last request period of the current window. The maximum service delay is 24, and only the last job instance meets its request deadline. However, using VDS, the maximum delay is 13, and all 3 job instances are serviced in their own request periods. EWDF does not consider m_i ,




Figure 6. Example service delays for VDS versus EWDF.

but only window-size, $k_i T_i$, to decide the scheduling priority. Though EWDF is suitable for the relaxed window-constrained model, it is not suitable for the original model.

3.4. VDS Properties

This section describes some of the key properties of VDS. These are summarized as follows:

- If a feasible schedule exists, such that at any time no virtual deadlines are missed, then VDS ensures that the maximum delay of each job is bounded.
- If a feasible schedule exists, it follows that VDS guarantees each job a minimum share of service in a finite interval. This is based on the assumption that each job is serviced at the granularity of a fixed-sized time slot, Δ (i.e., $\forall i, C_i = \Delta$), and all request periods are multiples of such a time slot (i.e., $\forall i, T_i = q_i \Delta \mid q_i \in \mathbb{Z}^+$).
- If the minimum required utilization, U_{min} , is less than or equal to 1.0, and service times are all constant, then a feasible schedule is guaranteed using VDS.
- The algorithmic complexity of the VDS algorithm is a linear function of the number of jobs needing service, in the worst case.

Lemma 1. *If a feasible VDS schedule exists, the current window-constraint (m'_i, k'_i) of job J_i always satisfies the condition that $k'_i \geq m'_i$.*

Proof. The proof is by contradiction. We will show that if there exists a job J_i , whose current window-constraint is such that $k'_i < m'_i$, then there is a service violation in the VDS schedule.

If at some time there exists the condition $k'_i = m'_i - 1$, then in the previous request period, $k'_i = m'_i$, and J_i was not serviced. If we let t be the time at the beginning of the last Δ time units of the previous request period, then $T'_i = \Delta$ and J_i 's virtual deadline is:

$$\begin{aligned} Vd_i(t) &= \frac{k'_i}{m'_i} T_i + (t + T'_i - T_i) \\ &= T_i + t + T'_i - T_i = t + T'_i = t + \Delta; \end{aligned}$$

We know that J_i was not serviced in the interval $[t, t + \Delta]$, so there must be a violation according to the VDS algorithm.

Hence, by contradiction, if a feasible VDS schedule exists, the current window-constraint (m'_i, k'_i) of job J_i always satisfies the condition that $k'_i \geq m'_i$. \square

Delay Bound

Theorem 1. *If a feasible schedule exists, the maximum delay of service to a job, $J_i \mid 1 \leq i \leq n$, is $(k_i - m_i + 1)T_i - C_i$.*

Proof. From Lemma 1, we know that if a feasible VDS schedule exists, the current window-constraint (m'_i, k'_i) of job J_i at any time satisfies the condition $k'_i \geq m'_i$. Hence, if no instance of J_i has been serviced by the $(k_i - m_i + 1)$ th period of the current window, then $k'_i = m'_i = m_i$. An instance must be served during this period, otherwise $k'_i < m'_i$ in next period. This implies the worst case delay for J_i is $(k_i - m_i + 1)T_i - C_i$ in a feasible VDS schedule. \square

Service Share

Theorem 2. *If there is a feasible VDS schedule, every job has at least m_i instances serviced in each $k_i T_i$ window of real-time. Hence, the minimum service share of each job is $\frac{m_i C_i}{k_i T_i}$ in every request window.*

Proof. Again from Lemma 1, we know that if a feasible VDS schedule exists, the condition $k'_i \geq m'_i$ must hold. Now, in the last request period of a given window, $k'_i = 1$ and $m'_i \leq 1$ is true. If $m'_i = 1 = k'_i$, then an instance of J_i must be serviced in this last period of the window. If $m'_i \leq 0$ in the last period of a given window, then J_i has already been served at least m_i out of k_i times before the window has ended. Hence, each job, J_i , receives at least $\frac{m_i C_i}{k_i T_i}$ service in every request window. \square

Feasibility Test

Theorem 3. *If $U_{min} = \sum_{i=1}^n \frac{m_i C_i}{k_i T_i} \leq 1.0$, $C_i = \Delta$ and $T_i = q_i \Delta, \forall i \mid q_i \in \mathbb{Z}^+$ then VDS guarantees a feasible schedule according to the relaxed window-constrained model.*

Proof. For brevity we do not provide a rigorous proof. However, it involves a reduction to an equivalent EDF scheduling problem. Note that EDF 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.

In the equivalent EDF schedule, we must guarantee that n periodic jobs are each serviced for C_i units of time, every period $\frac{k_i T_i}{m_i}$. Now, if $\sum_{i=1}^n \frac{C_i}{(k_i T_i)/m_i} \leq 1.0$ then EDF guarantees all jobs will be serviced for C_i time units every period, $\frac{k_i T_i}{m_i}$. With VDS, we require a feasible schedule to have a minimum utilization of $\frac{m_i C_i}{k_i T_i}$. This is the same utilization as that in the equivalent EDF schedule.

In meeting the utilization requirement, VDS must guarantee every serviced instance of J_i (of which there must be at least m_i such instances) meets its virtual deadline with respect to the current time, t . Let us assume that $t = 0$ initially. At the beginning of the first request window, J_i 's virtual deadline is set to $\frac{k_i T_i}{m_i}$. This is the same as the deadline of the first instance of J_i in the equivalent EDF scheduling problem. Now, with VDS, virtual deadlines increase over time. Hence, if EDF can guarantee service to the first instance of J_i by time $t = \frac{k_i T_i}{m_i}$ then the first instance serviced by VDS must have a virtual deadline greater than or equal to this time when it is actually serviced.

The worst-case virtual deadline of each serviced job instance will not be earlier than the equivalent deadline in an EDF schedule. With the relaxed window-constrained scheduling model, job instances are not discarded after their request periods, so we need only select a minimum of m_i such instances for each J_i by the corresponding virtual deadlines. That is, at least one instance of J_i is serviced in a request window by the virtual deadline with respect to the current time.

The requirement that $C_i = \Delta$ and $T_i = q_i \Delta, \forall i \mid q_i \in \mathbb{Z}^+$ is imposed because we assume VDS makes scheduling decisions at the granularity of Δ -sized time-slots. This allows VDS to emulate the preemptive nature of EDF. \square

Algorithmic Complexity

Theorem 4. *The complexity of the VDS algorithm is $O(n)$, where n is the number of jobs requiring service.*

Proof. The VDS algorithm is based on virtual deadline ordering. The cost of ordering such deadlines can be $O(\log(n))$ if a heap structure is used. However, when VDS either services a job or switches to a new request period, it must update the corresponding virtual deadline. In the worst-case all n jobs require their virtual deadlines to be recalculated at the same time. This is an $O(1)$ operation on a per-job basis, implying that the scheduling complexity is $O(n)$ for n jobs. \square

4. Experimental Evaluation

4.1. Simulations

This section evaluates the performance of VDS, via a series of simulations comprising a total of 1,300,000 randomly generated job sets. We assume that all jobs in each

set are periodic with unit processing time, $\Delta = 1$, although they may have different request periods, $q_i \Delta \mid q_i \geq 1$. Each job J_i has a new instance arrive every request period, T_i , and a scheduling decision is made once every unit-length time slot, Δ . A range of minimum utilization factors, U_{min} , up to 1.3 are derived by randomly choosing the number of jobs in a set, as well as values for job request periods and window-constraints. Utilization factors are incremented in steps of 0.1, resulting in 13 such values with 100,000 jobs sets in each case. Scheduling is performed for each job set over its hyper-period, to capture all possible window-constraint violations. In each test case, VDS is compared to several other algorithms, and violations are determined for both the original and relaxed window-constrained scheduling problems.

Performance Metrics: The following metrics are defined to measure the performance of each algorithm:

- V_{test_s} : This is the number of simulation tests that violate the service requirements of each job, according to the *relaxed* window-constrained scheduling problem. That is, if there is any job J_i that has less than m_i instances serviced in any window of $k_i T_i$ real-time, the corresponding test violates the service requirements. It should be noted that one test consists of a schedule for all jobs in a single set, serviced over their entire hyper-period.
- V_{test_d} : This is the number of simulation tests that violate the service requirements of each job, according to the *original* window-constrained scheduling problem. That is, if there is any job J_i that has less than m_i job instances meeting their request deadlines in any window of $k_i T_i$ real-time, the corresponding test violates the service requirements.
- V_s : This is the total violation rate of all jobs, in all tests, that fail to be serviced at least m_i times in any window of $k_i T_i$ real-time.
- V_d : This is the total violation rate of all jobs, in all tests, that fail to meet at least m_i deadlines in any window of $k_i T_i$ real-time.

The violation rate of each job J_i is calculated as the ratio of the number of windows with violations in the hyper-period, to the number of windows in the hyper-period. For each J_i , the number of real-time windows in the hyper-period is $lcm(k_i T_i, \forall i) / k_i T_i$.

Original Window-Constrained Scheduling Problem: In the original window-constrained scheduling problem, each job instance must be serviced in its current request period, otherwise it will be late. If we assume late job instances are simply discarded, the number of instances that meet deadlines must be the same as the number that are serviced. In this case, a window-based service constraint is equivalent to a window-based deadline constraint. Therefore, $V_d = V_s$ and $V_{test_d} = V_{test_s}$.

U_{\min}	V_{test_d}, V_{test_s}			V_d, V_s		
	DWCS	EDF-Pfair	VDS	DWCS	EDF-Pfair	VDS
(0.0-0.1)	0	0	0	0	0	0
(0.1-0.2)	0	0	0	0	0	0
(0.2-0.3)	0	0	0	0	0	0
(0.3-0.4)	0	0	0	0	0	0
(0.4-0.5)	0	0	0	0	0	0
(0.5-0.6)	0	0	0	0	0	0
(0.6-0.7)	5	0	0	0.011045	0	0
(0.7-0.8)	130	0	0	1.146656	0	0
(0.8-0.9)	1206	0	0	12.50002	0	0
(0.9-1.0)	14555	77	14	340.4671	4.679056	0.6
(1.0-1.1)	100000	100000	100000	83917.59	79749.3281	102407.2
(1.1-1.2)	100000	100000	100000	220502.2	195115.125	260838.6
(1.2-1.3)	100000	100000	100000	326949.8	281281.25	378124.8

U_{\min}	VDS		EWDF		VDS		EWDF	
	V_{test_s}	V_{test_d}	V_{test_s}	V_{test_d}	V_s	V_d	V_s	V_d
(0.0-0.1)	0	0	0	0	0	0	0	0
(0.1-0.2)	0	0	0	0	0	0	0	0
(0.2-0.3)	0	0	0	0	0	0	0	0
(0.3-0.4)	0	0	0	0	6	0	0	0.05
(0.4-0.5)	0	1	0	272	0	0.002	0	3.3
(0.5-0.6)	0	28	0	3649	0	0.2	0	74.6
(0.6-0.7)	0	888	0	19429	0	13.5	0	804.5
(0.7-0.8)	0	9125	0	52097	0	192.5	0	5861.2
(0.8-0.9)	0	37422	0	77643	0	2190.1	0	31481.2
(0.9-1.0)	0	72610	0	89413	0	14991.5	0	122458.5
(1.0-1.1)	100000	100000	100000	100000	94860.29	138155.2	67690.34	336293.7
(1.1-1.2)	100000	100000	100000	100000	238094.8	292439.6	168082.1	385539.3
(1.2-1.3)	100000	100000	100000	100000	347534.4	403402.9	246201.8	421908.6

Figure 7. Comparisons of service violations for (a) the original, and (b) the relaxed window-constrained scheduling problem.

Figure 7(a) shows results for VDS versus DWCS and the EDF-Pfair algorithm, with respect to the original window-constrained scheduling problem. The latter EDF-Pfair algorithm is a form of EDF-based pfair scheduling, as described by Mok and Wang [12]. It can be seen that, in underload cases, VDS results in fewer violations than the other scheduling algorithms. Moreover, VDS only starts to show violations when the utilization is up to 0.9, when there are only 14 out of 100,000 tests which fail. Similarly, the violation rate for VDS is very small. Although the EDF-Pfair algorithm performs well, it is not as good as VDS. However, DWCS results in violations when the minimum utilization factor is only 0.6. Likewise, the number of violating test cases, and the violation rate are much larger with DWCS than VDS.

Relaxed Window-Constrained Scheduling Problem: For the relaxed window-constrained scheduling problem, each instance of job J_i can legitimately be serviced in the current window of size $k_i T_i$, even if a corresponding request deadline has passed. This means there can be less job instances meeting deadlines than are actually serviced. Therefore, $V_s \leq V_d$ and $V_{test_s} \leq V_{test_d}$.

Figure 7(b) shows results for VDS versus EWDF, with respect to the relaxed window-constrained scheduling problem. In this case, VDS and EWDF are able to guarantee no service violations up to 100% utilization. In the overload cases, VDS has more violations than EWDF, because it tries to provide (proportionally) fair service to every job. That is, VDS attempts to provide each job with at least $m_i C_i$ units of service time every $k_i T_i$, even though this is not possible. However, compared to EWDF, VDS has (1) better delay properties, as it attempts to service job instances earlier, and (2) has fewer deadline violations.

4.2. CPU Scheduling Experiments in Linux

We have implemented VDS as part of a CPU scheduler in the Linux 2.4.18 kernel, to evaluate its performance in a

working system. A Dell precision 330 workstation, with a single 1.4Ghz Pentium 4 processor, 256KB L2 cache and 512MB RDRAM is used to compare VDS and DWCS schedulers. The experimental setup is similar to that in prior studies involving DWCS in the Linux kernel [13]. In the results that follow, we used the Pentium timestamp counter to accurately measure elapsed clock cycles and, hence, scheduling performance.




Figure 8. Violations using VDS versus DWCS CPU schedulers in the Linux kernel.

Figure 8 compares the performance of VDS and DWCS in a real system, in terms of average violations per task³. In these experiments, a violation occurs when fewer than m out of k consecutive deadlines are met for periodic, preemptive CPU-bound tasks. Each task runs in an infinite loop but can be preempted every clock tick, or *jiffy*, to allow the scheduler to execute. In effect, one can think of a task as an infinite sequence of sub-tasks, each requiring one jiffy's

³For scheduling purposes, Linux treats both threads and processes as tasks.

worth (about $10mS$ on an Intel x86) of service every request period.

It should be noted that the x -axis of Figure 8 does not represent a linear scale. Rather, each data point represents the utilization, $U_{min} = \sum_{i=1}^n \frac{m_i C_i}{k_i T_i}$. These values are derived from a combination of up to $n = 8$ tasks, with randomly generated scheduling parameters m_i , k_i and T_i for each task. Since each task executes for one jiffy between scheduling points (discounting any system-processing overheads), we can assume that service times are all unit-length. As can be seen, when the utilization is less than 1.0, there are almost no window-constraint violations using VDS compared to DWCS. As expected, violations occur for both algorithms when U_{min} exceeds 1.0.

5. Related Work

Window-constrained scheduling is a form of weakly-hard service [3, 4], that is similar to the (m, k) -firm scheduling [6]. Hamdaoui and Ramanathan [6] were the first to introduce the notion of (m, k) -firm deadlines, in which statistical service guarantees are applied to jobs. Their algorithm uses a “distance-based” priority scheme to increase the priority of a job in danger of missing more than m deadlines over a window of k requests for service. In contrast, VDS uses a virtual deadline scheme to derive a job’s priority and ensure deterministic service guarantees. Similar to (m, k) -firm scheduling is the work by Koren and Shasha on ‘skip over’ scheduling [9]. Skip over scheduling allows certain job instances to be skipped, but may unnecessarily miss servicing a job instance when there are resources available.

There are also examples of (m, k) -hard schedulers [2] but most such approaches require off-line feasibility tests, to ensure predictable service. Additionally, our VDS algorithm is targeted at a specific window-constrained problem that requires explicit service of a minimum number (m_i) of instances of each job J_i in a window of $k_i T_i$ time units, such that strong delay bounds are met.

Other related research includes pinwheel scheduling [7, 5, 1] but all time intervals, and hence request periods, are of a fixed size. In essence, the generalized pinwheel scheduling problem is equivalent to determining a schedule for a set of n jobs $\{J_i \mid 1 \leq i \leq n\}$, 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. Both our previous DWCS algorithm, and VDS, can be thought of as special cases of pinwheel scheduling. With VDS, service guarantees are provided over non-overlapping windows of k_i deadlines spaced apart by T_i time units. However, VDS guarantees feasibility when resources are 100% utilized, even when k_i is finite and different jobs have arbitrary request periods.

Finally, Jeffay and Goddard’s Rate-Based Execution (RBE) model [8] attempts to service jobs at an average rate of x times every y time units. In essence, this is similar to the window-constrained guarantee using VDS, that ensures a minimum service time of $m_i C_i$ every window of $k_i T_i$ units of real-time (given that C_i is a constant for each and every job, J_i). However, RBE does not consider that a certain fraction of job instances can be late or discarded.

6. Conclusions and Future Work

We originally addressed the window-constrained scheduling problem in our previous work on the DWCS algorithm. DWCS attempts to guarantee that no more than x out of y deadlines are missed for real-time periodic CPU tasks, or consecutive packets in delay-constrained streams, when all request periods are identical and resources are 100% utilized. This paper describes an extension to our previous work on window-constrained scheduling. We propose a relaxed version of the window-constrained problem, in which m out of k job instances must be serviced within a finite window of time, possibly after their real-time deadlines. Without loss of generality and applicability, feasible schedules are possible for a wider range of jobs under the relaxed window-constrained scheduling problem. This is because m out of k jobs may be serviced late, as long as they are serviced in a bounded, and specific, window of time.

The Virtual Deadline Scheduling (VDS) algorithm described in this paper is specifically aimed at providing service to real-time jobs, based on the relaxed window-constrained scheduling problem. That is, rather than guaranteeing at least m out of k real-time deadlines are met for consecutive instances of each job, it is only necessary to provide service to m instances every window of k request periods. With VDS, at least m out of k job instances are serviced by their *virtual* deadlines, even when resources are 100% utilized. Although such deadlines may be after their corresponding real-time deadlines, VDS is able to limit the extent to which job instances are serviced late. As a result, VDS limits the delay imposed on each instance of a job. Additionally, VDS is capable of outperforming DWCS and other similar algorithms for the original window-constrained scheduling problem when jobs have arbitrary request periods. In many cases, this makes VDS a more flexible algorithm, since it places fewer restrictions on the service specifications of window-constrained jobs.

Future work involves the development and analysis of algorithms such as VDS and DWCS, to provide end-to-end service guarantees across multi-hop networks. Such algorithms will feature as part of our work on the construction of a scalable distributed system of end-hosts for processing and delivering live data streams in real-time.

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] 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.
- [3] G. Bernat, A. Burns, and A. Llamosi. Weakly-hard real-time systems. *IEEE Transactions on Computers*, 50(4):308–321, April 2001.
- [4] G. Bernat and R. Cayssials. Guaranteed on-line weakly-hard real-time systems. In *Proceedings of the 22nd IEEE Real-Time Systems Symposium*, December 2001.
- [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] M. Hamdaoui and P. Ramanathan. A dynamic priority assignment technique for streams with (m,k)-firm deadlines. *IEEE Transactions on Computers*, April 1995.
- [7] R. Holte, A. Mok, L. Rosier, I. Tulchinsky, and D. Varvel. The pinwheel: A real-time scheduling problem. In *Proceedings of the 22nd Hawaii International Conference of System Science*, pages 693–702, Jan 1989.
- [8] K. Jeffay and S. Goddard. A theory of rate-based execution. In *Proceedings of the 20th IEEE Real-Time Systems Symposium (RTSS)*, December 1999.
- [9] 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.
- [10] 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.
- [11] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. TAG: A tiny aggregation service for ad-hoc sensor networks. In *Proceedings of Operating Systems Design and Implementation*. USENIX, December 2002.
- [12] A. K. Mok and W. Wang. Window-constrained real-time periodic task scheduling. In *Proceedings of the 22st IEEE Real-Time Systems Symposium*, 2001.
- [13] R. West, I. Ganey, and K. Schwan. Window-constrained process scheduling for linux systems. In *the Third Real-Time Linux Workshop*, November 2001.
- [14] R. West and C. Poellabauer. Analysis of a window-constrained scheduler for real-time and best-effort packet streams. In *Proceedings of the 21st IEEE Real-Time Systems Symposium*, December 2000.
- [15] 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.
- [16] R. West, Y. Zhang, K. Schwan, and C. Poellabauer. Dynamic window-constrained scheduling of real-time streams in media servers. Accepted for publication in *IEEE Transactions on Computers*, 2004.