

# An Analysis of Fixed-Priority Schedulability on a Multiprocessor

TR-050201

(revised 6/23/2005)

Theodore P. Baker  
Department of Computer Science  
Florida State University  
Tallahassee, FL 32306-4530  
e-mail: baker@cs.fsu.edu

## Abstract.

A new feasibility test for preemptive scheduling of periodic or sporadic real-time tasks on a single-queue  $m$ -server system allows for arbitrary fixed task priorities and arbitrary deadlines. For the special case when deadline equals period and priorities are rate monotonic, any set of tasks with maximum individual task utilization  $u_{\max}$  and minimum individual task utilization  $u_{\min}$  is feasible if the total utilization does not exceed  $m(1 - u_{\max})/2 + u_{\min}$ .

**Keywords:** deadline monotonic, fixed priority, multiprocessor, rate monotonic, real time, scheduling, symmetric multiprocessing, utilization bound

## 1. Introduction

Starting at least as early as the Safeguard anti-ballistic missile system and continuing up to the present, high performance real-time embedded systems have relied on multiprocessor architectures. With the trend toward multi-core architectures in the current and next generation of microprocessors embedded applications of multiprocessors are likely to become much more common.

The understanding of real-time multiprocessor scheduling has lagged behind that of single-processor scheduling. Over the three decades since Liu and Layland's 1973 seminal analysis of rate monotonic and deadline scheduling (Liu and Layland, 1973), the theory of fixed-priority single-processor scheduling has been refined, extended, and generalized, to the point that it is now very well understood (Sha et al., 2004). During those same decades comparatively little attention was paid to the possibility of extending the analysis to multiple processors.

Much of the analysis of multiprocessor scheduling that has been done has focussed on a partitioned model, in which tasks are assigned statically to processors (Dhall and Liu, 1978; Oh and Baker, 1998; Lopez et al., 2001; Lopez et al., 2000). The alternative, global scheduling, was shown by Dhall and Liu to have very poor worst-case performance. A task set may have processor utilization arbitrarily close to 1 and still not be schedulable on  $m$  processors using rate monotonic or earliest-deadline-first scheduling. By comparison, even though optimal partitioning is NP complete, heuristic partitioning algorithms can guarantee schedulability to a much higher processor utilization level than 1. For example, using rate monotonic local scheduling and a simple first-fit-decreasing partitioning heuristic, a utilization level of at least  $m(2^{1/2} - 1)$  is always feasible for an  $m$ -processor system (Oh and Baker, 1998). This is not only much better than the worst case for global scheduling; it is not far from the limit of  $(m + 1)/2$  that applies to all fixed-job-priority multiprocessor scheduling algorithms, partitioned or not (Andersson et al., 2001).

Recently, progress has been made in understanding global multiprocessor scheduling, based on a re-evaluation of the result of Dhall and Liu. Their worst-case example has two kinds of tasks: “heavy” ones, with high ratio of computation time to deadline, and “light” ones, with low ratio of computation time to deadline. It is the mixing of those two kinds of tasks that causes a problem. A scheduling policy that segregates the heavy tasks from the light ones, on disjoint sets of CPU’s, would have no problem with Dhall’s example. Examination of further examples leads one to conjecture that such a hybrid scheduling policy would not miss any deadlines until a fairly high level of CPU utilization is achieved, and might even permit the use of simple utilization-based schedulability tests.

The preemptive scheduling of periodic tasks on multiprocessors was examined in (Andersson et al., 2001), and it was shown that any system of independent periodic tasks for which the utilization of every individual task is at most  $m/(3m - 2)$  can be scheduled successfully on  $m$  processors using rate monotonic scheduling if the total utilization is at most  $m^2/(3m - 1)$ . In (Baruah and Goossens, 2003) there is a similar result, showing that a total utilization of at least  $m/3$  can be achieved if the individual task utilizations do not exceed  $1/3$ . Andersson, Baruah, and Johnsson proposed a hybrid scheduling algorithm, called RM-US[ $m/(3m - 2)$ ], which gives higher priority to tasks with utilizations above  $m/(3m - 2)$ , that is able to successfully schedule *any* set of independent periodic tasks with total utilization up to  $m^2/(3m - 1)$ .

This paper further advances the understanding of global fixed-priority multiprocessor schedulability. The main contribution is a new analysis concept, called the  $(\mu, k - 1)$  *busy interval*. By analyzing the

workload of a  $(\mu, k-1)$  busy interval one can relax several of the assumptions made in the prior analyses cited above, and derive a more general schedulability test. The principal consequences are:

1. Applicability to all fixed priority assignments, rather than just rate monotonic.
2. Applicability to tasks with arbitrary deadlines, rather than just those where deadline equals period.
3. Applicability of the RM-US hybrid model to systems with higher utilization levels, by setting the cut-off between “heavy” and “light” at any desired utilization level, rather than just  $1/3$ .

These results provide a means of verifying the feasibility of task sets on multiprocessor systems that make use of global scheduling, at non-trivial utilization levels and without arbitrary constraints on priorities and deadlines. By reducing the gap in guaranteed-feasible utilization levels between the partitioned and global scheduling approaches, they also suggest that perhaps global scheduling should be given more serious consideration for real-time systems. (The broader question of whether global scheduling is preferable to partitioned scheduling, either theoretically or pragmatically, remains a subject for further research.)

The rest of the paper presents the derivation of the theory. Section 2 defines the problem formally, and outlines the overall approach. Section 3 derives a lower bound on the workload contributions of competing tasks in an interval where a task misses a deadline. Section 4 defines the notion of  $(\mu, k-1)$  busy interval and derives an upper bound on the workload contributions of competing tasks in any such interval. Section 5 combines the upper and lower bounds on workload to obtain schedulability tests, including a red lower bound on the minimum achievable utilization for rate monotonic scheduling. Section 6 compares the performance of those tests on some randomly generated task sets. Section 7 provides an upper bound on the minimum achievable utilization for rate monotonic scheduling. Section 8 reviews in more detail the connections to prior work, including an application to determining an “optimal” value of  $\lambda$  for RM-US $[\lambda]$  scheduling. Section 9 summarizes and concludes.

## 2. Definition of the Problem

Suppose one is given a set of  $N$  simple independent sporadic tasks  $\tau_1, \dots, \tau_N$ , where each task  $\tau_i$  has minimum inter-release time (called *period* for short)  $T_i$ , worst case computation time  $c_i$ , and relative deadline  $d_i$ , where  $c_i \leq d_i$ , and  $c_i \leq T_i$ . Each task generates a sequence of *jobs*, each of whose release

time is separated from that of its predecessor by at least  $T_i$ . No special assumptions are made about the first release time of each task.

Time is represented by the domain of rational numbers. Square brackets and parentheses are used to distinguish whether time intervals include their endpoints. For example the time interval  $[t_1, t_2)$  contains the time values greater than or equal to  $t_1$  and less than  $t_2$ . All of the intervals  $[t_1, t_2]$ ,  $[t_1, t_2)$ ,  $(t_1, t_2]$  and  $(t_1, t_2)$  are said to be of *length*  $t_2 - t_1$ .

The objective of this paper is to formulate a simple test for schedulability of a task set, expressed in terms of the periods, deadlines, and worst-case computation times of the tasks, such that if the test is passed no deadlines will be missed. The problem is approached by analyzing the minimum processor load that is needed over an interval of time to cause a missed deadline.

**DEFINITION 1.** Let  $S = \{\tau_1, \dots, \tau_N\}$  be any task set. A *release-time assignment* for  $S$  is a function  $r : \{1, \dots, N\} \times \mathbb{N} \rightarrow \text{Time}$ . The value  $r(i, j)$  is interpreted as the release time of the  $j$ th job of  $\tau_i$ . All task releases are required to be separated by at least the task period, *i.e.*,  $r(i, j) + T_i \leq r(i, j + 1)$ .

The jobs of each task must be executed sequentially, and all jobs are scheduled on  $m$  identical processors according to a global preemptive fixed-priority policy, where task  $\tau_i$  always has priority over task  $\tau_{i+1}$ . Here “global” means that jobs are assigned to processors dynamically, so that whenever there are  $m$  or fewer jobs ready they will all be executing, and whenever there are more than  $m$  jobs ready there will be  $m$  jobs executing, all with priority higher than or equal to the priorities of the jobs that are not executing.

Since it is not the intent of this paper to compare the efficiency of global versus partitioned scheduling, and for the sake of simplicity of analysis, the abstract computational model does not include any execution time penalty for preemption or for interprocessor task migration. In a real system there will be some penalty for interrupting the execution of a processor, and there may be some penalty for reloading the cache of the new processor with instructions/data of the task (if the cache of the old processor has not already been overwritten by intervening other tasks). The task migration penalty is variable, and in practical applications will add to the margin of error in schedulability analysis that already exists due to variability in task execution times and cache effects due to preemption that occur already on a single processor. However, if one prefers to take the task migration penalty into account explicitly, the analysis presented in this paper could be modified to account for the worst-case preemption and migration costs.

Migration of a job  $J_1$  from processor  $P_1$  to processor  $P_2$  can only occur when the following two events happen in sequence:

1.  $J_1$  must be preempted on  $P_1$ . That is, while  $J_1$  is running, some job  $J_2$  with higher priority than  $J_1$  must be released.
2.  $P_2$  must become available for  $J_1$  before  $P_1$  does. That is, whatever job  $J'_2$  (with higher priority than  $J_1$ ) was running on  $P_2$  when  $J_1$  was preempted by  $J_2$  must complete.

The overhead cost of migration can be accounted for by associating it with either or both of these two events. Assuming the main component of the overhead is paid in lost time on processor  $P_2$ , supposed it is all charged against the execution time of  $J_2$ . That is, it is as if the execution of  $J_2$  were prolonged by the amount of time it takes to migrate  $J_1$  to  $P_2$ . This cost can only be charged once to each job  $J_2$ , since each job completes only once and so can have at most one (migrated) successor on the same processor.

DEFINITION 2. Given a task set  $S$  and a release-time assignment  $r$ , the *work*  $W_i$  done by task  $\tau_i$  over a time interval  $[t - \Delta, t)$  is the actual amount of computation time that is used by jobs of  $\tau_i$  in the interval, and the *load* due to task  $\tau_i$  is  $W_i/\Delta$ . (This definition differs slightly from some other work on demand analysis by counting only work that is *actually executed* in the interval.) Wherever the notation  $W_i$  is used the release time assignment and time interval will be clear from context.

For any integer  $k$ ,  $1 \leq k \leq N$ , the *level  $k$  work* with release-time assignment  $r$  is  $\sum_{i \leq k} W_i$ , and the *level  $k$  load* is  $\sum_{i \leq k} W_i/\Delta$ .

DEFINITION 3. For a given task set and release time assignment, a *first missed deadline* is a time  $t$  at which some task misses a deadline and such that no task misses a deadline before  $t$ .

If a task set misses a deadline for some release time assignment then it has a unique first missed deadline for that release time assignment. If one can find a lower bound on the processor load over an interval leading up to every first missed deadline, and one can guarantee that a given set of tasks could not possibly generate so much load in such an interval, that would be sufficient to serve as a proof of schedulability.

### 3. Lower Bound on Load

One can establish a lower bound on the level  $k - 1$  load of the interval ending at a first missed deadline of task  $\tau_k$  and starting with the release time of the corresponding job by observing that, since the job does not complete by the end of the interval, the lengths of all the subintervals in which the job does not execute, plus the execution time of the job, must exceed the length of the interval between its release time and deadline. This fact is well known and is the basis of the prior analysis in (Phillips et al., 1997) and others. It is illustrated in Figure 1 for the case where  $m = 3$  and  $d_k \leq T_k$ . The diagonally shaded rectangles indicate blocks of time during which  $\tau_k$  executes. The dotted rectangles indicate times during which all  $m$  processors must be busy executing other jobs that contribute to the load for this interval. It is easy to see that the total level  $k - 1$  work of the interval  $[t - d_k, t)$  must be at least  $m(d_k - x)$ , where  $x < c_k$  is the amount of time that  $\tau_k$  executes in the interval.




Figure 1. All processors must be busy whenever  $\tau_k$  is not executing.




Figure 2. More than one job of  $\tau_k$  may execute in a  $\tau_k$ -busy interval if  $d_k \geq T_k$ .

To allow for the possibility that  $d_k \geq T_k$ , one needs to consider intervals that may include more than one job of  $\tau_k$ . Figure 2 shows the release times and deadlines of two such jobs  $\tau_{k,1}$  and  $\tau_{k,2}$ . The job  $\tau_{k,1}$  is delayed by two blocks of higher priority interference. This interference is not enough to cause  $\tau_{k,1}$  to

miss its deadline, but (because jobs of the same task must be executed sequentially) it delays the start of  $\tau_{k,2}$  enough to cause that job to miss its deadline.

DEFINITION 4. A job is *backlogged* at a time  $t$  if it is released before time  $t$  and has nonzero execution time remaining at time  $t$ . A task is backlogged if it has a backlogged job.

For any task  $\tau_k$ , a time interval  $[t', t)$  is  $\tau_k$ -*busy* if  $t$  is a missed deadline for  $\tau_k$  and there are backlogged jobs of  $\tau_k$  continually throughout the interval  $(t', t)$ .

LEMMA 5. For a given a set of tasks  $S$  and a given release-time assignment  $r$ , if  $\tau_k$  is backlogged at time  $t$  then there is a unique  $\tau_k$ -busy interval  $[t - \Delta, t)$  such that:

1. There are no backlogged jobs of  $\tau_k$  at time  $t - \Delta$ .
2. There is a job of  $\tau_k$  released at time  $t - \Delta$ .

*Proof.* Let  $t'$  be the latest time before  $t$  at which  $\tau_k$  is not backlogged. There must be such a time, since  $\tau_k$  is not backlogged at the system start time. It must be that  $t'$  is a release time of  $\tau_k$ . The interval  $[t', t)$  satisfies the definition of  $\tau_k$ -busy. The value  $t'$  is unique, since  $\tau_k$  is not backlogged at time  $t'$  and  $\tau_k$  is backlogged at all times from  $t'$  through  $t$ .  $\square$

The unique interval guaranteed by Lemma 5 is called the *maximal*  $\tau_k$ -busy interval ending at  $t$ . If  $d_k \leq T_k$  this interval cannot be longer than  $d_k$  unless  $\tau_k$  has missed a deadline prior to  $t$ , but if  $d_k > T_k$  the interval may be arbitrarily longer than  $d_k$ . For example, consider a system of  $m + 1$  tasks, such that  $T_1 = \dots = T_{m+1} = 1$ ,  $d_1 = \dots = d_m = 1$ ,  $d_{m+1} = 2$ ,  $c_1 = \dots = c_m = \epsilon$ , and  $c_{m+1} = 1$ . The first  $m$  tasks create a block of interference of duration  $\epsilon$  for every release of  $\tau_{m+1}$ , so that each successive job of  $\tau_{m+1}$  completes later by  $\epsilon$ . The first job of task  $\tau_{m+1}$  to miss its deadline will be the  $j$ th job, where  $j$  is the least integer greater than  $1/\epsilon$ . It will miss its deadline at time  $jT_{m+1} + d_{m+1} = j + 1$ , since  $j\epsilon > d_{m+1} - T_{m+1} = 1$ . One can make  $j$  arbitrarily large by choosing  $\epsilon$  small enough.

LEMMA 6 (lower bound on load). If  $t$  is a first missed deadline of  $\tau_k$  then the maximal  $\tau_k$ -busy interval  $[t - \Delta, t)$  has level  $(k-1)$  load greater than  $m(1 - \frac{c_k}{\min\{T_k, d_k\}})$ .

*Proof.* Let  $x$  be the amount of time  $\tau_k$  executes in the interval  $[t - \Delta, t)$ . Since  $\tau_k$  is continually backlogged over the interval, the only times that  $\tau_k$  does not execute are the times that all  $m$  processors are executing jobs of tasks with higher priority than  $\tau_k$ . The level  $(k-1)$  work of  $[t - \Delta, t)$  must be at least equal to the work of these higher priority tasks, *i.e.*,

$$\sum_{i < k} \frac{W_i}{\Delta} \geq \frac{m(\Delta - x)}{\Delta} = m\left(1 - \frac{x}{\Delta}\right) \quad (1)$$

Let  $j$  be the number of jobs of  $\tau_k$  that execute in the interval. The amount  $x$  of time that  $\tau_k$  executes in the interval is bounded as follows:

$$x < jc_k \quad (2)$$

Since  $\tau_k$  is not backlogged at the start of the interval, all of the  $j$  jobs are released on or after  $t - \Delta$  and not later than  $t - d_k$  (because  $t$  is a missed deadline), and the release times are separated by at least  $T_k$ , so:

$$(j - 1)T_k + d_k \leq \Delta \quad (3)$$

It follows from (1-3) that

$$\sum_{i < k} \frac{W_i}{\Delta} > m\left(1 - \frac{jc_k}{(j - 1)T_k + d_k}\right)$$

Let  $f : \mathbb{N} \rightarrow \text{Time}$  be the function defined by  $f(j) = \frac{jc_k}{(j-1)T_k + d_k}$ . The objective is to find an upper bound for  $f(j)$ , subject to the available constraints. There are two cases to consider:

1. If  $d_k \leq T_k$  then  $\frac{c_k}{d_k} \geq \frac{c_k}{T_k}$ , and  $f$  is monotonically non-increasing with respect to  $j$ . To see this, take any two fractions  $\frac{u}{v} \geq \frac{w}{x}$ , where  $u, v, w$ , and  $x$  are positive numbers. It follows that

$$\begin{aligned} ux &\geq vw \\ uv + ux &\geq uv + vw \\ u(v + x) &\geq v(u + w) \\ \frac{u}{v} &\geq \frac{u + w}{v + x} \end{aligned}$$

Letting  $u = jc_k$ ,  $v = (j - 1)T_k + d_k$ ,  $w = c_k$ , and  $x = T_k$ , it follows that

$$f(j) = \frac{u}{v} \geq \frac{u + w}{v + x} = f(j + 1)$$



That is,  $f$  is monotonically non-increasing.

Since  $f$  is monotonically non-increasing with respect to  $j$  and  $j \geq 1$ , it follows that

$$f(j) \leq f(1) = \frac{c_k}{d_k}$$

2. If  $d_k > T_k$  then  $\frac{c_k}{d_k} < \frac{c_k}{T_k}$ ,  $f$  is increasing with respect to  $j$ , and

$$f(j) \leq \lim_{j \rightarrow \infty} f(j) = \frac{c_k}{T_k}$$

Putting the above two cases together, it follows that:

$$\sum_{i < k} \frac{W_i}{\Delta} > m \left( 1 - \frac{c_k}{\min\{d_k, T_k\}} \right)$$

□

#### 4. Upper Bound on Load

The next step is to derive an upper bound on the level  $(k-1)$  load of an interval leading up to a first missed deadline of task  $\tau_k$ . If one can find such an upper bound  $\beta_k$ , it will follow that schedulability of a task system can be guaranteed by checking that  $\beta_k$  is less than the minimum level  $(k-1)$  load needed to cause a missed deadline of  $\tau_k$ . The upper bound will be defined as the sum of individual upper bounds on the load due to each task that can preempt  $\tau_k$  in the interval.

The work  $W_i$  done by task  $\tau_i$  in any interval  $[t - \Delta, t)$  is clearly bounded by the length  $\Delta$  of the interval and may include:

1. a portion of the execution times of zero or more jobs that are released before  $t - \Delta$  but are unable to complete by that time, which are called *carried-in* jobs;
2. the full execution times  $c_i$  of zero or more jobs that are released on or after time  $t - \Delta$  and complete by time  $t$ ;
3. a portion of the execution time of at most one that is released at some time  $t - \delta$ ,  $0 < \delta \leq \Delta$ , but is unable to complete by time  $t$ .

To bound the size of the carried-in contribution of  $\tau_i$  the maximal  $\tau_k$ -busy interval is extended downward as far as possible while still maintaining the level  $(k-1)$  load that caused  $\tau_k$  to miss its deadline at time  $t$ .

**DEFINITION 7.** A time interval  $[t', t)$  is  $(\mu, k-1)$  *busy* if the level  $(k-1)$  load is greater than  $\mu$ . It is a *maximal*  $(\mu, k-1)$  busy interval if it is  $(\mu, k-1)$  busy and there is no  $t'' < t'$  such that  $[t'', t)$  is also  $(\mu, k-1)$  busy.

**LEMMA 8.** *If  $t$  is a first missed deadline of  $\tau_k$  and  $0 < \mu \leq m(1 - \frac{c_k}{\min\{T_k, d_k\}})$  then there is a unique maximal  $(\mu, k-1)$  busy interval  $[t - \hat{\Delta}, t)$ , and  $\hat{\Delta} \geq d_k$ .*

*Proof.* By Lemma 5, there is a unique maximal  $\tau_k$ -busy interval  $[t', t)$ . Since  $t$  is a missed deadline for  $\tau_k$ , the length of this  $\tau_k$ -busy interval is at least  $d_k$ . By Lemma 6, the level  $(k-1)$  load of this interval is greater than  $\mu$ . Therefore, the set of all starting points  $t'' \leq t'$  of  $(\mu, k-1)$  busy intervals  $[t'', t)$  is non-empty. This set must have a minimal member, since there are no backlogged jobs at the start time of the system. Let  $\hat{\Delta} = t - t''$  for this minimum value  $t''$  and the lemma is satisfied.

□

**DEFINITION 9.** Given a task set  $S$ , a release-time assignment  $r$ , a task  $\tau_k$  that has a first missed deadline at time  $t$ , and a value  $\mu$ , the unique interval  $[t - \hat{\Delta}, t)$  that is guaranteed by Lemma 8 is called the  $(\mu, k-1)$  *busy interval* of  $\tau_k$ . From this point on, let  $[t - \hat{\Delta}, t)$  denote such an interval.

The next step in the analysis is to find an upper bound on the load  $\frac{W_i}{\hat{\Delta}}$  due to each task  $\tau_i$  in a  $(\mu, k-1)$  busy interval. The only interesting cases are those where  $\tau_i$  can make a contribution to the level  $(k-1)$  load, that is, where  $i < k$ .

**LEMMA 10 (upper bound on load).** *If  $t$  is a first missed deadline of task  $\tau_k$ ,  $0 < \mu \leq m(1 - \frac{c_k}{\min\{T_k, d_k\}})$ , and  $[t - \hat{\Delta}, t)$  is the corresponding  $(\mu, k-1)$  busy interval, then the contribution  $W_i/\hat{\Delta}$  of each task  $\tau_i$ ,  $i < k$ , to the load of the interval is strictly bounded above by the function  $\beta_{\mu, k}(i)$  defined to be the lesser of the value 1 and the value given by the expression in Table I.*

*Proof.*

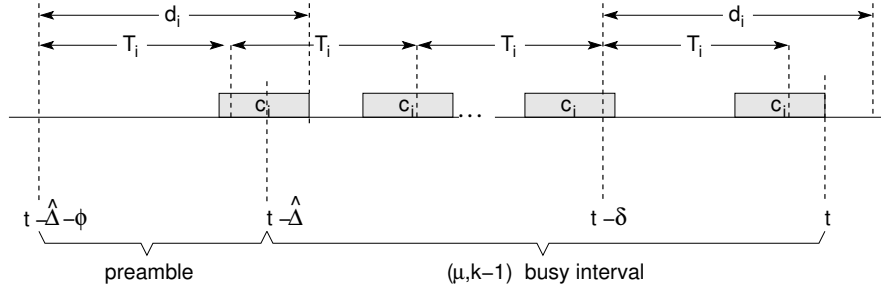
Table I. Cases for  $\beta_{\mu,k}(i)$ , the upper bound on  $\frac{W_i}{\Delta}$ .

Case	$\beta_{\mu,k}(i)$
$\frac{m-\mu}{m-1} \geq \frac{c_i}{T_i}$	$\min\{1, \frac{c_i}{T_i}(1 + \frac{T_i - c_i}{d_k})\}$
$\frac{m-\mu}{m-1} < \frac{c_i}{T_i}$	$\min\{1, \frac{c_i}{T_i}(1 + \frac{T_i - c_i}{d_k}) + \frac{d_i}{d_k}(\frac{c_i}{T_i} - \frac{m-\mu}{m-1})\}$

Let  $j$  be the number of jobs of  $\tau_i$  that execute in the interval. If  $j = 0$  the lemma is satisfied trivially. Therefore, it is only necessary to consider the case where  $j \geq 1$ .

Let  $\epsilon > 0$  be the amount of time that the last of these  $j$  jobs of  $\tau_i$  executes in the interval, and let  $t - \delta$  be the release time of this job. Observe that  $\epsilon \leq \delta$  and  $\epsilon \leq c_i$ .

Let  $t - \hat{\Delta} - \phi$  be the release time of the first job of  $\tau_i$  that is released before  $t - \hat{\Delta}$  and executes in the interval, if such exists; otherwise, let  $\phi = 0$ .

Figure 3. Preamble of  $(\mu, k-1)$  busy interval.

If  $\phi > 0$  the interval  $[t - \hat{\Delta} - \phi, t - \hat{\Delta})$  is non-empty and must be  $\tau_i$ -busy. Call this the *preamble* with respect to  $t_i$  of  $[t - \hat{\Delta}, t)$ . In this case reasoning similar to that of Lemma 6 can be used to bound the amount of work that  $\tau_i$  may carry from the preamble into  $[t - \hat{\Delta}, t)$ , as follows.

Let  $x$  be the total amount of time spent executing jobs of  $\tau_i$  in the preamble and let  $y$  be the sum of the lengths of all the subintervals within the preamble where all  $m$  processors are simultaneously executing jobs that preempt  $\tau_i$ . Since  $\tau_i$  can execute when and only when there are less than  $m$  processors executing jobs that preempt  $\tau_i$ ,  $y = \phi - x$ . It follows that the total amount of level  $i$  work executed in the preamble must be at least  $my + x$ . Putting this together with the fact that the interval  $[t - \hat{\Delta}, t)$  is  $(\mu, k-1)$  busy, one can conclude that the level  $(k-1)$  work of the entire interval  $[t - \hat{\Delta} - \phi, t)$  is at least  $\mu\hat{\Delta} + my + x$ .

By the definition of maximal  $(\mu, k-1)$  busy interval, the level  $(k-1)$  load of the interval  $[t - \hat{\Delta} - \phi, t)$  is less than  $\mu$ . It follows that if  $\phi > 0$  then

$$\begin{aligned}\mu\hat{\Delta} + my + x &< \mu(\phi + \hat{\Delta}) \\ m(\phi - x) + x = my + x &< \mu\phi \\ x(1 - m) &< \phi(\mu - m) \\ x &> \frac{m - \mu}{m - 1}\phi\end{aligned}$$

If  $\phi = 0$  the preamble is empty and so it follows that in all cases  $x > \frac{m-\mu}{m-1}\phi$ .

By subtracting out  $\frac{m-\mu}{m-1}\phi$  as a lower bound on the work  $x$  of  $\tau_i$  done in the preamble, one can obtain the following bound on the load due to  $\tau_i$  in the  $(\mu, k-1)$  busy interval  $[t - \hat{\Delta}, t)$ :

$$\frac{W_i}{\hat{\Delta}} \leq \frac{c_i(j-1) + \epsilon - \frac{m-\mu}{m-1}\phi}{\hat{\Delta}} \quad (4)$$

Because of the minimum separation constraint,

$$\begin{aligned}(j-1)T_i + \delta &\leq \hat{\Delta} + \phi \\ (j-1) &\leq \frac{\hat{\Delta} + \phi - \delta}{T_i}\end{aligned}$$

and so

$$\frac{W_i}{\hat{\Delta}} \leq \frac{c_i \frac{\hat{\Delta} + \phi - \delta}{T_i} + \epsilon - \frac{m-\mu}{m-1}\phi}{\hat{\Delta}}$$

The expression on the right of the above inequality is increasing with respect to  $\epsilon$ . Since  $\epsilon \leq c_i$ , the value of the expression is never greater than when  $\epsilon = c_i$ .

The same expression is decreasing with respect to  $\delta$ . By definition,  $\delta \geq \epsilon$ . It follows that the maximum value of the expression on the right side of the inequality is never greater than when  $\delta = \epsilon = c_i$ .

$$\frac{W_i}{\hat{\Delta}} \leq \frac{c_i \frac{\hat{\Delta} + \phi - c_i}{T_i} + c_i - \frac{m-\mu}{m-1}\phi}{\hat{\Delta}} = \frac{c_i}{T_i} \left(1 + \frac{T_i - c_i}{\hat{\Delta}}\right) + \frac{\phi \left(\frac{c_i}{T_i} - \frac{m-\mu}{m-1}\right)}{\hat{\Delta}} \quad (5)$$

The value of the expression on the right side of the above inequality may be increasing or decreasing with respect to  $\phi$ , depending on whether  $\frac{c_i}{T_i} \leq \frac{m-\mu}{m-1}$ .

**Case 1:** If  $\frac{c_i}{T_i} \leq \frac{m-\mu}{m-1}$  the value of the expression on the right of inequality (5) is non-increasing with respect to  $\phi$ , and since  $\phi \geq 0$ , the global maximum is achieved when  $\phi = 0$ . It follows that

$$\frac{W_i}{\hat{\Delta}} < \frac{c_i}{T_i} \left(1 + \frac{T_i - c_i}{\hat{\Delta}}\right) \leq \frac{c_i}{T_i} \left(1 + \frac{T_i - c_i}{d_k}\right) \quad (6)$$

**Case 2:** If  $\frac{c_i}{T_i} > \frac{m-\mu}{m-1}$  the value of the expression on the right of inequality (5) is increasing with respect to  $\phi$ . Since there are no missed deadlines prior to  $t$  and the job released at time  $t - \hat{\Delta} - \phi$  does not complete by  $t - \hat{\Delta}$ ,  $\phi < d_i$ . It follows that

$$\frac{W_i}{\hat{\Delta}} < \frac{c_i}{T_i} \left(1 + \frac{T_i - c_i}{\hat{\Delta}}\right) + \frac{d_i \left(\frac{c_i}{T_i} - \frac{m-\mu}{m-1}\right)}{\hat{\Delta}}$$

Since  $\frac{c_i}{T_i} > \frac{m-\mu}{m-1}$  and  $\hat{\Delta} \geq d_k$ ,

$$\frac{W_i}{\hat{\Delta}} < \frac{c_i}{T_i} \left(1 + \frac{T_i - c_i}{d_k}\right) + \frac{d_i \left(\frac{c_i}{T_i} - \frac{m-\mu}{m-1}\right)}{d_k} \quad (7)$$

In addition, since  $W_i$  is the actual execution time of task  $\tau_i$  in the interval and a task may not execute on more than one processor at a time, we also have  $\frac{W_i}{\Delta} \leq 1$  in all cases. Taking this into account, the upper bounds for  $\frac{W_i}{\Delta}$  derived in each of the above cases correspond to the expressions in Table I.

□

## 5. Schedulability Tests

Based on the above analysis, one can now prove the following theorem, which provides a sufficient condition for schedulability.

**THEOREM 11** ( $O(N^3)$  schedulability test). *A set of sporadic tasks  $S = \{\tau_1, \dots, \tau_N\}$  is schedulable on  $m$  processors using preemptive fixed-priority scheduling if, for every task  $\tau_k$ ,  $k = m + 1, \dots, N$ , there exists a positive value  $\mu \leq m \left(1 - \frac{c_k}{\min\{T_k, d_k\}}\right)$  such that*

$$\sum_{i=1}^{k-1} \beta_{\mu,k}(i) \leq \mu \quad (8)$$

where  $\beta_{\mu,k}(i)$  is as defined in Table I.

*Proof.* The proof is by contradiction. Suppose there are a task set  $S$  and a release time assignment  $r$  for which some task  $\tau_k$  has a first missed deadline at time  $t$ . By the priority ordering,  $k > m$ . Let  $[t - \hat{\Delta}, t)$  be the  $(\mu, k-1)$  busy interval guaranteed by Lemma 8. By the definition of  $(\mu, k-1)$  busy,

$$\sum_{i=1}^{k-1} \frac{W_i}{\hat{\Delta}} > \mu$$

By Lemma 10,  $\frac{W_i}{\hat{\Delta}} \leq \beta_{\mu,k}(i)$ , for  $i = 1, \dots, k-1$ . It follows that

$$\sum_{i=1}^{k-1} \beta_{\mu,k}(i) \geq \sum_{i=1}^{k-1} \frac{W_i}{\hat{\Delta}} > \mu$$

The above is a contradiction of (8).

□

To use the above condition as a schedulability test, it might seem necessary to consider all possible values of  $\mu$  for each  $k$ . On the contrary, the only values of  $\mu$  that need to be considered are the upper bound and the points at which the function  $\beta_{\mu,k}(i)$  is discontinuous with respect to the parameter  $\mu$ . That is, at the points

$$\mu_i = m - \frac{c_i}{T_i}(m-1)$$

for  $i = 1, \dots, k$ , and

$$\mu_{\max} = m \left(1 - \frac{c_k}{\min\{T_k, d_k\}}\right)$$

The computational complexity of checking (8) for all such values of  $\mu$  for each value of  $k$  is  $O(N^3)$ . If one is willing to sacrifice some precision for a faster test, one can check fewer values of  $\mu$ , resulting in an  $O(N^2)$  schedulability test. There are many possible heuristics for choosing a subset of the values of  $\mu$  to consider. Any one of the values of  $\mu$  checked in the  $O(N^3)$  test would be sufficient. If only one value of  $\mu$  is to be checked, using the largest possible value seems to give the best result in most cases.

**COROLLARY 12** ( $O(N^2)$  test). *A set of sporadic tasks  $\tau_1, \dots, \tau_N$  is schedulable on  $m$  processors if for every task  $\tau_k$ ,  $k = m+1, \dots, N$ ,*

$$\sum_{i=1}^{k-1} \beta_{\mu,k}(i) \leq \mu \tag{9}$$

where  $\lambda_k = \frac{c_k}{\min\{T_k, d_k\}}$  and  $\mu = m(1 - \lambda_k)$ .

*Proof.*

The proof is by direct application of Theorem 11 for  $\mu = m(1 - \lambda_k)$ .

□

If one is willing to sacrifice some more precision for a still faster test, there is an  $O(N)$  test.

**COROLLARY 13** ( $O(N)$  schedulability test). *A set of sporadic tasks  $\tau_1, \dots, \tau_N$  is schedulable on  $m$  processors if*

$$\sum_{i=1}^{N-1} \min\left\{1, \frac{c_i}{T_i} \left(1 + \frac{T_i - c_i}{d_{\min}}\right)\right\} \leq m(1 - \lambda_{\max}) \quad (10)$$

where  $\lambda_i = \frac{c_i}{\min\{T_i, d_i\}}$ ,  $\lambda_{\max} = \max\{\lambda_i \mid i = 1, \dots, N\}$ , and  $d_{\min} = \min\{d_i \mid i = 1, \dots, N\}$ .

*Proof.*

Corollary 13 is proved by application of Theorem 11. Consider any task  $\tau_k$ . Let  $\mu = m(1 - \lambda_{\max})$ . It follows that  $\mu \leq m(1 - \lambda_k)$ , and  $\frac{m-\mu}{m-1} > \frac{c_i}{T_i}$  for all  $i$ . Therefore, since  $k \leq N$  and  $d_k \geq d_{\min}$ , if (10) is satisfied then

$$\sum_{i=1}^{k-1} \beta_{\mu, k}(i) \leq \sum_{i=1}^{N-1} \min\left\{1, \frac{c_i}{T_i} \left(1 + \frac{T_i - c_i}{d_{\min}}\right)\right\} \leq \mu$$

By Theorem 11, the task set must be schedulable. □

If one assumes the deadline of each task is equal to its period Theorem 11 also yields a lower bound on the minimum achievable utilization for rate monotonic scheduling.

**COROLLARY 14** (utilization test). *A set of sporadic tasks, all with deadline equal to period and utilization less than or equal to one, is guaranteed to be schedulable on  $m$  processors using preemptive rate monotonic scheduling if*

$$\sum_{i=1}^N \frac{c_i}{T_i} \leq \frac{m}{2}(1 - u_{\max}) + u_{\min} \quad (11)$$

where  $u_{\max} = \max\{\frac{c_i}{T_i} \mid i = 1, \dots, N\}$  and  $u_{\min} = \min\{\frac{c_i}{T_i} \mid i = 1, \dots, N\}$ .

*Proof.*

Since deadline equals period,  $\lambda_i = \frac{c_i}{T_i} = u_i$  for all  $i$ . The  $\tau_k$  be any task. Let  $\mu = m(1 - u_{\max})$ . It follows that  $\mu \leq m(1 - \lambda_k)$ , and  $\frac{m-\mu}{m-1} > \frac{c_i}{T_i}$  for all  $i$ . Since  $c_i < T_i \leq d_k$  and  $d_k = T_k$ ,

$$\min\left\{1, \frac{c_i}{T_i}\left(1 + \frac{T_i - c_i}{d_k}\right)\right\} = \frac{c_i}{T_i}\left(1 + \frac{T_i - c_i}{T_k}\right)$$

By the rate monotonic ordering of task priorities,  $T_i \leq T_k$  for  $i < k$ , and so

$$\frac{c_i}{T_i}\left(1 + \frac{T_i - c_i}{T_k}\right) \leq \frac{c_i}{T_i}\left(1 + \frac{T_i - c_i}{T_i}\right) = 2\frac{c_i}{T_i} - \left(\frac{c_i}{T_i}\right)^2 < 2\frac{c_i}{T_i}$$

Therefore,

$$\sum_{i=1}^{k-1} \beta_{\mu,k}(i) = \sum_{i=1}^{k-1} \frac{c_i}{T_i}\left(1 + \frac{T_i - c_i}{T_k}\right) \leq \sum_{i=1}^{k-1} 2\frac{c_i}{T_i} \leq 2\sum_{i=1}^N \frac{c_i}{T_i} - 2\sum_{i=k}^N \frac{c_i}{T_i}$$

If condition (11) is satisfied it follows that, for  $k = 1, \dots, N$ ,

$$\sum_{i=1}^{k-1} \beta_{\mu,k}(i) \leq 2\left(\frac{m}{2}(1 - u_{\max}) + u_{\min}\right) - 2u_{\min} = \mu$$

and by Theorem 11, the task set must be schedulable.

□

Note that (Baker, 2003a) and (Baker, 2003b) claimed a stronger result, in which the  $u_{\min}$  term of the utilization bound above is replaced by  $u_{\max}$ . The proof given for that stronger result had a flaw, based on failure to properly distinguish  $\lambda_k$  from  $\lambda_{max}$ .

Theorem 11 and its corollaries provide three general fixed-priority schedulability tests, one of complexity  $O(N^3)$ , one of complexity  $O(N^2)$ , and one of complexity  $O(N)$ , and a utilization-bound test for rate monotonic scheduling in the case where deadline equals period.



## 6. Simulations

It is natural to wonder how the performances of these tests compare. In order to address this question, simulation experiments were conducted, using pseudo-randomly generated task sets.

The performance measure chosen for these simulations is the ratio  $A/B$ , where  $A$  is the number of task sets that a given test verifies as schedulable and  $B$  is the number of task sets tested. (Note that  $B$  is not the number of *schedulable* task sets, since the only known necessary and sufficient algorithms for totally determining general schedulability are computationally impractical.)

Simulations were conducted for various numbers of processors, task sizes, and priority assignment rules, and with various algorithms for randomly generating the task periods, deadlines, and execution times. In the interest of brevity, just a few results are presented here.

The procedure was to generate a random set of  $m + 1$  tasks, run all the schedulability tests on that set, then add another randomly generated task to the set, run all the schedulability tests on the new set, *etc.*, until the set grew to a size that it missed a deadline when simulated with all tasks released together at time zero. The procedure was then repeated. This method of generating random task sets produces a fairly uniform distribution of total utilizations. For example, the histogram in Figure 4 is of the total utilization over one collection of 1,000,000 task sets. This collection of task sets had individual task utilizations approximately uniformly distributed over the range 0.05 to 0.2, with a few high and low outliers due to the use of whole integers for all time values. The distribution is sparse at the low end because sets with four or fewer tasks were thrown out of the sample.

The histograms in Figures 5 and 4 show the performance of the schedulability tests on two representative experiments, corresponding to the same collection of task sets as Figure 4. The plots show the fraction of the total number of task sets that each schedulability test reports to be schedulable, for various utilization levels, expressed as percentages. The line labeled “Sim” is provided as a coarse upper bound on the fraction of the task sets that might be schedulable, obtained by simulating the execution when all the tasks are initially released at time zero. This coarse bound is used because there is no known computationally feasible algorithm for determining with certainty whether or not each task set is schedulable.

The higher-complexity tests clearly are able to identify a larger number of schedulable task sets at the higher utilization levels. In all experiments, the  $O(N)$  test and utilization bound tests performed much worse than the  $O(N^2)$  and  $O(N^3)$  tests. The  $O(N^3)$  test performed consistently better than the  $O(N^2)$




Figure 4. Distribution of total utilizations in experiments with four processors.




Figure 5. Deadline equals period and rate monotonic priorities.

test, but the margin was not large and the performance was essentially identical on many experiments, including those with wider variations in individual task utilizations and deadlines.




Figure 6. Random deadlines and deadline monotonic priorities.

## 7. An Upper Bound on Minimum Achievable Utilization

Corollary 14 provides a lower bound on the minimum achievable utilization for rate monotonic scheduling. The following theorem provides an upper bound.

**THEOREM 15** (upper bound on minimum achievable RM utilization). *There exist task sets that are not feasible with preemptive RM scheduling on  $m$  processors and have utilization arbitrarily close to  $u_{\max} + m \ln(\frac{2}{1+u_{\max}})$ , where  $u_{\max}$  is the maximum single-task utilization.*

*Proof.*

The goal of the proof is to show that there is an infinite sequence of task sets  $S'_{m,1}, S'_{m,2}, S'_{m,3}, \dots$ , each of which is not schedulable on  $m$  processors and for which

$$\lim_{k \rightarrow \infty} U(S_{m,k}) = u_{\max} + m \ln\left(\frac{2}{1+u_{\max}}\right)$$

Instead of constructing such a sequence of unschedulable task sets directly, it will be sufficient to construct an infinite sequence of *barely schedulable* task sets  $S_{m,1}, S_{m,2}, S_{m,3}, \dots$ , whose utilizations converge to the desired limit. Then, for each  $k$ , an unschedulable task set  $S'_{m,k}$  can be obtained from  $S_{m,k}$  by increasing the execution time  $c_N$  of the lowest-priority task by an amount that converges to zero

for sufficiently large  $k$ . Since the limit of the amounts by which the task execution times are increased is zero, the limit of the utilizations of the sequence of unschedulable task sets  $S'_{m,k}$  is the same as the limit of the utilizations of the barely schedulable task sets  $S_{m,k}$ . Therefore, to prove the theorem it is sufficient to construct the sequence  $S_{m,1}, S_{m,2}, S_{m,3}, \dots$  of barely schedulable task sets.

The construction of the barely schedulable task sets and the proof of the limiting utilization are derived from (Liu and Layland, 1973). The differences are that there are  $m$  processors instead of one, and the utilization of the longest-period task is bounded by  $u_{\max}$ .

Let  $p_i = \left(\frac{2}{1+u_{\max}}\right)^{\frac{i}{k}}$  for  $i = 1, \dots, k+1$ . The periods of the  $N = mk + 1$  tasks in  $S_{m,k}$  are defined to be

$$T_{(i-1)m+1} = T_{(i-1)m+2} = \dots = T_{i \cdot m} = p_i \text{ for } i = 1, \dots, k$$

$$T_N = p_{k+1}$$

and the execution times of the tasks are defined to be

$$c_{(i-1)m+1} = c_{(i-1)m+2} = \dots = c_{i \cdot m} = p_{i+1} - p_i \text{ for } i = 1, \dots, k$$

$$c_N = 2p_1 - p_{k+1} = T_N - 2 \sum_{i=1}^k (p_{i+1} - p_i)$$

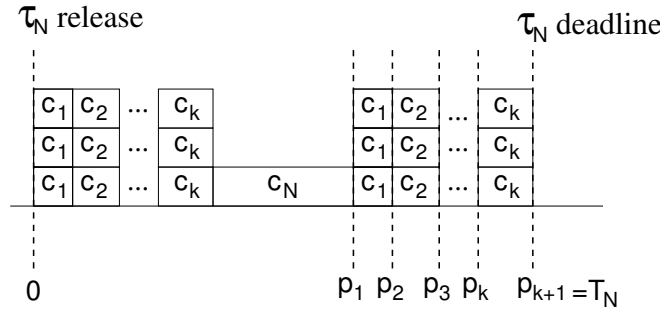


Figure 7. Task set that is barely schedulable.

As shown schematically in Figure 7, when all the tasks are released together at time zero they execute as  $k$  blocks of  $m$ , with task  $\tau_N$  being forced to execute during the time the other tasks are idle. Such a task set is barely schedulable. Any increase in the execution time of task  $\tau_N$  will cause it to miss its deadline. The utilization of such a task set is

$$U(S_{m,k}) = \frac{2p_1 - p_{k+1}}{p_{k+1}} + \sum_{i=1}^k m \frac{p_{i+1} - p_i}{p_i}$$

$$\begin{aligned}
&= 2\left(\prod_{i=1}^k \frac{p_i}{p_{i+1}}\right) - 1 + m\left(\sum_{i=1}^k \frac{p_{i+1}}{p_i}\right) - mk \\
&= 2\left(\prod_{i=1}^k \left(\frac{1+u_{\max}}{2}\right)^{\frac{1}{k}}\right) - 1 + m\left(\sum_{i=1}^k \left(\frac{2}{1+u_{\max}}\right)^{\frac{1}{k}}\right) - mk \\
&= 2\left(\frac{1+u_{\max}}{2}\right) - 1 + mk\left(\frac{2}{1+u_{\max}}\right)^{\frac{1}{k}} - mk \\
&= u_{\max} + mk\left(\left(\frac{2}{1+u_{\max}}\right)^{\frac{1}{k}} - 1\right)
\end{aligned}$$

L'Hôpital's Rule can be applied to find the limit of the above expression for large  $k$ , which is

$$\lim_{k \rightarrow \infty} U(S_{m,k}) = u_{\max} + m \ln\left(\frac{2}{1+u_{\max}}\right)$$

□

Of course, the actual minimum achievable RM utilization is somewhere between the lower bound given by Corollary 14 and the upper bound given by Theorem 15. The values of  $u_{\max} + m \ln\left(\frac{2}{1+u_{\max}}\right)$  (upper) and  $\frac{m}{2}(1 - u_{\max})$  (lower) are plotted in Figure 8, for  $m$  between 1 and 32. Though this paper only shows that the minimum achievable RM utilization lies between these two bounding functions, the author suspects (for reasons explained in the next section) that the actual minimum achievable utilization may correspond exactly to the higher of these two bounds.

## 8. Relation to Prior Work

In (Andersson et al., 2001), a periodic task set  $\{\tau_1, \tau_2, \dots, \tau_N\}$  is defined to be a *light system on  $m$  processors* if it satisfies the following properties:

1.  $\sum_{i=1}^N \frac{c_i}{T_i} \leq \frac{m^2}{3m-2}$
2.  $\frac{c_i}{T_i} \leq \frac{m}{3m-2}$ , for  $1 \leq i \leq n$ .

They then proved the following theorem:

**THEOREM 16** (Andersson, Baruah, Jonsson). *Any periodic task system that is light on  $m$  processors is scheduled to meet all deadlines on  $m$  processors by the preemptive rate monotonic scheduling algorithm.*




Figure 8. Upper and lower bounds on minimum achievable utilization

The above result is narrower but slightly tighter than Corollary 14. If one takes  $u_{\max} = m/(3m - 2)$ , it follows that the system of tasks is schedulable to meet deadlines if

$$\sum_{i=1}^N \frac{c_i}{T_i} \leq \frac{m}{2} \left(1 - \frac{m}{3m - 2}\right) = \frac{m^2 - m}{3m - 2}$$

In (Baruah and Goossens, 2003) there is the following similar result.

**COROLLARY 17** (Baruah & Goossens). *A set of tasks, all with deadline equal to period, is guaranteed to be schedulable on  $m$  processors using rate monotonic scheduling if  $\frac{c_i}{T_i} \leq 1/3$  for  $i = 1, \dots, N$  and*

$$\sum_{i=1}^N \frac{c_i}{T_i} \leq m/3$$

This is a special case of Corollary 14. If one takes  $u_{\max} = 1/3$ , it follows that the system of tasks is schedulable to meet deadlines if

$$\sum_{i=1}^{N-1} \frac{c_i}{T_i} \leq \frac{m}{2}(1 - 1/3) + u_{\min} \leq m/3$$

The results presented in this paper generalize and extend the above cited results in the following ways:

1. Theorem 11 can be applied to tasks with arbitrary deadlines. This is important for systems where some tasks have tighter deadlines, due to bounded jitter requirements, and other tasks have looser deadlines, due to buffering.
2. Theorem 11 can be applied to any set of sporadic tasks, without an arbitrary upper bound on individual task utilizations.
3. Corollary 14 can be applied to task systems with  $u_{\max} > \frac{m}{3m-2}$ , and for task systems with  $u_{\max} < \frac{m-2}{3m-2}$  it provides a higher utilization bound than  $\frac{m^2}{3m-2}$ .

In (Andersson et al., 2001), the following hybrid scheduling algorithm is proposed:

ALGORITHM 1. RM-US[ $\lambda$ ]

**(heavy task rule)** If  $\frac{c_i}{T_i} > \lambda$  then schedule  $\tau_i$ 's jobs at maximum priority.

**(light task rule)** If  $\frac{c_i}{T_i} \leq \lambda$  then schedule  $\tau_i$ 's jobs according to their normal rate monotonic priorities.

They then proved that Algorithm RM-US[ $m/(3m-2)$ ] correctly schedules on  $m$  processors any periodic task system whose total utilization is at most  $m^2/(3m-2)$ . The proof is based on the observation that the upper bound on total utilization guarantees the number of heavy tasks cannot exceed  $m$ . The essence of the argument is that Algorithm RM-US[ $m/(3m-2)$ ] can do no worse than scheduling each of the heavy tasks on its own processor, and then scheduling the remainder (which must be light on the remaining processors) using RM.

Corollary 14 extends the analysis of RM-US[ $\lambda$ ] to other values of  $\lambda$ . Let  $\xi$  be the number of tasks with utilization greater than or equal to  $\lambda$ . If  $\xi < m$ , the algorithm can do no worse than scheduling the  $\xi$  highest utilization tasks on dedicated processors and scheduling the  $N - \xi$  lowest-utilization tasks

on the remaining  $m - \xi$  processors. The corollary guarantees that the latter tasks can be scheduled unless their combined utilization exceeds  $\frac{m-\xi}{2}(1 - \lambda)$ . That is, no deadline can be missed unless the total utilization, including both heavy and light tasks, exceeds

$$\frac{m - \xi}{2}(1 - \lambda) + \xi\lambda$$

If no restrictions are placed on the task set, the optimum value of  $\lambda$  for use with RM-US and the utilization bound test is the one that maximizes with respect to  $\lambda$  the minimum with respect to  $\xi$  of the above expression. Since the value of the expression is decreasing with respect to  $\xi$  for  $\lambda < 1/3$  and is increasing with respect to  $\xi$  for  $\lambda > 1/3$ , the minimum is maximized when  $\lambda = 1/3$ , in which case the worst-case guaranteed feasible utilization is  $\frac{m}{3}$ . Of course, the value  $\lambda = 1/3$  is optimal only with respect to verifying schedulability using the utilization bound test, which is probably not tight.

It is argued in (Lundberg, 2002) that the true optimum value of  $\lambda$  for RM-US[ $\lambda$ ] is approximately 0.3748225282. The basis of the argument is an assertion that a task experiences the maximum competing load when there is ‘block interference similar to that shown in Figure 1, and argues that the worst-case block interference occurs with task sets similar to the family of examples in the proof of Theorem 15, whose limiting utilization is  $\lambda + m \ln(\frac{2}{1+\lambda})$ . If Lundberg’s reasoning can be made rigorous, it would follow that  $u_{\max} + m \ln(\frac{2}{1+u_{\max}})$  is the actual worst-case RM utilization bound.

The schedulability tests presented in Section 5 also suggest the following generalization of the RM-US idea:

#### ALGORITHM 2. DM-Hybrid

Choose  $N - \xi$  (where  $0 \leq \xi < m$ ) tasks with the lowest value of  $\frac{c_i}{\min\{d_i, T_i\}}$  and call those the *light* tasks; call the rest of the tasks the *heavy* tasks. Schedule the  $\xi$  heavy tasks maximum priority, and the other tasks according to deadline-monotonic priorities.

The number  $\xi$  of heavy tasks is chosen so that the remaining  $N - \xi$  tasks set can be verified as schedulable on the remaining  $m - \xi$  processors, according to whatever schedulability test is available. For example, if one uses the  $O(N^3)$  test, one would choose the smallest  $\xi$  such that the  $N - \xi$  lowest-utilization tasks pass the  $O(N^3)$  test for schedulability on  $m - \xi$  processors. If such a  $\xi$  exists, it follows that the task set is schedulable.



## 9. Conclusions

Efficiently computable schedulability tests have been given for general fixed-priority scheduling on a homogeneous multiprocessor system, with arbitrary deadlines. These improve on previously known multiprocessor RM schedulability conditions by relaxing the assumptions of rate monotonic priorities and deadline being equal to period.

For the case where period equals deadline this analysis gives a simple lower bound on the minimum achievable utilization. That is, a system of independent periodic or sporadic tasks can be scheduled by RM scheduling to meet all deadlines if the total utilization is at most  $\frac{m}{2}(1 - u_{\max}) + u_{\min}$ , where  $u_{\max}$  is the maximum of the individual task utilizations and  $u_{\min}$  is the minimum. This result can be used to verify the RM schedulability of systems of tasks with sufficiently low individual processor utilization, or combined with a hybrid scheduling policy to verify the schedulability of systems with a few high-utilization tasks. It can be applied statically, or applied dynamically as an admission test. This improves on previously known utilization-based multiprocessor RM schedulability tests, by allowing both higher total utilizations and higher individual task utilizations. In addition to the lower bound on the minimum achievable RM utilization, an upper bound of  $u_{\max} + m \ln\left(\frac{2}{1+u_{\max}}\right)$  has been derived. However, the exact minimum achievable RM utilization remains an open question.

The existence of these schedulability tests makes verification of single-queue (global) fixed-priority multi-processor schedulability an option for certain classes of task sets. This may be of immediate interest for real-time applications on symmetric multiprocessing operating systems that support global scheduling as the default. Removing one of the reasons often given for favoring a queue-per-processor (partitioned) approach to multiprocessor scheduling also opens the way for additional research.

Further study is needed into the comparative strengths of global versus partitioned scheduling. One question is about the average-case performance. The worst-case utilization bounds seem very close, but is that also true of the average case? A second question is about the implementation overhead. Global scheduling has higher overhead in at least two respects: the contention delay and the synchronization overhead for a single dispatching queue is higher than for per-processor queues; the cost of resuming a task may be higher if it is on a different processor (due to interprocessor interrupt handling and cache reloading) than on the processor where it last executed. The latter cost can be quite variable, since it depends on the actual portion of a task's memory that remains in cache when the task resumes execution, and how much of that remnant will be referenced again before it is overwritten. It seems that

only experimentation with actual implementations can determine how serious are these overheads, and how they balance against any advantages global scheduling may have for on-time completion of tasks.

### Acknowledgement

The author is thankful to the anonymous reviewers for their constructive comments, which improved the quality of this paper. He is especially thankful to the reviewer who observed that most of the results in the paper apply to arbitrary fixed priority assignments, and the reviewer who suggested adding some simulation results.

### References

- Andersson, B., S. Baruah, and J. Jonsson: 2001, 'Static-priority scheduling on multiprocessors'. In: *Proc. 22nd IEEE Real-Time Systems Symposium*. London, UK, pp. 193–202.
- Baker, T. P.: 2003a, 'An Analysis of deadline-monotonic scheduling on a multiprocessor'. Technical Report TR-030301, Florida State University Department of Computer Science, Tallahassee, FL.
- Baker, T. P.: 2003b, 'Multiprocessor EDF and deadline monotonic schedulability analysis'. In: *Proc. 24th IEEE Real-Time Systems Symposium*. pp. 120–129.
- Baruah, S. and J. Goossens: 2003, 'Rate-monotonic scheduling on uniform multiprocessors'. *IEEE Trans. Computers* **52**(7), 966–970.
- Dhall, S. K. and C. L. Liu: 1978, 'On a real-time scheduling problem'. *Operations Research* **26**(1), 127–140.
- Liu, C. L. and J. W. Layland: 1973, 'Scheduling algorithms for multiprogramming in a hard real-time environment'. *Journal of the ACM* **20**(1), 46–61.
- Lopez, J. M., J. L. Diaz, and D. F. Garcia: 2001, 'Minimum and maximum utilization bounds for multiprocessor RM scheduling'. In: *Proc. 13th Euromicro Conf. Real-Time Systems*. Delft, Netherlands, pp. 67–75.
- Lopez, J. M., J. L. Diaz, M. Garcia, and D. F. Garcia: 2000, 'Worst-Case Utilization Bound for EDF Scheduling on Real-Time Multiprocessor Systems'. In: *Proc. 12th Euromicro Conf. Real-Time Systems*. pp. 25–33.
- Lundberg, L.: 2002, 'Analyzing Fixed-Priority Global Multiprocessor Scheduling'. In: *Proc. 8th IEEE Real-Time and Embedded Technology and Applications Symposium*. San Jose, CA, USA, pp. 145–153, IEEE Computer Society.
- Oh, D. I. and T. P. Baker: 1998, 'Utilization Bounds for  $N$ -Processor Rate Monotone Scheduling with Stable Processor Assignment'. *Real Time Systems* **15**(2), 183–193.
- Phillips, C. A., C. Stein, E. Torng, and J. Wein: 1997, 'Optimal time-critical scheduling via resource augmentation'. In: *Proc. 29th Annual ACM Symposium on Theory of Computing*. El Paso, TX, pp. 140–149, ACM.
- Sha, L., T. Abdelzaher, K. E. Årzén, A. Cervin, T. P. Baker, A. Burns, G. Buttazzo, M. Caccamo, J. Lehoczky, and A. K. Mok: 2004, 'Real Time Scheduling Theory: A Historical Perspective'. *Real-Time Systems* **28**(2–3), 101–155.