

# Aperiodic Task Scheduling for Hard-Real-Time Systems

BRINKLEY SPRUNT

*Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213*

LUI SHA

*Department of Computer Science and Software Engineering Institute, Carnegie Mellon University,  
Pittsburgh, PA 15213*

JOHN LEHOCZKY

*Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213*

**Abstract.** A real-time system consists of both aperiodic and periodic tasks. Periodic tasks have regular arrival times and hard deadlines. Aperiodic tasks have irregular arrival times and either soft or hard deadlines. In this article, we present a new algorithm, the Sporadic Server algorithm, which greatly improves response times for soft deadline aperiodic tasks and can guarantee hard deadlines for both periodic and aperiodic tasks. The operation of the Sporadic Server algorithm, its performance, and schedulability analysis are discussed and compared with previously published aperiodic service algorithms.

## 1. Introduction: the real-time scheduling problem

Real-time systems are used to control physical processes that range in complexity from automobile ignition systems to controllers for flight systems and nuclear power plants. In these systems, the correctness of system functions depends upon not only the results of computation but also on the times at which results are produced.

A real-time task is generally placed into one of four categories based upon its arrival pattern and its deadline. If meeting a given task's deadline is critical to the system's operation, then the task's deadline is considered to be *hard*. If it is desirable to meet a task's deadline but occasionally missing the deadline can be tolerated, then the deadline is considered to be *soft*. Tasks with regular arrival times are called *periodic tasks*. A common use of periodic tasks is to process sensor data and update the current state of the real-time system on a regular basis. Periodic tasks, typically used in control and signal-processing applications, have hard deadlines. Tasks with irregular arrival times are *aperiodic tasks*. Aperiodic tasks are used to handle the processing requirements of random events such as operator requests. An aperiodic task typically has a soft deadline. Aperiodic tasks that have hard deadlines are called *sporadic tasks*. We assume that each task has a known worst-case execution time. In summary, we have

*Hard and soft deadline periodic tasks.* A periodic task has a regular interarrival time equal to its period and a deadline that coincides with the end of its current period. Periodic tasks usually have hard deadlines, but in some applications the deadlines can be soft.

*Soft deadline aperiodic tasks.* An aperiodic task is a stream of jobs arriving at irregular intervals. Soft deadline aperiodic tasks typically require a fast average response time.

*Sporadic tasks.* A sporadic task is an aperiodic task with a hard deadline and a *minimum* interarrival time (Mok 1983). Note that without a minimum interarrival time restriction, it is impossible to guarantee that a sporadic task's deadline would always be met.

To meet the timing constraints of the system, a scheduler must coordinate the use of all system resources using a set of well-understood real-time scheduling algorithms that meet the following objectives:

Guarantee that tasks with hard timing constraints will always meet their deadlines.

Attain a high degree of schedulable utilization for hard deadline tasks (periodic and sporadic tasks). Schedulable utilization is the degree of resource utilization at or below which all hard deadlines can be guaranteed. The schedulable utilization attainable by an algorithm is a measure of the algorithm's utility: the higher the schedulable utilization, the more applicable the algorithm is for a range of real-time systems.

Provide fast average response times for tasks with soft deadlines (aperiodic tasks).

Ensure scheduling stability under transient overload. In some applications, such as radar tracking, an overload situation can develop in which the computation requirements of the system exceed the schedulable resource utilization. A scheduler is said to be stable if during overload it can guarantee the deadlines of critical tasks even though it is impossible to meet all task deadlines.

The quality of a scheduling algorithm for real-time systems is judged by how well the algorithm meets these objectives.

This article develops advanced algorithms to schedule aperiodic tasks. For soft deadline aperiodic tasks, the goal is to provide fast average response times. For hard deadlines aperiodic tasks (sporadic tasks), the goal is to guarantee that their deadlines will always be met. The new algorithms presented here meet both of these goals and are still able to guarantee the deadlines of hard deadline periodic tasks. For simplicity, we assume that periodic tasks have hard deadlines and constant execution times.<sup>1</sup> In Section 2 we review related work on scheduling periodic and aperiodic tasks. Section 3 introduces the Sporadic Server (SS) algorithm for scheduling aperiodic tasks and illustrates its operation with several examples. This section also addresses the schedulability issues of the SS algorithm and compares its performance with previous algorithms. Section 4 addresses the scheduling of sporadic tasks and discusses the use of a *deadline monotonic* priority assignment for scheduling sporadic tasks with short deadlines and long interarrival times. Finally, Section 5 presents a summary and conclusions.

## 2. Background and related work

### 2.1. Scheduling periodic tasks

A well-understood scheduling algorithm for guaranteeing the hard deadlines of periodic tasks in a multiprogrammed real-time system is Liu and Layland's (1973) rate monotonic scheduling algorithm. Under this algorithm, fixed priorities are assigned to tasks based upon the rate of their requests (i.e., a task with a relatively short period is given a relatively high priority). Liu and Layland proved that this algorithm is the optimum static preemptive scheduling algorithm for periodic tasks with hard deadlines. The algorithm guarantees that  $n$  periodic tasks can always be guaranteed to meet their deadlines if the resource utilization of the tasks is less than  $n(2^{1/n}-1)$ , which converges to  $\ln 2$  (=69%) for large  $n$ . The rate monotonic algorithm can be used as a basis for developing a family of scheduling algorithms that addresses a wide range of practical problems. The rate monotonic algorithm has the following favorable qualities for scheduling real-time systems:

*High schedulable utilization.* Although Liu and Layland show a low scheduling bound for the rate monotonic algorithm, this bound is pessimistic and represents the absolute worst-case conditions. Lehoczky, Sha, and Ding (1987) performed an exact characterization and stochastic analysis for a randomly generated set of periodic tasks scheduled by the rate monotonic algorithm and found that the average scheduling bound is usually much better than the worst case. They concluded that a good approximation to the threshold of schedulability for the rate monotonic algorithm is 88%. In fact, with the period transformation method, the utilization threshold can, in principle, be arbitrarily close to 100% (Sha and Goodenough 1988). As an example of the high degree of schedulable utilization attainable with the rate monotonic algorithm, a schedulable utilization level of 99% was achieved for the Navy's Inertial Navigation System (Borger 1987).

*Stability under transient overload.* Another concern for scheduling algorithms is transient overload, the case where stochastic execution times can lead to a desired utilization greater than the schedulable utilization bound of the task set. To handle transient overloads, Sha, Lehoczky, and Rajkumar (1986) describe a period transformation method for the rate monotonic algorithm that can guarantee that the deadlines of critical tasks can be met.

*Aperiodic tasks.* A real-time system often has both periodic and aperiodic tasks. Lehoczky, Sha, and Strosnider (1987) developed the Deferrable Server algorithm, which is compatible with the rate monotonic scheduling algorithm and provides a greatly improved average response time for soft deadline aperiodic tasks over polling or background service algorithms while still guaranteeing the deadlines of periodic tasks.

*Resource sharing.* Although determining the schedulability of a set of periodic tasks that use semaphores to enforce mutual exclusion has been shown to be NP-hard (Mok 1986), Sha, Rajkumar, and Lehoczky (1987) have developed a priority inheritance protocol and derived a set of sufficient conditions under which a set of periodic tasks that share resources using this protocol can be scheduled using the rate monotonic algorithm.

*Low scheduling overhead.* Since the rate monotonic algorithm assigns a static priority to each periodic task, the selection of which task to run is a simple function. Scheduling algorithms that dynamically assign priorities may incur a larger overhead because task priorities have to be adjusted in addition to selecting the highest-priority task to execute.

Thus, we use the rate monotonic algorithm as the basis for scheduling soft deadline aperiodic tasks and sporadic tasks. In the next section, we review the related work on aperiodic scheduling.

## 2.2. Scheduling aperiodic tasks

The scheduling problem for aperiodic tasks is very different from that for periodic tasks. Scheduling algorithms for aperiodic tasks must be able to guarantee the deadlines for hard deadline aperiodic tasks and provide good average response times for soft deadline aperiodic tasks even though the occurrence of the aperiodic requests are nondeterministic. The aperiodic scheduling algorithm must also accomplish these goals without compromising the hard deadlines of the periodic tasks.

Two common approaches for servicing soft deadline aperiodic requests are background processing and polling tasks. Background servicing of aperiodic requests occurs whenever the processor is idle (i.e., not executing any periodic tasks and no periodic tasks pending). If the load of the periodic task set is high, then utilization left for background service is low, and background service opportunities are relatively infrequent. Polling consists of creating a periodic task for servicing aperiodic requests. At regular intervals, the polling task is started and services any pending aperiodic requests. However, if no aperiodic requests are pending, the polling task suspends itself until its next period and the time originally allocated for aperiodic service is not preserved for aperiodic execution but is instead used by periodic tasks. Note that if an aperiodic request occurs just after the polling task has suspended, then the aperiodic request must wait until the beginning of the next polling task period or until background processing resumes before being serviced. Even though polling tasks and background processing can provide time for servicing aperiodic requests, they have the drawback that the average wait and response times for these algorithms can be long, especially for background processing.

Figures 2 and 3 illustrate the operation of background and polling aperiodic service using the periodic task set presented in Figure 1. The rate monotonic algorithm is used to assign priorities to the periodic tasks yielding a higher priority for task A. In each of these examples, periodic tasks A and B both become ready to execute at time = 0. Figures 2 and 3 show the task execution from time = 0 until time = 20. In each of these examples, two aperiodic requests occur: the first at time = 5 and the second at time = 12.

The response time performance of background service for the aperiodic requests shown in Figure 2 is poor. Since background service only occurs when the resource is idle, aperiodic service cannot begin until time = 16. The response time of the two aperiodic requests are 12 and 6 time units respectively, even though both requests each need only 1 unit of time to complete.

The response time performance of polling service for the aperiodic requests shown in Figure 3 is better than background service for the first request but not for the second. For this example, a polling server is created with an execution time of 1 time unit and a period of 5 time units which, using the rate monotonic algorithm, makes the polling server the highest-priority task. The polling server's first period begins at time = 0. The lower part of Figure 3 shows the capacity (available execution time) of the polling server as a

Periodic Tasks:



	Execution Time	Period	Priority
 Task A	4	10	High
 Task B	8	20	Low

Figure 1. Periodic task set for Figures 2, 3, 4, and 5.

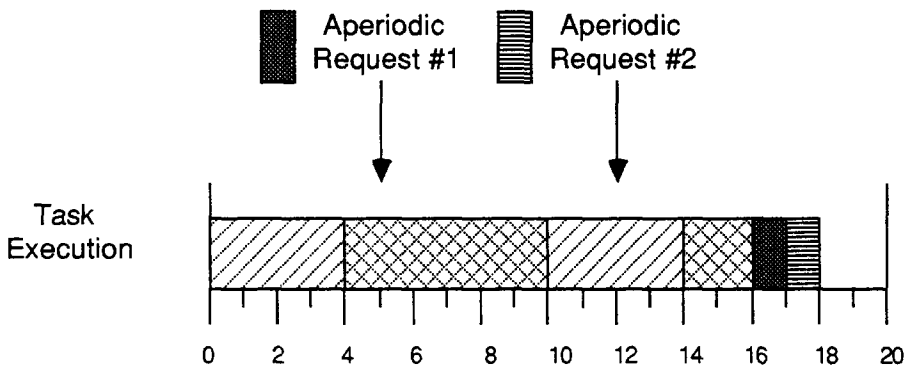


Figure 2. Background aperiodic service example.

function of time. As can be seen from the upward arrow at time = 0 on the capacity graph, the execution time of the polling server is discarded during its first period because no aperiodic requests are pending. The beginning of the second polling period coincides with the first aperiodic request, and so, the aperiodic request receives immediate service. However, the second aperiodic request misses the third polling period (time = 10) and must wait until the fourth polling period (time = 15) before being serviced. Also note that since the second aperiodic request only needs half of the polling server's capacity, the remaining half is discarded because no other aperiodic tasks are pending. Thus, this example demonstrates how polling can provide an improvement in aperiodic response time performance over background service but is not always able to provide immediate service for aperiodic requests.

The Priority Exchange (PE) and Deferrable Server (DS) algorithms, introduced by Lehoczky, Sha, and Strosnider (1987), overcome the drawbacks associated with polling and background servicing of aperiodic requests. As with polling, the PE and DS algorithms create a periodic task (usually of high priority) for servicing aperiodic requests. However, unlike polling, these algorithms preserve the execution time allocated for aperiodic service if, upon the invocation of the server task, no aperiodic requests are pending. These algorithms can yield improved average response times for aperiodic requests because of their ability




Figure 3. Polling aperiodic service example.

to provide immediate service for aperiodic tasks. In particular, the DS algorithm has been shown capable of providing an order of magnitude improvement in the responsiveness of asynchronous class messages for real-time token rings (Strosnider 1988). These algorithms are called *bandwidth preserving algorithms* because they provide a mechanism for preserving the resource bandwidth allocated for aperiodic service if, upon becoming available, the bandwidth is not immediately needed. The PE and DS algorithms differ in the manner in which they preserve their high-priority execution time.

The DS algorithm maintains its aperiodic execution time for the duration of the server's period. Thus, aperiodic requests can be serviced at the server's high priority at any time as long as the server's execution time for the current period has not been exhausted. At the beginning of the DS's period, the server's high-priority execution time is replenished to its full capacity.

The DS algorithm's method of bandwidth preservation is illustrated in Figure 4 using the periodic task set of Figure 1. For this example, a high-priority server is created with an execution time of 0.8 time units and a period of 5 time units. At time = 0, the server's execution time is brought to its full capacity. This capacity is preserved until the first aperiodic request occurs at time = 5, at which point it is immediately serviced, exhausting the server's capacity by time = 6. At time = 10, the beginning of the third server period, the server's execution time is brought to its full capacity. At time = 12, the second aperiodic request occurs and is immediately serviced. Notice that although the second aperiodic request only consumes half the server's execution time, the remaining capacity is preserved, not discarded as in the polling example. Thus, the DS algorithm can provide better aperiodic

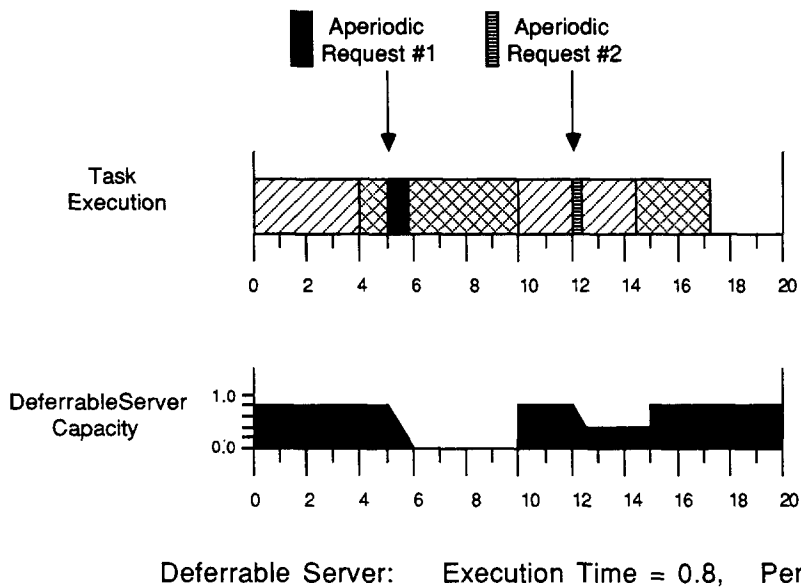


Figure 4. Deferrable server example.

responsiveness than polling because it preserves its execution time until it is needed by an aperiodic task.

Unlike the DS algorithm, the PE algorithm preserves its high-priority execution time by exchanging it for the execution time of a lower-priority periodic task. At the beginning of the PE server's period, the server's high-priority execution time is replenished to its full capacity. If the highest-priority execution time available is aperiodic time (as is the case at the beginning of the PE server's period) and aperiodic tasks are pending, then the aperiodic tasks are serviced. Otherwise, the highest-priority pending periodic task is chosen for execution and a priority exchange occurs. The priority exchange converts the high-priority aperiodic time to aperiodic time at the assigned priority level of the periodic task. When a priority exchange occurs, the periodic task executes at the priority level of the higher-priority aperiodic time, and aperiodic time is accumulated at the priority level of the periodic task. Thus, the periodic task advances its execution time, and the aperiodic time is not lost but preserved, albeit at a lower priority. Priority exchanges will continue until either the high-priority aperiodic time is exhausted or an aperiodic request occurs, in which case the aperiodic time is used to service the aperiodic request. Note that this exchanging of high-priority aperiodic time for low-priority periodic time continues until either the aperiodic time is used for aperiodic service or until the aperiodic time is degraded to the priority level of background processing (this complete degradation will occur only when no aperiodic requests arrive early enough to use the aperiodic time). Also, since the objective of the PE algorithm is to provide a low average response time for aperiodic requests, aperiodic requests win all priority ties. At all times the PE algorithm uses the highest-priority execution time available to service either periodic or aperiodic tasks.

The PE algorithm's method of bandwidth preservation is demonstrated in Figure 5 using the periodic task set of Figure 1. In this example, a high-priority PE server is created with

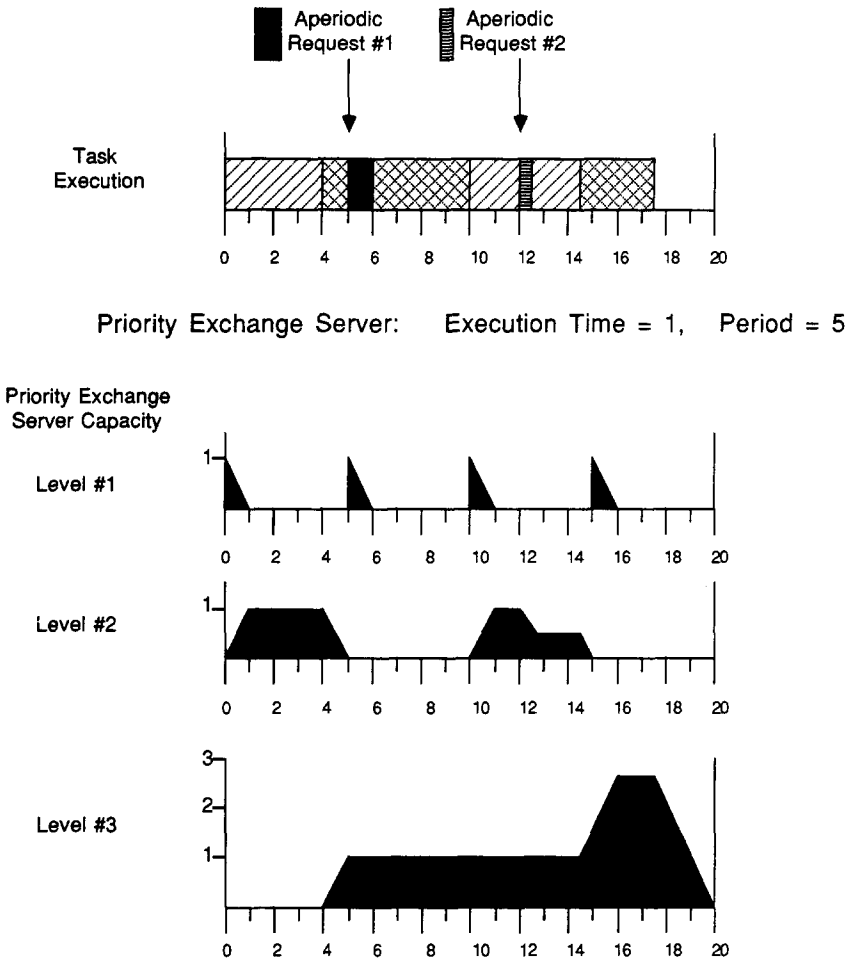


Figure 5. Priority Exchange Server example.

an execution time of 1 time unit and a period of 5 time units. Since the PE algorithm must manage aperiodic time across all priority levels, the capacity of the PE server as a function of time consists of three graphs: one for each priority level. The PE server's priority is priority level 1, which corresponds to the highest priority level, followed by priority 2 for periodic task A and priority 3 for periodic task B. At time = 0, the PE server is brought to its full capacity, but no aperiodic tasks are pending and a priority exchange occurs between priorities 1 and 2. The PE server gains aperiodic time at priority 2, and periodic task A executes at priority 1. At time = 4, task A completes and task B begins. Since no aperiodic tasks are pending, another exchange takes place between priority 2 and priority 3. At time = 5, the server's execution time at priority 1 is brought to its full capacity and is used to provide immediate service for the first aperiodic request. At time = 10, the server's priority 1 execution time is brought to full capacity and is then exchanged down to priority 2.



At time = 12, the server's execution time at priority 2 is used to provide immediate service for the second aperiodic request. At time = 14.5, the remaining priority 2 execution time is exchanged down to priority 3. At time = 15, the newly replenished server time at priority 1 is exchanged down to priority 3. Finally, at time = 17.5, the remaining PE server execution time at priority 3 is discarded because no tasks, periodic or aperiodic, are pending. Thus, the PE algorithm can also provide improved response times for aperiodic tasks compared to the polling algorithm.

The PE and DS algorithms differ in their complexity and in their effect upon the schedulability bound for periodic tasks. The DS algorithm is a much simpler algorithm to implement than the PE algorithm, because the DS algorithm always maintains its high-priority execution time at its original priority level and never exchanges its execution time with lower priority levels as does the PE algorithm. However, the DS algorithm does pay schedulability penalty (in terms of a lower schedulable utilization bound) for its simplicity. Both algorithms require that a certain resource utilization be reserved for high-priority aperiodic service. We refer to this utilization as the server size,  $U_s$ , which is the ratio of the server's execution time to the server's period. The server size and type (i.e., PE or DS) determine the scheduling bound for the periodic tasks,  $U_p$ , which is the highest periodic utilization for which the rate monotonic algorithm can always schedule the periodic tasks. Following are the equations developed by Lehoczky, Sha, and Strosnider (1987) for  $U_p$  in terms of  $U_s$  as the number of periodic tasks approaches infinity for the PE and DS algorithms:

$$\text{PE:} \quad U_p = \ln \frac{2}{U_s + 1} \quad \text{DS:} \quad U_p = \ln \frac{U_s + 2}{2U_s + 1}$$

These equations show that for a given server size,  $U_s$  ( $0 < U_s < 1$ ), the periodic schedulability bound,  $U_p$ , for the DS algorithm is lower than it is for the PE algorithm. These equations also imply that for a given periodic load, the server size,  $U_s$  for the DS algorithm is smaller than that for the PE algorithm. For example, with  $U_p = 60\%$ , the server size for the PE algorithm is 10% compared to a server size for the DS algorithm of 7%.

The focus of this article is to develop a general algorithm for scheduling both soft and hard deadline aperiodic tasks. Such an algorithm should provide a responsiveness comparable to that attainable with the PE and DS algorithms for soft deadline aperiodic tasks. For hard deadline aperiodic tasks, the algorithm should provide a general technique for allocating resource utilization to guarantee that the deadline will always be met.

By comparing the PE and DS algorithms we can identify the relative merits of each. The advantage of the DS algorithm over the PE algorithm is that it is conceptually much simpler, and thus easier to implement. The PE algorithm must manage aperiodic time across all priority levels in the system, whereas the DS algorithm maintains its execution time at its original priority level. The simple bandwidth preservation technique of the DS algorithm also implies that, on average, the priority of the server's execution time will be higher than it would be for an equivalently-sized PE server because the DS algorithm preserves its high-priority execution time at its original priority level for the duration of its period. In contrast, the PE algorithm must either use its high-priority execution time for aperiodic service or trade it for lower-priority periodic time. However, the PE algorithm

has a server size advantage over the DS algorithm. This advantage of the PE algorithm is even greater when multiple servers are executing at different priority levels. A better algorithm for soft deadline aperiodic tasks would have the advantages of both these algorithms while overcoming their limitations. Section 3 describes and gives examples of such an algorithm, the Sporadic Server (SS) algorithm. Section 3.4 presents the results of a simulation study that compares the response time performance of the polling, DS, PE, and SS algorithms.

One important class of aperiodic tasks not generally supported by previous aperiodic service algorithms consists of aperiodic tasks with hard deadlines. To guarantee hard deadlines for aperiodic tasks, a high-priority server can be created with enough execution time to guarantee that an aperiodic task can meet a hard deadline. However, a minimum interarrival time restriction must be placed upon the aperiodic task to ensure that the server will always have sufficient execution time to meet the aperiodic task's deadline. This minimum interarrival time for the aperiodic task must be equal to or greater than the period of its server task. As long as the aperiodic task does not request service more frequently than the minimum interarrival time, its hard deadline can be guaranteed. The SS algorithm can be used to provide such a guarantee, but only if the aperiodic task's deadline is equal to or greater than its minimum interarrival time. For the cases when the deadline is shorter than the minimum interarrival time, a priority assignment based upon the deadline of the aperiodic task rather than on its minimum interarrival time is necessary. This is referred to as a *deadline monotonic priority assignment* and requires a different schedulability analysis than that used for a rate monotonic priority assignment. The necessity of deadline monotonic priority assignment for hard deadline aperiodic tasks and the required schedulability analysis is discussed in Section 4.

### 3. Scheduling soft deadline aperiodic tasks

In this section, we investigate the scheduling of soft deadline aperiodic tasks. Our objective is to provide fast average response times for aperiodic tasks while guaranteeing the hard deadlines of periodic tasks. In this section, we introduce a new algorithm, determine its schedulability bound, and compare its response time performance to previous aperiodic service algorithms.

#### 3.1. The sporadic server algorithm

Although both the DS and PE algorithms have been shown to provide good average response times for aperiodic tasks (Lehoczky, Sha, and Strosnider 1987), these algorithms can be improved in several ways. The DS and PE algorithms have comparable aperiodic response time performance, but each algorithm has advantages over the other, as was discussed in Section 2.2. The advantage of the DS algorithm over the PE algorithm is that the DS algorithm always maintains its high-priority server execution time at its original priority level. As such, the DS algorithm has a much simpler implementation than the PE algorithm because the DS algorithm does not have to manage the exchanging of aperiodic execution

time between priority levels as is necessary with the PE algorithm. However, the PE algorithm has the advantage of a larger server size than the DS algorithm. A better algorithm would have an aperiodic response time performance comparable to the DS and PE algorithms, the low implementation complexity of the DS algorithm, and the larger server size of the PE algorithm. The Sporadic Server (SS) algorithm has these qualities.

The SS algorithm, like the DS and PE algorithms, creates a high-priority task for servicing aperiodic tasks. The SS algorithm preserves its server execution time at its high priority level until an aperiodic request occurs. The SS algorithm differs from the DS and PE algorithms in the way it replenishes its server execution time. The DS and PE algorithms periodically replenish their server execution time to full capacity. The SS algorithm only replenishes its server execution time after some or all of the execution time is consumed by aperiodic task execution. This method of replenishing server execution time sets the SS algorithm apart from the DS and PE algorithms and is central to understanding the operation of the SS algorithm. The following terms are used to explain the SS algorithm's method of replenishing server execution time:

- $P_S$       The task priority level at which the system is currently executing.
- $P_i$       One of the priority levels in the system. Priority levels are consecutively numbered in priority order, with  $P_1$  being the highest priority level and  $P_2$  being the next highest.
- Active      This term is used to describe a period of time with respect to a particular priority level. A priority level  $P_i$  is considered to be *active* if the current priority of the system,  $P_S$ , is equal to or higher than the priority of  $P_i$ .
- Idle      This term has the opposite meaning of the term *active*. A priority level  $P_i$  is *idle* if the current priority of the system  $P_S$ , is less than the priority of  $P_i$ .
- $RT_i$       The replenishment time for priority level  $P_i$ . This is the time at which consumed execution time for the sporadic server of priority level  $P_i$  will be replenished.

Replenishment of consumed sporadic server execution time for a sporadic server executing at priority level  $P_i$  proceeds as follows:

If the server has execution time available, the replenishment time,  $RT_i$ , is set when priority level  $P_i$  becomes active. The value of  $RT_i$  is set equal to the current time plus the period of  $P_i$ . If the server capacity has been exhausted, the next replenishment time can be set when the server's capacity becomes nonzero and  $P_i$  is active.

Replenishment of any consumed server execution time is scheduled to occur at  $RT_i$  if either the priority level  $P_i$  becomes idle or the server's execution time has been exhausted. The amount to be replenished is equal to the amount of server execution time consumed.

### 3.2. SS algorithm example

The SS algorithm will be illustrated by comparing its operation to that of the DS and PE algorithms using a simple periodic task set. The task set is composed of two periodic tasks:  $\tau_1$  and  $\tau_2$ . Task  $\tau_1$  has the shorter period and thus is assigned a higher priority than task  $\tau_2$  by the rate monotonic algorithm. For these examples, both  $\tau_1$  and  $\tau_2$  begin their periods at time = 0. The periodic task set parameters are as follows:

Task	Exec time	Period	Utilization (%)
$\tau_1$	2	10	20.0
$\tau_2$	6	14	42.9

For this periodic task set, the maximum server sizes were determined for the DS, PE, and SS algorithms. For each algorithm, the server's period was chosen to be 5 units of time. This implies that the aperiodic server has the highest priority, followed by priorities of  $\tau_1$  and  $\tau_2$ . The initial period for each of the algorithms begins at time = 0. The server size characteristics for the DS, PE, and SS algorithms are as follows:

Algorithm	Exec time	Period	Server size (%)
DS	1.00	5	20.0
PE	1.33	5	26.7
SS	1.33	5	26.7

Figures 6, 7, and 8 show the behavior of the DS, PE, and SS algorithms for this task set. The upper part of these figures shows the task execution order, and the lower part shows the server capacity as a function of time. In each of these examples, two aperiodic requests occur. Both requests are for aperiodic tasks that require 1.33 units of execution time. The first aperiodic request occurs at time = 1 and the second occurs at time = 8.

Figure 6 shows the behavior of the DS algorithm for this task set. At time = 0, the server's execution time is brought to its full capacity of 1.00 unit and  $\tau_1$  begins execution. The server's capacity is preserved until the first aperiodic request occurs and is serviced at time = 1. At time = 2, the server's execution time is exhausted and  $\tau_1$  resumes execution. At time = 3,  $\tau_1$  completes execution and  $\tau_2$  begins execution. At time = 5, the server's execution time is brought to its full capacity of 1.00 unit and service for the first aperiodic request resumes, consuming 0.33 units of server execution time. At time = 5.33, the service for the first request is completed and  $\tau_2$  resumes execution. The response time for the first aperiodic request is 4.33 units of time. At time = 8, the second aperiodic request occurs and is serviced using the remaining 0.66 units of server execution time. At time = 8.66, aperiodic service is suspended and  $\tau_2$  resumes execution. At time = 10,  $\tau_2$  completes execution, the server's execution time is brought to its full capacity of 1.00 unit, and service for the second aperiodic request is resumed. At time = 10.66, service for the

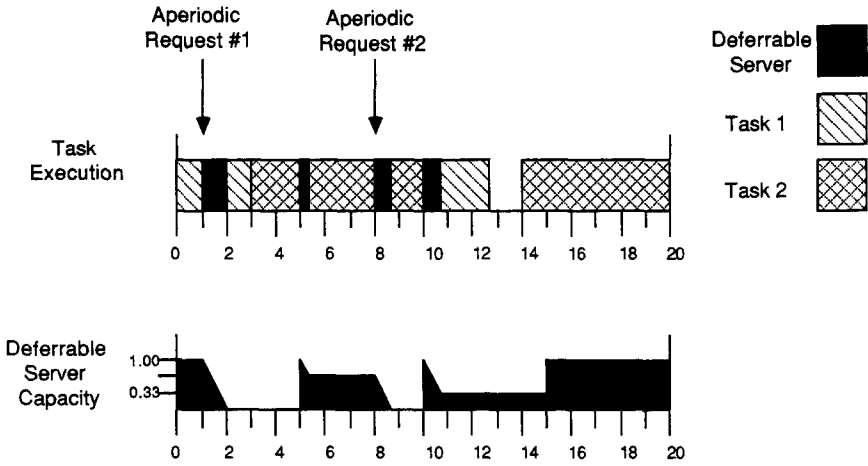


Figure 6. Deferrable Server example.

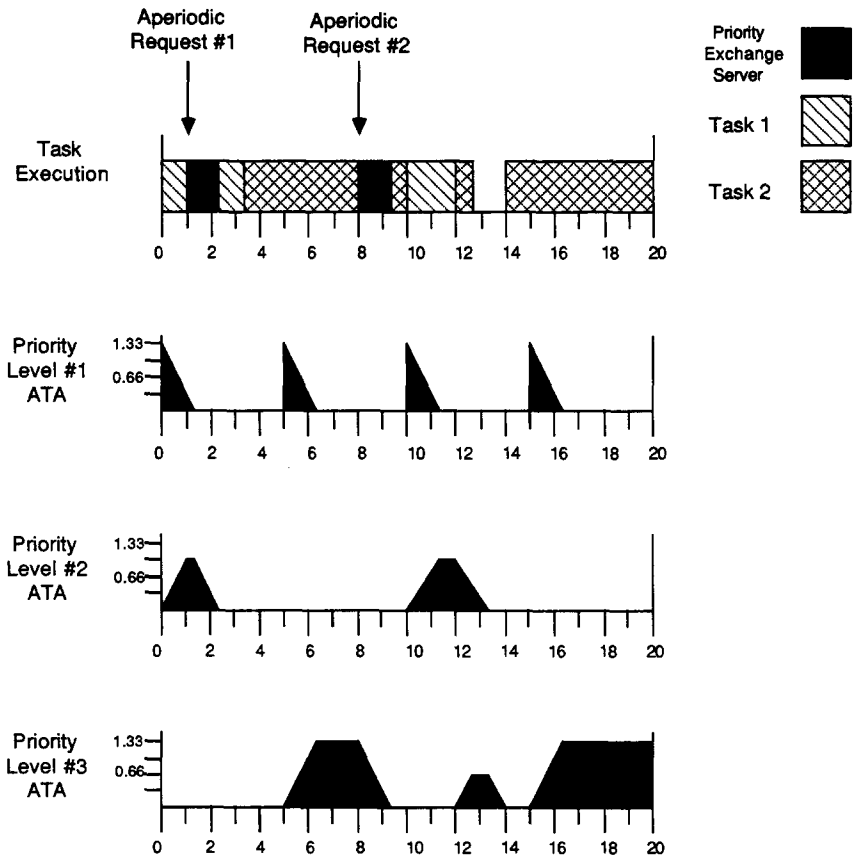


Figure 7. Priority Exchange Server example.

second aperiodic request is completed (leaving 0.33 units of server execution time) and  $\tau_1$  begins execution. The response time for the second aperiodic request is 2.66 units of time. At time = 14,  $\tau_2$  begins execution. At time = 15, the server's execution time is brought to its full capacity of 1.00 unit. At time = 20,  $\tau_2$  completes execution.

Figure 7 shows the behavior of the PE algorithm for this task set. Since the PE algorithm exchanges server execution time with lower-priority periodic tasks, three time lines are shown in Figure 7 for the aperiodic time available (ATA) of priority levels 1 (the highest), 2, and 3 (the lowest). At time = 0,  $\tau_1$  begins execution and the server time of priority level 1 is brought to its full capacity of 1.33 units. Since no aperiodic tasks are pending, the PE server's priority 1 execution time is exchanged with the periodic execution time of priority level 2. This exchange continues until the first aperiodic request occurs at time = 1, at which point the remaining priority 1 server time is used to service the aperiodic task. At time = 1.33, the priority 1 server time is exhausted and aperiodic service is continued using priority 2 server time. At time = 2.33, service for the first aperiodic request is complete, the PE server time is completely exhausted, and  $\tau_1$  resumes execution. At time = 3.33,  $\tau_1$  completes execution and  $\tau_2$  begins execution. At time = 5, the server time of priority level 1 is brought to its full capacity of 1.33 units. Since no aperiodic tasks are pending, the priority 1 server time is traded down to priority 3 server time by time = 6.33. At time = 8, the second aperiodic request occurs and is serviced using the priority 3 server time. At time = 9.33, service for the second aperiodic request is complete, the priority 3 server time is completely exhausted, and  $\tau_2$  resumes execution. At time = 10, the server time of priority 1 is brought to its full capacity of 1.33 units, and  $\tau_2$  begins execution. Since no aperiodic requests occur between time = 10 and time = 20, the PE server time is exchanged for lower-priority aperiodic execution time or is discarded when no exchanges are possible (as from time = 12.66 to time = 14, when no periodic or aperiodic tasks are ready to execute).

The following server characteristics should be noted by comparing the operation of the DS and PE algorithms in Figures 6 and 7.

The DS algorithm always preserves its server execution time at its original priority level, whereas the PE algorithm must exchange or discard its server execution time if no aperiodic tasks are pending. This quality of the DS algorithm allows a less complex implementation than is necessary for the PE algorithm.

Since the PE algorithm has a larger server size, it was able to provide a better response time for the aperiodic requests.

Both algorithms, periodically replenish their server execution time at the beginning of the server period.

Figure 8 shows the behavior of the SS algorithm using a high-priority sporadic server for this task set. Since the sporadic server is the only task executing at priority level  $P_1$  (the highest priority level),  $P_1$  becomes active only when the sporadic server executes an aperiodic task. Similarly, whenever the sporadic server is not executing an aperiodic task,  $P_1$  is idle. Therefore,  $RT_1$  is set whenever an aperiodic task is executed by the sporadic server. Replenishment of consumed sporadic server execution time will occur one server period after the sporadic server initially services an aperiodic task.

The task execution in Figure 8 proceeds as follows. For this example, the sporadic server




Figure 8. High-priority Sporadic Server example.

begins with its full execution time capacity. At time = 0,  $\tau_1$  begins execution. At time = 1, the first aperiodic request occurs and is serviced by the sporadic server. Priority level  $P_1$  has become active and  $RT_1$  is set equal to 6. At time = 2.33, the servicing of the first aperiodic request is completed, exhausting the server's execution time, and  $P_1$  becomes idle. A replenishment of 1.33 units of time is set for time = 6 (note the arrow in Figure 8 pointing from time = 1 on the task execution time line to time = 6 on the server capacity time line). The response time of the first aperiodic request is 1 unit of time. At time = 3.33,  $\tau_1$  completes execution and  $\tau_2$  begins execution. At time = 6, the first replenishment of server execution time occurs, bringing the server's capacity up to 1.33 units of time. At time = 8, the second aperiodic request occurs and  $P_1$  becomes active as the aperiodic request is serviced using the sporadic server's execution time.  $RT_1$  is set equal to 13. At time = 9.33, the servicing of the second aperiodic request completes,  $P_1$  becomes idle, and  $\tau_2$  is resumed. A replenishment of 1.33 units of time is set for time = 13 (note the arrow in Figure 8 pointing from time = 8 on the task execution time line to time = 13 on the server capacity time line). At time = 13, the second replenishment of server execution time occurs, bringing the server's capacity back up to 1.33 units of time.


By comparing Figures 6, 7, and 8, the following advantages of the SS algorithm are seen:

The SS algorithm has a low implementation complexity comparable to the DS algorithm, because it maintains its server execution time at its original priority level and does not exchange server execution time with lower priority levels as the PE algorithm does. The SS algorithm has the same server size advantage over the DS algorithm as the PE algorithm does.

The SS algorithm may also have a run-time advantage over the DS and PE algorithms. A run-time overhead is incurred periodically for the DS and PE algorithms to replenish their server execution time. This overhead is paid whether or not any aperiodic tasks were serviced during the last server period. The SS algorithm only pays a replenishment overhead if some of its server execution time has been consumed.

To better understand the SS algorithm's replenishment method, the previous high-priority sporadic server example presented in Figure 8 is augmented with examples of an equal- and a medium-priority sporadic server. For these examples, presented in Figures 9 and 10, all aperiodic requests require 1 unit of execution time.

Figure 9 shows the task execution and the task set characteristics for the equal-priority sporadic server example. The sporadic server and  $\tau_1$  both execute at priority level  $P_1$  and  $\tau_2$  executes at priority level  $P_2$ . At time = 0,  $\tau_1$  begins execution,  $P_1$  becomes active, and  $RT_1$  is set to 10. At time = 2, the first aperiodic request occurs and is serviced by the sporadic server. At time = 3, service is completed for the first aperiodic request,  $\tau_2$  begins execution,  $P_1$  becomes idle, and a replenishment of 1 unit of server execution time is set for time = 10. At time = 8, the second aperiodic request occurs and is serviced using the sporadic server,  $P_1$  becomes active, and  $RT_1$  is set to 18. At time = 9, service is completed for the first aperiodic request,  $\tau_2$  resumes execution,  $P_1$  becomes idle, and a replenishment of 1 unit of server execution time is set for time = 18. At time = 10,  $\tau_1$  begins execution and causes  $P_1$  to become active and the value of  $RT_1$  to be set. However, when  $\tau_1$  completes at time = 12 and  $P_1$  becomes idle, no sporadic server execution time has been consumed. Therefore, no replenishment time is scheduled even though the priority level of the sporadic server became active.




<u>Task</u>	<u>Exec Time</u>	<u>Period</u>	<u>Utilization</u>
SS	2	10	20.0%
$\tau_1$	2	10	20.0%
$\tau_2$	6	14	42.9%

Figure 9. Equal-priority Sporadic Server example.



Figure 9 illustrates two important properties of the SS algorithm. First,  $RT_i$  can be determined from a time that is *less* than the request time of an aperiodic task. This occurs for the first aperiodic request in Figure 9 and is allowed because  $P_1$  became active before and remained active until the aperiodic request occurred. Second, the amount of execution time replenished to the sporadic server is equal to the amount consumed. When the PE and DS algorithms replenish their server execution time, the server capacity is always brought to its maximum value regardless of how much server execution was used.

Figure 10 shows the task execution and task set characteristics for the medium-priority sporadic server example. The sporadic server executes at priority level  $P_2$ , between the priority levels of  $\tau_1$  ( $P_1$ ) and  $\tau_2$  ( $P_3$ ). At time = 0,  $\tau_1$  begins execution. At time = 1,  $\tau_1$  completes execution and  $\tau_2$  begins execution. At time = 4.5, the first aperiodic request occurs and is serviced using the sporadic server, making priority level  $P_2$  active. At time = 5,  $\tau_1$  becomes active and preempts the sporadic server. At this point, both priority levels  $P_1$  and  $P_2$  are active. At time = 6,  $\tau_1$  completes execution,  $P_1$  becomes idle, and the sporadic server is resumed. At time = 6.5, service for the first aperiodic request is completed,  $\tau_1$  resumes execution, and  $P_2$  becomes idle. A replenishment of 1 unit of sporadic server execution time is scheduled for time = 14.5. At time = 8, the second aperiodic request occurs and consumes 1 unit of sporadic server execution time. A replenishment of 1 unit of sporadic server execution time is set for time = 18.



<u>Task</u>	<u>Exec Time</u>	<u>Period</u>	<u>Utilization</u>
$\tau_1$	1.0	5	20.0%
SS	2.5	10	25.0%
$\tau_2$	6.0	14	42.9%

Figure 10. Medium-priority Sporadic Server example.

Figure 10 illustrates another important property of the SS algorithm. Even if the sporadic server is preempted and provides discontinuous service for an aperiodic request (as occurs with the first aperiodic request in Figure 10), only one replenishment time is necessary. Preemption of the sporadic server does not cause the priority level of the sporadic server to become idle, allowing the discontinuous consumption of sporadic server execution time to be replenished continuously. Note that one replenishment for the consumption of sporadic server execution time resulting from both aperiodic requests in Figure 10 is not permitted because the priority level of the sporadic server became idle between the completion of the first aperiodic request and the initial service of the second aperiodic request.

### 3.3. The schedulability of sporadic servers

In this section, we prove that, from a scheduling point of view, a sporadic server can be treated as a standard periodic task with the same period and execution time as the sporadic server. It is necessary to prove this claim because the sporadic server violates one of the basic assumptions governing periodic task execution as described by Liu and Layland (1973) in their analysis of the rate monotonic algorithm. Given a set of periodic tasks that is schedulable by the rate monotonic algorithm, this assumption requires that, once a periodic task is the highest-priority task that is ready to execute, it must execute.<sup>2</sup> If a periodic task *defers* its execution when it otherwise could execute immediately, then it may be possible that a lower-priority task will miss its deadline even if the set of tasks was schedulable. A sporadic server does not meet this Liu and Layland requirement because even though it may be the highest-priority task that is ready to execute, it will preserve its execution time if it has no pending requests for aperiodic service. This preservation of the server's execution time is equivalent to deferring the execution of a periodic task. A deferrable server also fails to meet the above requirement because of its preservation capabilities. To prove that a sporadic server can still be treated as normal periodic task, we will show that the sporadic server's replenishment method compensates for any deferred execution of the sporadic server. In contrast, we will also show how the replenishment method for a deferrable server fails in this respect.

To demonstrate how deferred execution can cause a lower-priority task to miss its deadline, the execution time of a periodic task will be compared to the execution of the deferrable server with the same period and execution time. Figure 11 presents the execution behavior of three periodic tasks. Let  $T$  represent the period of a periodic task and  $C$  represent its execution time. Task A, with  $T_A = 4$  and  $C_A = 1$ , has the highest priority. Task B, with  $T_B = 5$  and  $C_B = 2$ , has a medium priority. Task C, with  $T_C = 10$  and  $C_C = 3$ , has the lowest priority. Tasks A and B begin their first periods at time = 0, and task C begins its first period at time = 3. Note that no idle time exists during the first period of task C. This constrains task C to a maximum execution time of 3 units.

Figure 12 presents the execution of the task set of Figure 11, with task B replaced with deferrable server. Also presented in Figure 12 is the capacity of the deferrable server as a function of time. For this example, the following aperiodic requests are made to the deferrable server:

Request	Request Instant	Exec Time
1	1	1
2	3	1
3	5	2
4	10	2

Referring to the deferrable server execution graph in Figure 12, at time = 2, the service for the first aperiodic request completes and no other aperiodic requests are pending. At this point, the deferrable server *defers* its execution by preserving its remaining execution time until the second aperiodic request occurs at time = 3. Note that this deferred execution, followed by the servicing of the second aperiodic request from time = 3 to time = 4, has blocked task C from executing during this interval, whereas during the same interval in Figure 11, task C was not blocked. The third and fourth aperiodic requests are executed by the deferrable server during the same intervals as task B executes in Figure 11. Thus, task A and the deferrable server limit task C to a maximum of 2 units of execution time per period, whereas the original periodic task was able to complete 3 units of computation during the same period in Figure 11. If task C had needed 3 units of execution time during the period shown in Figure 12, it would have missed its deadline. It is this *invasive* quality of the deferrable server for lower-priority periodic tasks that results in the lower scheduling bound for the DS algorithm described in Section 2.2.

Figure 13 presents the execution of the task set of Figure 11 with task B replaced by a sporadic server. The third time line in Figure 13 is the capacity of the sporadic server as a function of time. The arrows in Figure 13 indicate the replenishment of consumed sporadic server execution time. The requests for aperiodic service are the same as the

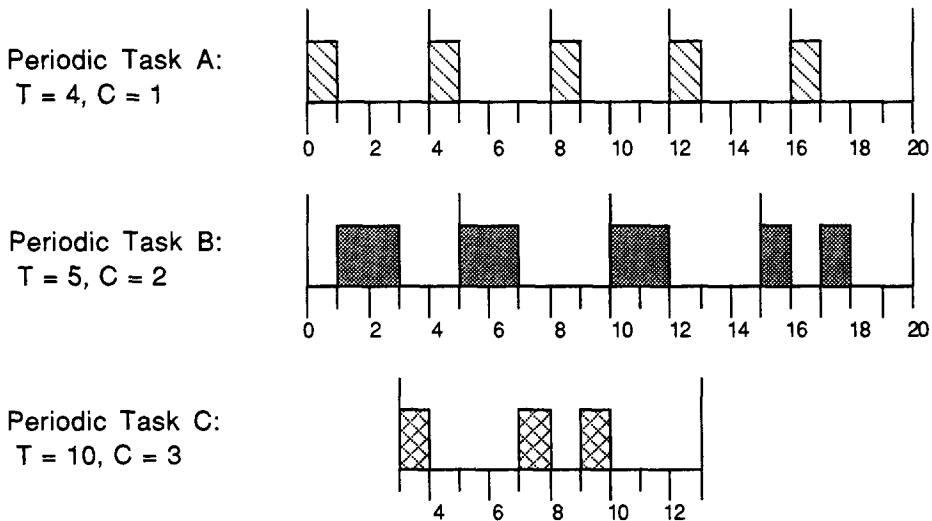


Figure 11. Periodic task example.




Figure 12. An example of deferred execution with the DS algorithm.




Figure 13. An example of deferred execution with the SS algorithm.

requests of Figure 12. Note that the sporadic server, like the deferrable server in Figure 12, blocks the execution of task C during time = 3 to time = 4. However, the sporadic server replenishment method prevents the execution time consumed during this interval from being replenished until time = 8. This allows task C to execute from time = 6 to time = 7, whereas the deferrable server was executing during this interval in Figure 12. The sporadic server, unlike the deferrable server, blocks task C from executing from time = 9 to time = 10. However, the replenishment of the server execution time consumed during this interval does not occur until time = 14. This allows task C to complete its execution from time = 11 to time = 12. Thus, in this example, the sporadic server allows task C to meet its deadline, whereas a deferrable server having the same period and execution time cannot. However, one should note that the sporadic server completes the third and fourth aperiodic requests 3 time units later than the deferrable server. Thus, for the same series of aperiodic requests, a sporadic server may provide a longer response time than an equivalently sized deferrable server. The next section presents a simulation study that investigates the relative performance of these algorithms.

Now that we have shown a specific example of how the replenishment method of the SS algorithm can compensate for deferred execution of the sporadic server, we need to prove that, in terms of schedulability, a sporadic server can always be treated as a periodic task having the same period and execution time. To do this, we first show that a sporadic server can behave in a manner identical to a periodic task with the same period and execution time. Next, we show that any execution of the sporadic server before the server's priority level becomes idle falls into one of three cases: (1) none of the server's execution time is consumed, (2) the server's execution time is completely consumed, or (3) the server's execution time is only partly consumed. In the first case, the server can be shown to be equivalent to a periodic task with a delayed request. In the second case, the server's execution is shown to be identical to that of a normal periodic task. In the third case, the execution behavior of the server is shown to be equivalent to two periodic tasks: one that executes normally and one that is delayed. Thus, all types of sporadic server execution are shown to be equivalent to the execution of one or more periodic tasks.

In order to explore the schedulability effects of sporadic servers upon a periodic task set, several terms and concepts developed by Liu and Layland (1973) are needed. A periodic task set is composed of  $n$  independent tasks,  $\tau_1$  through  $\tau_n$ , numbered in order of decreasing priority. A periodic task  $\tau_i$  is characterized by its computation time,  $C_i$ , and a period,  $T_i$ . At all times, the highest-priority task that is ready to execute is selected for execution, preempting any lower-priority tasks as necessary. The deadline of a periodic task occurs one period after it is requested. The *critical instant* for a periodic task is the instant at which a request for that task will have the longest response time. Liu and Layland state and prove the following theorem concerning the critical instant of a periodic task:

**Theorem 1.** Given a set of periodic tasks, the critical instant for any task occurs whenever the task is requested simultaneously with all higher-priority tasks.

The *critical zone* for a periodic task is the time interval between its critical instant and the completion of that request. A periodic task set is schedulable if the critical zone for each task is less than or equal to its deadline.

We now state and prove the following property of schedulable periodic task sets:

**Lemma 1.** Given a periodic task set that is schedulable with  $\tau_i$ , the task set is also schedulable if  $\tau_i$  is split into  $k$  separate periodic tasks,  $\tau_{i,1}, \dots, \tau_{i,k}$ , with  $\sum_{j=1}^k C_{i,j} = C_i$  and  $T_{i,j} = T_i$  for  $1 \leq j \leq k$ .

**Proof.** By the rate monotonic algorithm, all tasks  $\tau_{i,1}, \dots, \tau_{i,k}$  will be assigned the same priority because they have the same period. The rate monotonic algorithm schedules a periodic task set independently of the phasing between any two tasks. Suppose the requests for tasks  $\tau_{i,1}, \dots, \tau_{i,k}$  are all in sync with each other. The execution pattern of  $\tau_{i,1}, \dots, \tau_{i,k}$  will now be identical to the execution pattern of the original task,  $\tau_i$ , and therefore the task set is still schedulable for this phasing of  $\tau_{i,1}, \dots, \tau_{i,k}$ . Since the task set is schedulable for one phasing of  $\tau_{i,1}, \dots, \tau_{i,k}$  it is schedulable for them all. The lemma follows.

Next we establish that the execution behavior of a sporadic server can be identical to that of a periodic task with the same period and execution time.

**Lemma 2.** Given a schedulable periodic task set, replace a periodic task with a sporadic server having the same period and execution time. If the requests for the sporadic server are identical to that of the original periodic task, then the execution behavior of the sporadic server is identical to that of the original periodic task.

**Proof.** The periodic task and the sporadic server execute at the same priority level, and each request for the sporadic server and the periodic task requires the same execution time. Each request for the sporadic server completely consumes its execution time, and the consumed execution time is replenished before the next request. Therefore, the sporadic server's execution is identical to that of the original periodic task.

We now show that sporadic servers are equivalent to periodic tasks in terms of schedulability.

**Theorem 2.** A periodic task set that is schedulable with a task,  $\tau_i$ , is also schedulable if  $\tau_i$  is replaced by a sporadic server with the same period and execution time.

**Proof.** To prove this theorem, we will show that, for all types of sporadic server execution, the sporadic server exhibits an execution behavior that can also be represented by a combination of one or more periodic tasks.

Let the sporadic server's period be  $T_{SS}$  and execution time be  $C_{SS}$ . At time  $t_A$ , let the priority level of the server become active and let the sporadic server be at full capacity. Consider the interval of time beginning at  $t_A$  during which the priority level of the server remains active. The execution behavior of the server during this interval can be described by one of three cases:

1. None of the server's execution time is consumed.
2. All of the server's execution time is consumed.
3. The server's execution time is partly consumed.

*Case 1.* If the server is not requested during this interval, then it preserves its execution time. The server's behavior is identical to a periodic task for which the interarrival time between two successive requests for the periodic task is greater than its period. Referring to Theorem 1, if any request for a periodic task is delayed, it cannot lengthen the critical zone of any lower-priority periodic task because the lower-priority task would be able to execute during the delay, yielding a *shorter* critical zone. Since a critical zone of a lower-priority task cannot be made longer, the delay has no detrimental schedulability effect.

*Case 2.* If the execution time of the sporadic server is completely consumed, a replenishment will be scheduled to occur at  $t_A + T_{SS}$  to bring the server back to full capacity. By Lemma 2 the behavior of the sporadic server between  $t_A$  and  $t_A + T_{SS}$  is identical to that of a similar periodic task that is requested at  $t_A$ .

*Case 3.* If the server's execution time is not completely exhausted, then a replenishment will be scheduled to occur at  $t_A + T_{SS}$  for the consumed execution time and the server will preserve the unconsumed execution time until it is requested. Let the amount of execution time to be replenished be  $C_R$ .

Now consider a periodic task with a period of  $T_{SS}$  and an execution time of  $C_{SS}$  that is split into two periodic tasks,  $\tau_x$  and  $\tau_y$ , both with a period of  $T_{SS}$ . Let  $\tau_x$  have an execution time of  $C_R$  and let  $\tau_y$  have an execution time of  $C_{SS} - C_R$ . By Lemma 1, the splitting of the original periodic task can be done without affecting schedulability. Let requests for both  $\tau_x$  and  $\tau_y$  be in sync until  $t_A$ . At  $t_A$ , let  $\tau_x$  execute normally but delay  $\tau_y$ . As in case 1, the delay for  $\tau_y$  has no schedulability effect. The behavior of the two periodic tasks  $\tau_x$  and  $\tau_y$  from  $t_A$  to  $t_A + T_{SS}$  is identical to that of the sporadic server over the same time interval.

Since a sporadic server's execution in each of these cases can be represented by a periodic task or a combination of periodic tasks with a period and total execution time that is identical to that of the sporadic server, the theorem follows.

If a sporadic server's execution time is only partly consumed before a replenishment is scheduled, as described in case 3 of the proof for Theorem 2, the server must then manage two quantities of execution time: the execution time that will be replenished and the unconsumed execution time. This could be a concern for a sporadic server implementation if many successive executions of the sporadic server split the server's execution time into many small quantities. However, if the server is not requested for a while, these small quantities can be merged together. In fact, regardless of how many times a server's execution time has been split into smaller quantities, if the server is not requested for an amount of time equal to or greater than its period, then all of these quantities can be merged together. An example of this behavior is shown in Figure 14. The sporadic server is the highest-priority task in the system and has a period of 5 and an execution time of 2. The first four requests for the server require 0.5 units of execution time. Since the priority level of the sporadic server becomes idle between each of these requests, the execution time of the sporadic server has been split four ways, requiring four separate replenishment times. However, no requests are made to the server from time = 4 to time = 9, one server period after the last request. By time = 9, the server's capacity has been completely replenished and merged together.

In this section, we have shown that although a sporadic server can defer its execution time, it can still be treated as a periodic task. However, we also showed that this is not




Figure 14. An example of the splitting and merging of sporadic server execution time.

true for a deferrable server. The key difference between these two types of servers is that when a sporadic server defers its execution, it also defers the replenishment of its execution time once it is consumed. In contrast, a deferrable server replenishes its execution time independently of when its execution time is consumed.

### 3.4. Sporadic server performance

Simulations were conducted to compare the response time performance of the polling, DS, PE, and SS algorithms for soft deadline aperiodic tasks. For these simulations, a set of ten periodic tasks with periods ranging from 54 to 1,200 was chosen. Three periodic loads were simulated for this set of periods: 40%, 69%, and 88%. The 40% load represents a low periodic load. The 69% case corresponds to the maximum resource utilization for which the rate monotonic algorithm can always guarantee schedulability, independent of the particular task periods (Liu and Layland 1973). The 88% case represents the average case scheduling bound for the rate monotonic algorithm (Lehoczky, Sha, and Ding 1987). The server period for each of the polling, DS, PE, and SS algorithms was chosen to be 54. The server execution times were selected to be the maximum value for which all the periodic tasks remain schedulable. The aperiodic load for these simulations was varied across the range of resource utilization unused by the periodic tasks. The interarrival times for the aperiodic tasks were modeled using a Poisson arrival pattern with three different average interarrival times: 18, 36, and 54. The aperiodic service times were modeled using an exponential distribution.

Figures 15, 16, and 17 present the results of these simulations for each of the mean aperiodic interarrival times simulated (18, 36, and 54). In each of these figures, three graphs correspond to the different periodic loads simulated (40%, 69%, and 88%). In all the graphs of these figures, the average aperiodic response time of each algorithm is plotted as a function of the aperiodic load. The average aperiodic response time of each algorithm is presented relative to the response time of background aperiodic service. In other words, the average response time equivalent to background service has a value of 1.0 on all the graphs. An improvement in average aperiodic response time over background service corresponds to a value of less than 1.0. The lower the response time curve lies on these graphs, the better the algorithm is for improving aperiodic responsiveness.



As can be seen from each of the graphs presented in Figures 15, 16, and 17, the DS, PE, and SS algorithms can provide a substantial reduction in average aperiodic response time compared to background or polling aperiodic service. These algorithms provide the greatest improvement for short, frequent aperiodic tasks, as can be seen in the lower left-hand side of each graph. The performance of the SS algorithm in each of these graphs is comparable to the performance of the DS and PE algorithms.

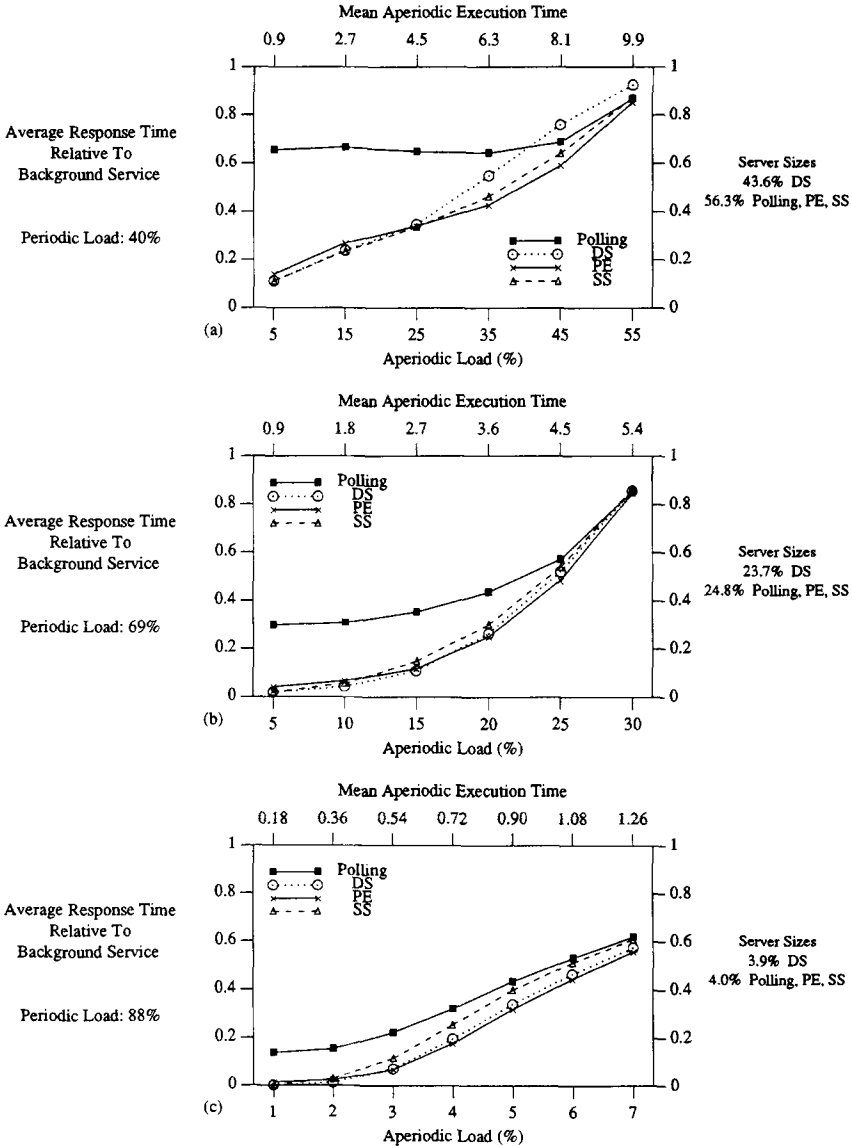


Figure 15. SS algorithm performance comparison, mean aperiodic interarrival time = 18.

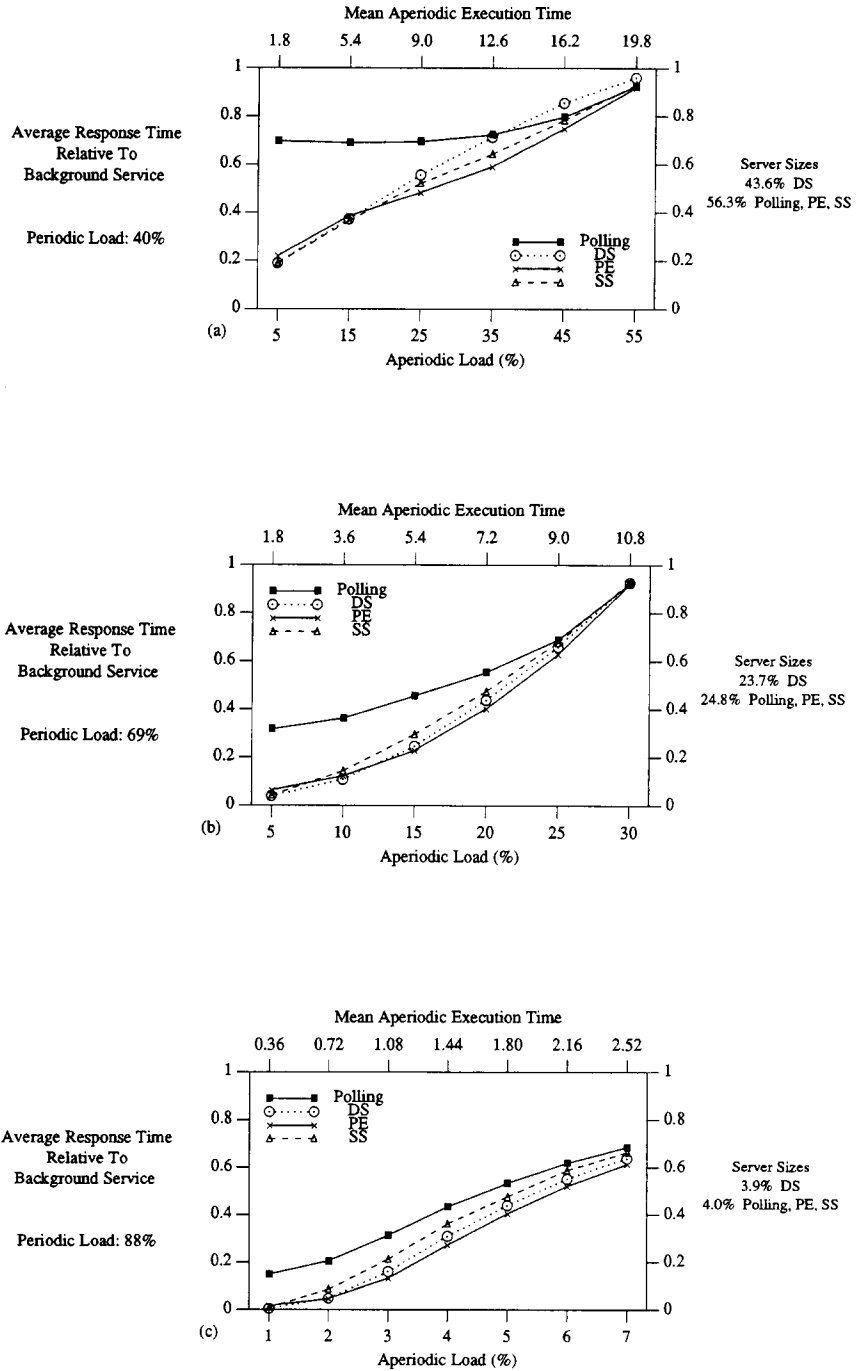


Figure 16. SS algorithm performance comparison mean aperiodic interarrival time = 36.

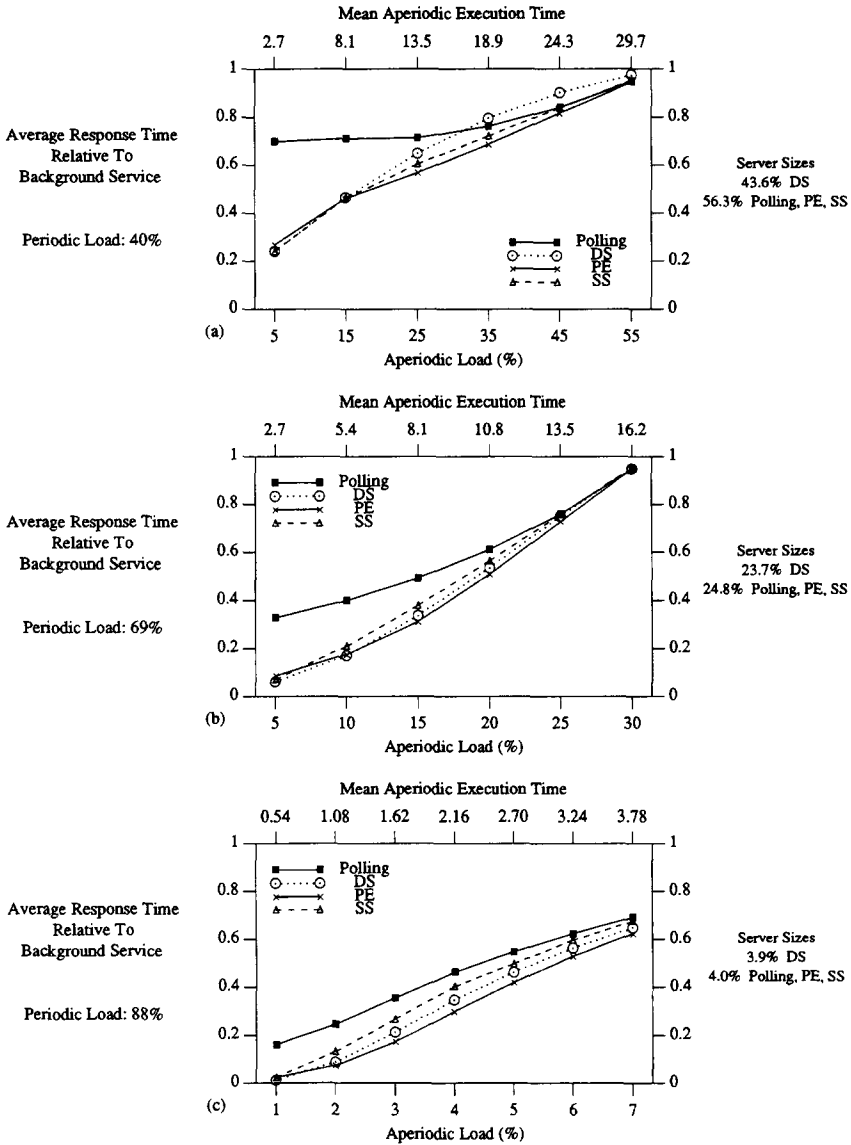


Figure 17. SS algorithm performance comparison, mean aperiodic interarrival time = 54.



Referring to the 40% periodic load case presented in graphs 15a, 16a, and 17a, the performance of the SS and PE algorithms becomes better than the performance of the DS algorithm as the aperiodic load increases. The performance of the DS algorithm even becomes worse than the performance of polling for high aperiodic load. This increase in the relative response time for the DS algorithm is due to the smaller server size of 43.6% compared to a server size of 56.3% for the polling, PE, and SS algorithms. The larger

server size of the SS algorithm is an advantage over the DS algorithm for a task set with a moderate periodic load and a high aperiodic load.

The 69% and 88% periodic load cases presented in graphs (b) and (c) of Figures 15, 16, and 17 show that the average response time of the SS algorithm becomes slightly higher than the response time of the DS and PE algorithms for moderate to high aperiodic load. The behavior is not attributable to a difference in server size, because as the periodic load increases, the difference in server sizes diminishes. The difference between the DS server size and server size for the polling, PE, and SS algorithms is about 1% for a periodic load of 69% and decreases to about 0.1% for a periodic load of 88%. The cause for the slightly higher response time of the SS algorithm is its more restrictive replenishment method, as discussed in Section 3.3. The DS and PE algorithms periodically replenish their server execution time independent of when the server's execution time was consumed. The SS algorithm replenishes consumed server execution time in a sporadic manner that depends upon when the execution time is consumed. On the average, once server execution time is exhausted, aperiodic tasks must wait longer for the SS algorithm to replenish its server execution time than with the DS or PE algorithms.

#### 4. Scheduling sporadic tasks

One important type of aperiodic task that is not generally supported by previous aperiodic service algorithms is a sporadic task, which is an aperiodic task with a hard deadline and a minimum interarrival time. To guarantee a hard deadline for a sporadic task, a high-priority server can be created to exclusively service the sporadic task. This server preserves its execution time at its original priority level until it is needed by the sporadic task. To guarantee that the server will always have sufficient execution time to meet the sporadic task's deadline, a minimum interarrival time restriction must be placed upon the sporadic task. This minimum interarrival time must be equal to or greater than the period of the server task. As long as the sporadic task does not request service more frequently than the minimum interarrival time, its hard deadline can be guaranteed. The SS algorithm can be used to provide a guarantee for a sporadic task as long as the aperiodic task's deadline is equal to or greater than its minimum interarrival time. However, for the cases when the deadline is shorter than the minimum interarrival time, a different priority assignment for the sporadic server and a different schedulability analysis are necessary. The priority of the sporadic server should be based upon the deadline of its associated sporadic task, not upon the maximum arrival rate of the sporadic task. This type of priority assignment, first considered by Leung and Whitehead in (1982), is referred to as deadline monotonic and requires a schedulability analysis different from that for a rate monotonic priority assignment. Apart from the assignment of server priority, the operation of a deadline monotonic sporadic server is identical to that of the SS algorithm as presented in Section 3.1. Section 4.1 demonstrates the necessity of a deadline monotonic priority assignment for a short deadline sporadic task and describes the schedulability analysis for task set where the priorities of the periodic tasks are assigned by the rate monotonic algorithm and the priorities of the sporadic servers are assigned by the deadline monotonic algorithm.

Periodic Task	Exec Time	Period	Utilization
	4	12	33.3%
	4	20	20.0%


Aperiodic Task	Exec Time	Minimum Interarrival Time	Utilization	Deadline
	8	32	25.0%	10

Figure 18. Periodic task set with short deadline aperiodic task.




#### 4.1. A simple example of the deadline monotonic priority assignment

The necessity for a deadline monotonic priority assignment for a short deadline sporadic task will be illustrated with a simple example that shows how a rate monotonic priority assignment cannot guarantee the sporadic task's hard deadline, while a deadline monotonic priority assignment can be used to guarantee that the hard deadline will always be met. The set of tasks presented in Figure 18 is used for this example.

Figure 19 presents the execution of these tasks using the rate monotonic priority assignment for the sporadic task. In Figure 19 a high priority is represented by a low number. Note that the rate monotonic algorithm assigns the lowest priority to the sporadic task because of its long minimum interarrival time. The task execution graph presented in Figure 19 assumes that requests for both periodic tasks and the sporadic task occur at time = 0. Both periodic tasks meet their deadlines, but the sporadic task completes only 2 of the required 8 units of execution time before it misses its deadline. Clearly, the rate monotonic algorithm is inappropriate for this task set.

Since a higher priority is needed for the short deadline sporadic task, one might consider treating it as a periodic task with a period of 10. This would give the sporadic task the highest priority and guarantee that its deadline would be met. However, this would cause other problems. The sporadic task only needs a resource utilization of 25% to meet its deadline. Treating it as a periodic task with a period of 10 means that a resource utilization of 80% (an execution time of 8 units divided by a period of 10 units) would be dedicated to a task that needs only 25%. This is a very wasteful scheduling policy. Also, if the priority of the sporadic task is increased in this manner, the total utilization required to schedule all three tasks is now 133.3% (33% + 20% + 80%), making this an infeasible schedule. A different method is needed that will assign priorities such that all their deadlines will be met using an efficient level of scheduled resource utilization.

Rate Monotonic Priority Assignment:

	Rate	Exec Time	Deadline	Priority
	12	4	12	1
	20	4	20	2
	32	8	10	3




Figure 19. Example of rate monotonic priority assignment for a short deadline sporadic task.

A deadline monotonic priority assignment can be used to efficiently schedule the task set presented in Figure 18. A sporadic server is created with an execution time of 8 units and a period of 32 units to service the short deadline sporadic task. The priority of the sporadic task is assigned based upon the deadline of the task it is servicing, not upon its period. In other words, the rate monotonic algorithm is used to assign priorities for the periodic tasks, and the priority of the sporadic task is assigned as if its rate of occurrence were equal to its deadline. The task execution graph for this priority assignment is presented in Figure 20. The sporadic task now has the highest priority and meets its deadline. The execution of the periodic tasks is delayed but both still meet their deadlines.

Now that we have shown that a deadline monotonic priority assignment can be used to guarantee that a sporadic task will meet its deadline, we need to be able to perform a schedulability analysis that will indicate whether a deadline monotonic priority assignment for the sporadic server with a rate monotonic priority assignment for the periodic tasks is feasible for a given task set. The deadline monotonic priority assignment raises the priority of the sporadic server above the priority that would be assigned by the rate monotonic algorithm. As such, the sporadic server can be considered a low-priority task that is allowed to block a higher-priority task from executing. A similar problem can occur when periodic

Deadline Monotonic Priority Assignment:

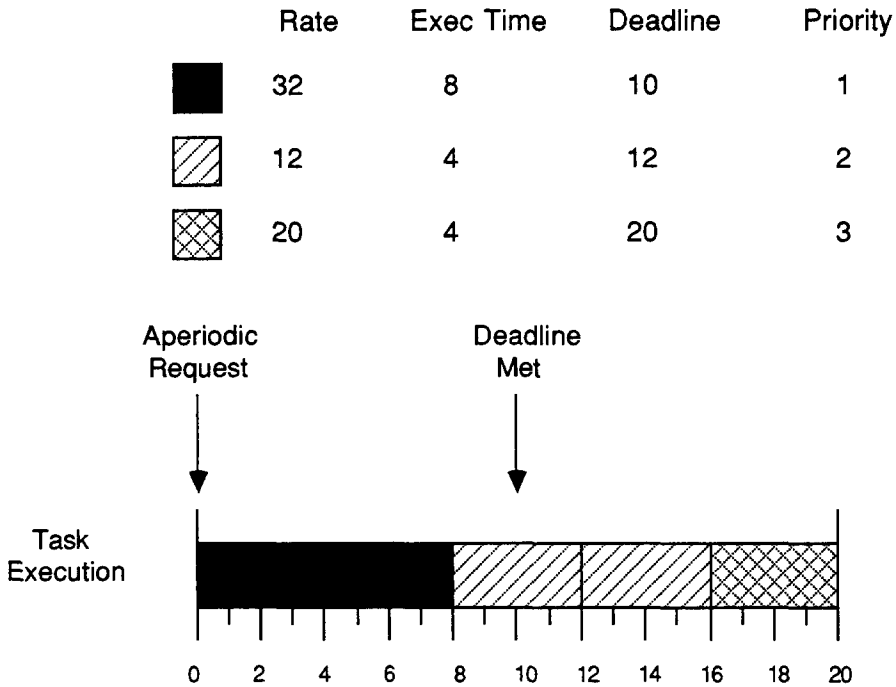


Figure 20. Example of deadline monotonic priority assignment for a short deadline sporadic task.

tasks that share data using critical sections are scheduled using the rate monotonic algorithm. A low-priority task that has already entered a critical section can block the execution of a high-priority task that needs to enter the critical section. The blocking of the high-priority task continues as long as the critical section is blocked by the lower-priority task. Sha, Rajkumar, and Lehoczky (1987) have developed the priority ceiling protocol for a set of periodic tasks that share data using critical sections.<sup>3</sup> This protocol establishes a set of sufficient conditions under which the periodic tasks can be scheduled using the rate monotonic algorithm. The schedulability analysis equations developed for the priority ceiling protocols can be used to test the schedulability of sporadic servers with deadline monotonic priority assignments:

$$\forall i \quad 1 \leq i \leq n, \quad \frac{C_1}{T_1} + \frac{C_2}{T_2} + \dots + \frac{C_i}{T_i} + \frac{B_i}{T_i} \leq i(2^{1/i} - 1) \quad (1)$$

$$\forall i \quad 1 \leq i \leq n, \quad \min_{(k, l) \in R_i} \left[ \sum_{j=1}^{i-1} U_j \frac{T_j}{lT_k} \left\lceil \frac{lT_k}{T_j} \right\rceil + \frac{C_i}{lT_k} + \frac{B_i}{lT_k} \right] \leq 1 \quad (2)$$

where  $C_i$  and  $T_i$  are respectively the execution time and period of task  $\tau_i$ ;  $U_i = C_i/T_i$  is the utilization of task  $\tau_i$ ; and  $R_i = \{(k, l) \mid 1 \leq k \leq i, l = 1, \dots, \lfloor T_i/T_k \rfloor\}$ . The term  $B_i$  is the worst-case blocking time for task  $\tau_i$ .

Equations 1 and 2 were derived from equations developed for testing the schedulability of a set of periodic tasks whose priorities have been assigned using the rate monotonic algorithm. Equation 1 was derived using the worst-case utilization bound equation for scheduling periodic tasks developed by Liu and Layland (1973), which, under the absolute worst-case conditions, provides a sufficient condition for determining schedulability of a rate monotonic priority assignment. Equation 2 was derived from an equation developed by Lehoczky, Sha, and Ding (1987) that provides necessary and sufficient conditions for determining schedulability of a rate monotonic priority assignment. Both Equations 1 and 2 provide sufficient conditions under which a periodic task set that allows lower-priority tasks to block higher-priority tasks can be scheduled by the rate monotonic algorithm. However, Equation 1 represents the absolute worst-case conditions, and a much tighter characterization is provided by Equation 2.

The blocking term,  $B_i$ , in Equations 1 and 2 represents the amount of time that a lower-priority task can block a higher-priority task from executing. For a sporadic server with a deadline monotonic priority,  $B_i$  is used to represent the amount of time the sporadic server can block a periodic task that has a higher priority than the rate monotonic priority of the sporadic server. To check the schedulability of a periodic task set with a deadline monotonic sporadic server, the term  $B_i$  is set equal to the execution time of the sporadic server for all  $\tau_i$  with a priority less than or equal to the priority of the sporadic server and greater than the sporadic server's original rate monotonic priority. For all  $\tau_i$  with a priority less than the sporadic server's original rate monotonic priority, the sporadic server should be treated normally (i.e., treated as a normal sporadic server with a rate monotonic priority assignment). For all  $\tau_i$  with a priority greater than the priority of the sporadic server, the corresponding value of  $B_i$  is set to zero.

The use of Equations 1 and 2 will be demonstrated with the task set presented in Figure 18. Let  $\tau_1$  and  $\tau_2$  be the periodic tasks and let  $\tau_3$  be the short deadline sporadic task. The following are the parameters for each task:

<i>Task</i>	$C_i$	$T_i$	$B_i$
$\tau_1$	4	12	8
$\tau_2$	4	20	8
$\tau_3$	8	32	0

Evaluation of Equation 1 proceeds as follows:

$$\begin{aligned}
 i = 1, \quad & \frac{C_1}{T_1} + \frac{B_1}{T_1} \leq 1 \\
 & \frac{4}{12} + \frac{8}{12} \leq 1 \\
 & \frac{12}{12} \leq 1
 \end{aligned}$$



$$i = 2, \quad \frac{C_1}{T_1} + \frac{C_2}{T_2} + \frac{B_2}{T_2} \leq 2(2^{\frac{1}{2}} - 1)$$

$$\frac{4}{12} + \frac{4}{20} + \frac{8}{20} \leq 0.824$$

$$\frac{56}{60} \geq 0.824$$

The inequality for  $i = 1$  holds, but the inequality for  $i = 2$  does not; therefore it is not necessary to check the  $i = 3$  case and Equation 2 must now be used.

The evaluation of Equation 2 proceeds as follows:

$i = 1$ : Check if  $C_1 + B_1 \leq T_1$ . Since  $4 + 8 \leq 12$ , task  $\tau_1$  is schedulable.

$i = 2$ : Check whether either of the following two inequalities hold:

$(k,l)$

$$(1,1) \quad C_1 + C_2 + B_2 \leq T_1 \quad 4 + 4 + 8 \geq 12$$

$$(2,1) \quad 2C_1 + C_2 + B_2 \leq T_1 \quad 8 + 4 + 8 \leq 20$$

The inequality for  $(k,l) = (2,1)$  holds and therefore,  $\tau_2$  is schedulable.

$i = 3$ : Check whether any of the following inequalities hold:

$(k,l)$

$$(1,1) \quad C_1 + 2C_2 + C_3 + B_3 \leq T_1 \quad 4 + 8 + 8 + 0 \geq 12$$

$$(1,2) \quad 2C_1 + 2C_2 + C_3 + B_3 \leq 2T_1 \quad 8+8+8+0 \leq 24$$

$$(2,1) \quad 2C_1 + C_2 + C_3 + B_3 \leq 8T_2 \quad 8+4+8+0 \leq 20$$

$$(3,1) \quad 3C_1 + 2C_2 + C_3 + B_3 \leq 8T_2 \quad 12+8+8+0 \leq 32$$

Since all the inequalities except  $(k,l) = (1,1)$  hold,  $\tau_3$  is schedulable. Note that it would have been sufficient to stop checking after finding one inequality that holds for each value of  $i$ ; all the inequalities are listed here for completeness only.

## 5. Summary and conclusions

This article described the Sporadic Server (SS) algorithm, a general algorithm for scheduling soft and hard deadline aperiodic tasks in real-time systems. The SS algorithm creates a server with a given period and utilization. The server maintains its utilization until it is needed by an aperiodic task and replenishes any consumed utilization in a sporadic manner based upon when the utilization is consumed. For soft deadline aperiodic tasks, a high-priority sporadic server is created that is shared by the soft deadline aperiodic tasks. For hard deadline aperiodic tasks (sporadic tasks), an individual sporadic server is dedicated to each hard deadline aperiodic task in order to guarantee that its deadline will always be met.

The SS algorithm was shown to have the low implementation complexity advantage of the Deferrable Server (DS) algorithm and the server size advantage of the Priority Exchange (PE) algorithm, and to provide a response time performance for soft deadline aperiodic tasks that is comparable to that attainable with either the DS or PE algorithms. In terms of schedulability, it was shown that a sporadic server can be treated as a periodic task having the same period and execution time. For short deadline sporadic tasks, it was also shown that the SS algorithm can guarantee their deadlines by using a deadline monotonic priority assignment rather than a rate monotonic priority assignment. Finally, it was shown that the schedulability analysis formulas developed for the priority ceiling protocols can be used to determine the schedulability of deadline monotonic sporadic servers.

### Acknowledgments

The authors would like to thank Ragunathan Rajkumar for his comments and help with this work.

This work was sponsored in part by the Office of Naval Research under contract N00014-84-K-0734, in part by the Naval Ocean System Center under contract N66001-87-C-0155, in part by Southeastern Center for Electrical Engineering Education, and in part by the Federal Systems Division of IBM Corporation under University Agreement YA-278067.

### Notes

1. See Sprunt, Lehoczky, and Sha 1988 for information on taking advantage of the stochastic execution time of periodic tasks.
2. For the case when two or more periodic tasks of equal priority are the highest priority tasks that are ready to execute, the requirement is that one of these tasks must execute.
3. See Sprunt, Sha, Lehoczky 1988 for information regarding the interaction of the priority inheritance protocols and sporadic servers.

### References

- Borger, M. W. 1987. *VAXELN Experimentation: Programming a Real-Time Periodic Task Dispatcher using VAX-ELN Ada 1.1*. Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA 15213.
- Lehoczky, J. P., L. Sha, and Y. Ding. 1987. *The Rate Monotonic Scheduling Algorithm: Exact Characterization and Average Case Behavior*. Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213.
- Lehoczky, J. P., L. Sha, and J. K. Strosnider. 1987. Enhanced Aperiodic Responsiveness in Hard-Real-Time Environments. *Proc. IEEE Real-Time Systems Symposium*, San Jose, CA, pp. 261-270.
- Leung, J. Y.-T., and J. Whitehead. 1982. On the Complexity of Fixed-Priority Scheduling of Periodic Real-Time Tasks. *Performance Evaluation 2*: 237-250.
- Liu, C. L., and J. W. Layland. 1973. Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment. *JACM 20* (January): 46-61.
- Mok, A. K.-L. 1983. Fundamental Design Problems of Distributed Systems for the Hard-Real-Time Environment. Ph.D. Dissertation, M.I.T., Cambridge, MA 02139.
- Sha, L., and J. Goodenough. 1988. *Real-Time Scheduling Theory in Ada*. Software Engineering Institute, Carnegie Mellon University.
- Sha, L., J. P. Lehoczky, and R. Rajkumar. 1986. Solutions for Some Practical Problems in Prioritized Preemptive Scheduling. *Proc. IEEE Real-Time Systems Symposium*, New Orleans, LA, pp. 181-191.
- Sha, L., R. Rajkumar, and J. P. Lehoczky. 1987. *Priority Inheritance Protocols: An Approach to Real-Time Synchronization*. Technical Report CMU-CS-87-181, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213.
- Sprunt, B., Sha L., and Lehoczky, J. P. 1988. *Scheduling Sporadic and Aperiodic Events in a Hard Real-Time System*. Software Engineering Institute, Carnegie Mellon University. Submitted for publication.
- Sprunt, B., J. P. Lehoczky, and L. Sha. 1988. Exploiting Unused Periodic Time for Aperiodic Service using the Extended Priority Exchange Algorithm. *Proc. IEEE Real-Time Systems Symposium*, Huntsville, AL.
- Strosnider, J. K. 1988. Highly Responsive Real-Time Token Rings. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA 15213.