

New Strategies for Assigning Real-Time Tasks to Multiprocessor Systems

Almut Burchard, Jörg Liebeherr, *Member, IEEE*, Yingfeng Oh, and Sang H. Son, *Member, IEEE*

Abstract—Optimal scheduling of real-time tasks on multiprocessor systems is known to be computationally intractable for large task sets. Any practical scheduling algorithm for assigning real-time tasks to a multiprocessor system presents a trade-off between its computational complexity and its performance. In this study, new schedulability conditions are presented for homogeneous multiprocessor systems where individual processors execute the rate-monotonic scheduling algorithm. The conditions are used to develop new strategies for assigning real-time tasks to processors. The performance of the new strategies is shown to be significantly better than suggested by the existing literature. Under the realistic assumption that the load of each real-time task is small compared to the processing speed of each processor, it is shown that the processors can be almost fully utilized.

Index Terms—Hard real-time systems, multiprocessor systems, rate-monotonic scheduling, periodic tasks, task assignment scheme.

I. INTRODUCTION

THE distinguishing feature of real-time computer systems is their attempt to achieve both logical and temporal correctness of computations. A computation is temporally correct if it finishes within a specified time frame. In this sense, all time-constrained computer applications require a real-time computer system. Commonly, however, real-time computer systems are used if violations of temporal correctness may result in drastic consequences as, for example, in power plants, hospitals, or manufacturing and transportation systems.

In most real-time applications, the computer system is subject to arrivals of messages containing monitor and control information from many different sources. These messages can arrive at any time, however, messages arrive at most periodically. Each message arrival initiates a request for executing a computational task. The task must be completed before the arrival of the subsequent message from the same source. Thus, the earliest arrival time of the next message from the same source is the deadline for executing the task. We refer to tasks that are requested at most periodically and must finish execution before the end of the next period as (*periodic*) *real-time* tasks.

To maximize the number of real-time tasks that can be processed without timing violations, real-time computer systems use sophisticated scheduling algorithms to decide the order in which tasks are executed. A schedule for assigning

tasks to one or more processors is said to be *feasible* if the execution of each task can be completed before its deadline. Given a scheduling algorithm, a task set is said to be *schedulable* if the scheduling algorithm generates a feasible schedule. A feasible schedule is said to be *minimal* if there is no feasible schedule utilizing less processors. A scheduling algorithm is said to be *optimal* if for any set of tasks the algorithm finds a minimal schedule.

Scheduling algorithms can be divided into *fixed priority* and *dynamic priority* algorithms. In fixed priority algorithms, the priority of a task remains constant at all times, whereas in dynamic priority algorithms, the priority of a task may change during its execution. In their seminal work, Liu and Layland [9] showed that in single processor systems the dynamic priority *earliest-deadline-due* (EDD) algorithm, which assigns the highest priority to the task closest to the end of its period, is optimal among all scheduling algorithms. They also showed that the *rate-monotonic* (RM) algorithm, which assigns higher priorities to tasks with shorter periods, is optimal among all fixed priority scheduling algorithms. Liu and Layland derived conditions under which the respective algorithms yield feasible schedules. Such conditions are referred to as *schedulability conditions*. For the EDD algorithm, Liu and Layland presented necessary and sufficient schedulability conditions, however, for the RM algorithms, they gave only sufficient conditions. The necessary and sufficient schedulability conditions for the RM algorithm were recently presented by Lehoczky, et al. [7].

Due to its low computational overhead the RM algorithm is widely regarded as an appropriate algorithm for scheduling real-time tasks on uniprocessor systems. Recently proposed extensions to the RM algorithm have increased its practical relevance [11], [12], [13].

Even though real-time computer systems are expected to greatly benefit from multiprocessor technology, employing multiprocessor systems for real-time applications has shown to be difficult. A major obstacle is that scheduling algorithms are significantly more complex for multiprocessor systems than for uniprocessor systems, since the scheduling algorithm must not only specify an ordering of tasks, but must also determine the specific processor to be used. Leung and Whitehead [8] proved that finding a minimal schedule for a given set of real-time tasks in a multiprocessor system is NP-hard. Therefore, research efforts have focused on the development of suitable heuristic algorithms which can be efficiently implemented, yet, require only a limited number of additional processors as compared to an optimal algorithm.

There are two strategies for scheduling real-time tasks on a multiprocessor system. In a *global* scheme each occurrence of

Manuscript received Mar. 12, 1994.

A. Burchard is with the Department of Mathematics, Princeton University, Princeton, NJ 08544.

J. Liebeherr, Y. Oh, and S.H. Son are with the Department of Computer Science, University of Virginia, Charlottesville, VA 22903; e-mail: jorg@virginia.edu.

To order reprints of this article, e-mail: transactions@computer.org, and reference IEEECS Log Number C95146.

TABLE I
ASYMPTOTIC UPPER BOUNDS OF TASK ASSIGNMENT SCHEMES

Scheme	Existing Schemes				Proposed Schemes	
	RMNF [5]	RMFF [5]	FFDUF [2]	RMBF [10]	RMST	RMGT
$\lim_{N_{opt} \rightarrow \infty} N/N_{opt}$	2.67	2.33	2	2.33	$1/(1-\alpha)$	1.75

a real-time task may be executed on a different processor. In contrast, a *partitioning* scheme enforces that all occurrences of a particular task are executed on the same processor. Partitioning schemes have several advantages over global schemes. First, partitioning schemes are less complex since the overhead of multiprocessor scheduling merely consists in assigning tasks to processors. Note that this operation is performed only once for each task, that is, before the task is executed for the first time. Second, if the assignment of tasks to processors is completed, well-known uniprocessor scheduling algorithms can be used for each processor.

The performance of a partitioning scheme is determined by two factors; the *task assignment algorithm* which distributes tasks to the processors, and the scheduling algorithm which determines the order of task executions on each processor. For a given scheduling algorithm, an *optimal* task assignment algorithm achieves a feasible schedule for each processor with the least number of processors. However, the problem of finding an optimal assignment for fixed priority scheduling algorithms, in particular the RM algorithm, as well as for dynamic priority scheduling algorithms, in particular EDD, was shown to be NP-hard [8].

In this study, we are concerned with task assignment schemes for homogeneous multiprocessor systems where each processor executes the RM scheduling algorithm. This problem has been addressed in a number of studies¹ [2], [3], [5], [10]. Typically, the task assignment schemes apply variants of well-known heuristic bin-packing algorithms where the set of processors is regarded as a set of bins.² The decision whether a processor is full is determined by a schedulability condition. All existing task assignment schemes are based on the sufficient schedulability conditions for uniprocessor systems derived in [9] and variants of this condition [4]. Thus, the existing assignment schemes differ mainly in the choice of the bin-packing heuristic.

In [5], two heuristic assignment schemes are proposed, referred to as rate-monotonic next-fit (RMNF) and rate-monotonic first-fit (RMFF). The schemes are based on the *next-fit* and *first-fit* bin-packing heuristic, respectively. In both schemes, tasks are sorted in decreasing order of their periods before the assignment is started. Tasks are assigned to a so-called *current* processor until the schedulability condition is violated, in which case the current processor is marked full and a new processor is selected. RMFF first tries to accommodate a task in a processor marked as full before assigning it to the current processor. The first-fit decreasing-utilization factor (FFDUF)

method is a variation of the first-fit heuristic scheme. Here, tasks are sorted in the order of their load factor [2]. In [10], a *best-fit* bin-packing heuristic is used as the basis for the rate-monotonic best-fit (RMBF) scheme. Similarly to RMFF, RMBF attempts to assign tasks to processors that have been marked as full. However, in RMBF, the full processors are inspected in a specific order. As in [5], tasks are assumed to be sorted by their period.

In all studies, the performance of task assignment schemes is evaluated by providing worst case bounds for N/N_{opt} , where N is the number of processors required to schedule a task set with a given heuristic method, and N_{opt} is the number of processors needed by an optimal assignment. Unfortunately, bounds for the existing schemes are only available as asymptotic bounds, that is, as $\limsup_{N_{opt} \rightarrow \infty} N/N_{opt}$.³ In the left part of Table I, we summarize the asymptotic performance bounds of the discussed task assignment schemes. The computational complexity of all assignment schemes is bounded by $O(K \log K)$ for a set of K real-time tasks.

Our approach for developing task assignment schemes for multiprocessor systems is different from previous work. Rather than increasing the level of sophistication of the bin-packing heuristic, we focus on developing tighter schedulability conditions that allow us to assign more tasks to each processor. We show that the maximum achievable load on each processor is significantly higher than suggested by previous work. If the load factor of each task is small compared to the processing power of a processor—a very realistic assumption considering the state-of-the-art of hardware technology—we will show that each processor can be almost fully utilized. More precisely, the *rate-monotonic small-tasks (RMST) scheme* proposed in this study achieves an asymptotic bound of $\lim_{N_{opt} \rightarrow \infty} N/N_{opt} \leq 1/(1-\alpha)$,

where α is the maximal load factor of an individual task. For general task sets we propose the *rate-monotonic general-tasks (RMGT) scheme* which yields an asymptotic bound of $\lim_{N_{opt} \rightarrow \infty} N/N_{opt} \leq 1.75$. Different from previous work we also

derive bounds of the performance parameter N/N_{opt} for $N < \infty$. The computational complexity of both assignment schemes for assigning a set of K real-time tasks is bounded by $O(K \log K)$. We use simulations to study the average-case behavior of our assignment schemes. The simulation results indicate that on the average, the performance of the new schemes is significantly better than in the worst case. The simulations also show that the average performance of our assignment schemes is superior to any existing scheme.

1. We only consider so-called *off-line schemes* which assume that the task set is known a priori. In [1], we also investigate so-called *on-line schemes* which allow the task set to change dynamically.

2. The bin-packing problem is concerned with packing different-sized items into fixed-sized bins, using the least number of bins [6].

3. "lim su" refers to the smallest asymptotic upper bound, and "lim inf" refers to the largest asymptotic lower bound.

The remainder of this study is structured as follows. In Section II we present our model for real-time tasks and multiprocessor systems. In Section III we derive a tight schedulability condition for the RM scheduling algorithm in a uniprocessor system that improves on the results presented in [9]. We also prove a scheduling result for multiprocessor systems which can be interpreted as *dual* result to the uniprocessor schedulability condition in [9]. In Section IV we construct two simple assignment schemes, referred to as *RMST* and *RMGT*. With our theoretical results from Section III we can prove bounds for the number of processors required with these schemes. In Section V we present simulations to illustrate the average case performance of our schemes. In Section VI we conclude our study.

II. MODEL DESCRIPTION

We consider a real-time computer system consisting of a homogeneous multiprocessor system and a set of K real-time tasks. The multiprocessor and the task set are characterized as follows.

- A real-time task is denoted by $\tau_i = (C_i, T_i)$ ($i = 1, \dots, K$). T_i denotes the shortest time between two requests of task τ_i , and is also referred to as the *period* of τ_i . C_i denotes the maximum execution time of task τ_i . Since we assume that the multiprocessor system is homogeneous the execution time is identical on each processor. Each real-time task must complete execution before the next request of the same task. Thus, in the worst case, the execution of τ_i must be completed before T_i time units.
- The period and the maximum execution time of task τ_i satisfy $T_i > 0$ and $0 \leq C_i \leq T_i$ ($i = 1, \dots, k$). We will refer to $U_i = C_i/T_i$ as the *load factor* of the i th task, and to $U = \sum_{i=1}^K U_i$ as the *total load* of the task set. We use ρ_n to denote the *utilization* of the n th processor, that is, the sum of the load factors of the tasks assigned to processor n .
- Each of the K tasks is assigned to a processor using a *partitioning scheme*, that is, each execution of the same task is performed on the same processor.
- Throughout this paper, we assume that the *rate-monotonic* (RM) algorithm is used to schedule tasks on each processor. That is, task τ_i has precedence over task τ_j , if $T_i < T_j$. We assume that scheduling of tasks is preemptive, and that task execution can be resumed without loss after interruptions.

III. SCHEDULABILITY CONDITIONS

In this section we derive two sufficient schedulability conditions for processors which schedule tasks with the RM algorithm. The first result, presented in Theorem 1, is a simple modification of the result for uniprocessor systems by Liu and Layland [9]. Our result yields a higher utilization of the processor if the task periods satisfy certain constraints. For uniprocessor systems, the increase in utilization provided by Theorem 1 is limited. However, for multiprocessor scheduling we can divide a task set into subsets and use the sharpened condition of Theorem 1 on all but possibly one processor.

In our second result, stated in Theorem 2, we present a schedulability condition for the RM algorithm in multiprocessor systems. In fact, Theorem 2 can be interpreted as a dual result to the schedulability conditions for uniprocessors given in [9]. Both results coincide for the special case $K = 2$.

A result similar to our Theorem 2 was conjectured in [5], but not proven. A partial proof, yet incomplete, and needing additional assumptions, was given in [4].

A. Rate-Monotonic Scheduling in Uniprocessor Systems

The schedulability condition presented in the following theorem takes advantage of a special property of the RM scheduling algorithm. We show that we can increase the processor utilization if all periods in a task set have values that are close to each other.

THEOREM 1. *Given a real-time task set τ_1, \dots, τ_K . Define*

$$S_i := \log_2 T_i - \lfloor \log_2 T_i \rfloor \quad i = 1, \dots, K \quad (1)$$

and

$$\beta := \max_{1 \leq i \leq K} S_i - \min_{1 \leq i \leq K} S_i \quad (2)$$

(a) *If $\beta < 1 - 1/K$, and the total load satisfies*

$$U \leq (K-1) \left(2^{\beta/(K-1)} - 1 \right) + 2^{1-\beta} - 1 \quad (3)$$

then the task set is schedulable on one processor with the RM algorithm.

(b) *If $\beta \geq 1 - 1/K$, and the total load satisfies*

$$U \leq K \left(2^{1/K} - 1 \right) \quad (4)$$

then the task set is schedulable on one processor with the RM algorithm.

Both conditions are tight.

Note that inequality (4) is identical to the schedulability condition given by Liu and Layland [9]. Theorem 1 improves upon [9] when $\beta < 1 - 1/K$, since the strict convexity of the function $f(x) = x(2^{1/x} - 1)$ implies that

$$(K-1) \left(2^{\beta/(K-1)} - 1 \right) + 2^{1-\beta} - 1 > K \left(2^{1/K} - 1 \right). \quad (5)$$

Throughout the paper, we will use a simpler version of the condition in (3), given in the following corollary.

COROLLARY 1. *Given a set of real-time tasks τ_1, \dots, τ_K , and define β as in (2). If the total load satisfies*

$$U \leq \max \{ \ln 2, 1 - \beta \ln 2 \} \quad (6)$$

then the task set can be scheduled on one processor.

PROOF. Since both schedulability conditions (3) and (4) of Theorem 1 are strictly decreasing with respect to K , we have that

$$K \left(2^{1/K} - 1 \right) > \lim_{K \rightarrow \infty} K \left(2^{1/K} - 1 \right) = \ln 2 \quad (7)$$

and

$$(K-1) \left(2^{\beta/(K-1)} - 1 \right) + 2^{1-\beta} - 1 > \lim_{K \rightarrow \infty} (K-1) \left(2^{\beta/(K-1)} - 1 \right) + 2^{1-\beta} - 1 \quad (8)$$

$$= \beta \ln 2 + 2^{1-\beta} - 1 \quad (9)$$

$$> 1 - \beta \ln 2 \quad (10)$$

Schedulability now follows from Theorem 1. \square

The remainder of this subsection contains the proof of Theorem 1, and requires three lemmas. The first lemma, Lemma 1, due to Lehoczky, et al. [7], gives the necessary and sufficient schedulability for the RM algorithm in a uniprocessor system.

LEMMA 1. *Given a set of real-time tasks $\tau_1, \tau_2, \dots, \tau_K$. Assume the tasks are ordered with increasing period, $T_1 \leq \dots \leq T_K$. Then, a task τ_k always meets its deadline T_k under rate-monotonic scheduling, if and only if there exists a time $0 < t \leq T_k$ such that*

$$t \geq \sum_{i=1}^k \left\lceil \frac{t}{T_i} \right\rceil C_i \quad (11)$$

We will need the following special cases of Lemma 1. If $T_K \leq 2T_1$, then condition (11) reduces to

$$\exists j \leq k: T_j \geq \sum_{i=1}^{j-1} 2C_i + \sum_{i=j}^k C_i. \quad (12)$$

If the task set consists of only two tasks, condition (11) reduces to

$$\left\lceil \frac{T_2}{T_1} \right\rceil (T_1 - C_1) \geq C_2 \quad \text{or} \quad T_2 \geq \left\lceil \frac{T_2}{T_1} \right\rceil C_1 + C_2 \quad (13)$$

The next lemma states that the RM algorithm is distinguished by a special property, which also holds for EDD, however, not for general fixed-priority or dynamic scheduling algorithms. The proof is a simple application of Lemma 1, but the result is surprisingly powerful. Lemma 2 permits a generalization of schedulability conditions, which apply to task sets where the longest period is at most twice as long as the shortest period, to unrestricted task sets. Also, Lemma 2 implies that it is not necessary to assume that a task set is ordered by periods in order to apply the schedulability conditions in [5], [10], a fact overlooked in both references.

LEMMA 2. *Given a task set $\tau_1, \tau_2, \dots, \tau_K$, and a task $\tau = (C, T)$ with $T \leq T_i$ for $i = 1, \dots, K$. If τ and τ_1, \dots, τ_K cannot be scheduled together on one processor with the RM scheduling algorithm, then also $(2C, 2T), \tau_1, \dots, \tau_K$ cannot be scheduled.*

PROOF. We assume that the tasks are ordered such that $T_1 \leq \dots \leq T_K$. Denote by k the smallest index, such that (C, T) together with $(C_1, T_1), \dots, (C_k, T_k)$ is not schedulable.

There are two possible cases. If $2T > T_k$, then we can use schedulability condition (12) for both task sets in question. But clearly, the condition

$$2C + \sum_{i=1}^{j-1} 2C_i + \sum_{i=j}^k C_i > T_j \quad \text{for all } j \leq k \quad (14)$$

is equivalent to

$$\sum_{i=1}^{j-1} 2C_i + \sum_{i=j}^k C_i + 2C > T_j \quad \text{for all } 1 \leq j \leq K. \quad (15)$$

Also

$$C + \sum_{i=1}^k C_i > T \quad (16)$$

is equivalent to

$$\sum_{i=1}^k 2C_i + 2C > 2T \quad (17)$$

which shows the claim that $(2C, 2T), (C_1, T_1), \dots, (C_k, T_k)$ cannot be scheduled. If, on the other hand, $2T \leq T_k$, then Lemma 1 implies that for all $0 \leq t \leq T_k$

$$t < \left\lceil \frac{t}{T} \right\rceil C + \sum_{i=1}^k \left\lceil \frac{t}{T_i} \right\rceil C_i \quad (18)$$

$$\leq \left\lceil \frac{t}{2T} \right\rceil 2C + \sum_{i=1}^k \left\lceil \frac{t}{T_i} \right\rceil C_i \quad (19)$$

Again we have shown that $(2C, 2T), (C_1, T_1), \dots, (C_k, T_k)$ is not schedulable. \square

Lemma 3 is a corollary of Lemma 2 which applies to multiprocessor systems.

LEMMA 3. *Assume that the task set τ_1, \dots, τ_K cannot be scheduled on N processors. Then the task set τ'_1, \dots, τ'_K given by*

$$C'_i = U_i T'_i, \quad T'_i = 2^{S_i} \quad (20)$$

cannot be scheduled on N processors.

PROOF. For all tasks τ_i with $T_i > 2^{S_i}$ we replace (C_i, T_i) by $(2^{-m}C_i, 2^{-m}T_i)$, where m is selected such that $2^{-m}T_i < 2^{S_i}$. Note that by scaling t in the schedulability conditions of Lemma 1, the replacement does not change schedulability of the task set. Also, the values for U_i and S_i remain unchanged. So, we may assume that $T_i \leq 2^{S_i}$ for all i . If $T_i = 2^{S_i}$ for all i , we are done. Otherwise, we select τ_k such that $T_k = \min_i(T_i)$ and replace (C_k, T_k) by $(2C_k, 2T_k)$. Clearly, this does not change the load factor U_k . Lemma 2 implies that the resulting task set cannot be scheduled on N processors. We repeat this procedure until we arrive at a task set with $T_i = 2^{S_i}$ for all tasks. \square

PROOF OF THEOREM 1. We will show that any set of K tasks that cannot be scheduled on a single processor violates condition (3), if $\beta < 1 - 1/K$, and violates condition (4), if $\beta \geq 1 - 1/K$. To show that the bounds from Theorem 1 are tight we will construct a task set that cannot be scheduled on one processor, but whose total load is arbitrarily close to the bounds in (3) or (4). The proof will proceed in four steps.

1) Formulate Theorem 1 in terms of a problem of minimizing U as a function of its variables $\underline{C} = (C_1, C_2, \dots, C_K)$ and $\underline{T} = (T_1, T_2, \dots, T_K)$.

- 2) Fix the periods $\underline{T} = (T_1, T_2, \dots, T_K)$ and minimize U over the execution times $\underline{C} = (C_1, C_2, \dots, C_K)$. Use the result to express the execution times as functions of the periods.
 - 3) Transform the reduced minimization problem into a convex problem.
 - 4) Solve the convex minimization problem.
- 1) Assume that the task set $(C_1, T_1), \dots, (C_K, T_K)$ cannot be scheduled on one processor. Since the conditions given by (3) and (4) are strictly decreasing with K , they are certainly violated for a task set if they are violated for a subset. Hence, we may assume without loss of generality that all proper subsets of the task set *can* be scheduled on one processor. By Lemma 3, we can assume that

$$T_1 \leq \dots \leq T_K \leq 2^\beta T_1 \quad (21)$$

Since, by assumption, the proper subset $\tau_1, \dots, \tau_{K-1}$ can, but the complete task set τ_1, \dots, τ_K cannot be scheduled on a single processor, task τ_K , which has the lowest priority, misses its deadline. By the schedulability condition in (12), this is equivalent to

$$\sum_{i=1}^{j-1} 2C_i + \sum_{i=j}^K C_i > T_j \quad j = 1, \dots, K \quad (22)$$

We will minimize the total load U as a function of the execution times and the periods of all tasks under the conditions in (21) and (22). Thus, we have to solve the following problem.

$$\text{minimize } U(\underline{C}, \underline{T}) = \sum_{i=1}^K \frac{C_i}{T_i} \quad (23)$$

subject to

$$\sum_{i=1}^{j-1} 2C_i + \sum_{i=j}^K C_i \geq T_j \quad j = 1, \dots, K \quad (24)$$

$$0 \leq C_i \leq T_i \quad i = 1, \dots, K \quad (25)$$

$$T_1 \leq \dots \leq T_K \leq 2^\beta T_1 \quad \beta \leq 1 \quad (26)$$

We replaced “>” with “≥” in (24) and (25) to ensure that the minimum is attained at some point. Since the functional is continuous, this does not affect the minimal value of U . Note that $U(\underline{C}, \underline{T})$ in (23) is not a convex function of its arguments. Hence, the problem cannot be solved straightforwardly with standard (nonlinear) optimization methods.

- 2) We will show that $U(\underline{C}, \underline{T})$ takes its minimal value in a point where the conditions in (24) hold with equality. Suppose that a minimum, say U^* , for the objective function is assumed at $\underline{C}^* = (C_1^*, C_2^*, \dots, C_K^*)$, and $\underline{T}^* = (T_1^*, T_2^*, \dots, T_K^*)$. If for some $j > 1$, inequality (24) is strict, we set

$$\tilde{C}_i = \begin{cases} C_{j-1}^* - \varepsilon & \text{if } i = j-1 \\ C_j^* + \varepsilon & \text{if } i = j \\ C_i^* & \text{otherwise} \end{cases} \quad (27)$$

where ε is defined by

$$\varepsilon := \sum_{i=1}^{j-1} 2C_i^* + \sum_{i=j}^K C_i^* - T_j > 0 \quad (28)$$

Then side condition (24) is unchanged for $i \neq j$, and holds with equality for $i = j$. The total load at this point satisfies

$$U(\tilde{\underline{C}}, \underline{T}^*) = U^* - \varepsilon \left(\frac{1}{T_{j-1}^*} - \frac{1}{T_j^*} \right) \leq U^* \quad (29)$$

where we have used inequality (21). So we found a new minimum of the functional.

Similarly, if there is strict inequality in condition (24) for $j = 1$, we set

$$\tilde{C}_i = \begin{cases} C_1^* + \varepsilon & \text{if } i = 1 \\ C_K^* - 2\varepsilon & \text{if } i = K \\ C_i^* & \text{otherwise} \end{cases} \quad (30)$$

where

$$\varepsilon := \sum_{i=1}^K C_i^* - T_1^* > 0 \quad (31)$$

The total load satisfies

$$U(\tilde{\underline{C}}, \underline{T}^*) = U^* - \varepsilon \left(\frac{2}{T_K^*} - \frac{1}{T_1^*} \right) \leq U^*, \quad (32)$$

where we have used inequality (21). Again we found a new minimum of the functional.

Summarizing, we have shown that U takes its minimum in a point with

$$\sum_{i=1}^{j-1} 2C_i^* + \sum_{i=j}^K C_i^* = T_j^* \quad \text{for } j = 1, \dots, K \quad (33)$$

Subtracting equations in (33) for consecutive indices, and subtracting (33) for $j = 1$ from (33) for $j = K$, we obtain the following identities

$$C_j^* = T_{j+1}^* - T_j^* \quad \text{for } j = 1, \dots, K-1 \quad (34)$$

$$C_K^* = 2T_1^* - T_K^* \quad (35)$$

Note that the side conditions in (25) are satisfied automatically. Thus, we have reduced the problem to

$$\text{minimize } U(\underline{T}) = \sum_{i=1}^{K-1} \frac{T_{i+1} - T_i}{T_i} + \frac{2T_1 - T_K}{T_K} \quad (36)$$

subject to

$$T_1 \leq \dots \leq T_K \leq 2^\beta T_1 \quad (37)$$

- 3) Substituting

$$x_i := \log_2 \frac{T_{i+1}}{T_i} \quad i = 1, \dots, K-1 \quad (38)$$

$$x_K := \log_2 \frac{2T_1}{T_K} \quad (39)$$

we rewrite the minimization problem in (36) and (37) as

$$\text{minimize } U(\underline{x}) = \sum_{i=1}^K (2^{x_i} - 1) \quad (40)$$

$$\text{subject to } x_i \geq 0, \quad i = 1, \dots, K \quad (41)$$

$$x_K \geq 1 - \beta \quad (42)$$

$$\sum_{i=1}^K x_i = 1 \quad (43)$$

Relation (43) is a consequence of the definitions in (38) and in (39).

- 4) The minimization problem in (40)-(43) is a (strictly) convex problem, since the functional is a sum of convex functions and the side conditions describe a convex set. At this point, the minimization can be completed with a Lagrange multiplier method.

More directly, it follows from the strict convexity of the problem that there is a unique critical point at which the total load is minimized. This critical point must be symmetric under any symmetry of the problem. If we disregard condition (42), the problem is completely symmetric under permutation of the indices. Hence, in the minimum, all x_i^* must be equal, and side condition (43) demands that

$$x_i^* = 1/K, \quad \text{for } i = 1, \dots, K \quad (44)$$

Note that for $\beta \geq 1 - 1/K$, condition (42) is satisfied automatically. Reversing the transformations in (34)-(35) and (38)-(39), we obtain the following solution of the original problem in (23)-(26).

$$C_i^* = a 2^{i/K} (2^{1/K} - 1), \quad T_i^* = a 2^{i/K}, \quad \text{for } i = 1, \dots, K \quad (45)$$

where $a > 0$ is any number. On this task set, $U(\underline{C}, \underline{T})$ takes the minimal value

$$U^* = K(2^{1/K} - 1) \quad (46)$$

given on the right hand side of (4). The bound is tight since the task set given in (45) can be scheduled on a single processor. However, if any of the execution times C_i^* is replaced by $\tilde{C}_i > C_i^*$, then the resulting task set cannot be scheduled.

If $\beta \leq 1 - 1/K$, the minimum satisfies

$$x_K^* = 1 - \beta \quad (47)$$

Problem (40)-(43) is symmetric under permutation of x_1, \dots, x_{K-1} . Hence, in the minimum, all values for x_i ($i \neq K$) must be equal. From (43) and (47) we obtain the solution

$$x_i^* = \frac{\beta}{K-1} \quad i = 1, \dots, K-1 \quad (48)$$

Transforming back with (34)-(35) and (38)-(39) gives the final solution

$$C_i^* = a 2^{\beta i/(K-1)} (2^{\beta/(K-1)} - 1) \quad T_i^* = a 2^{\beta i/(K-1)} \quad i = 1, \dots, K-1 \quad (49)$$

$$C_K^* = a 2^{\beta K/(K-1)} (2^{1-\beta} - 1) \quad T_K^* = a 2^{\beta K/(K-1)} \quad (50)$$

where $a > 0$ is any number. For this task set, $U(\underline{C}, \underline{T})$ takes the minimal value

$$U^* = (K-1)(2^{\beta/(K-1)} - 1) + 2^{1-\beta} - 1. \quad (51)$$

The bound in (3) is tight since the task set given in (49) and (50) can be scheduled on one processor; but if any of the execution times C_i^* is replaced by $\tilde{C}_i > C_i^*$, the resulting task set cannot be scheduled. This completes the proof. \square

B. Rate-Monotonic Scheduling in Multiprocessor Systems

In the previous subsection we were concerned with scheduling real-time tasks on a uniprocessor system that employs the RM scheduling algorithm. Theorem 1 addresses the question: When can a set of tasks be scheduled on *one* processor? For the answer we found the worst case task set which can still be scheduled on a single processor. For the corresponding result in multiprocessor systems we have to find the minimal number of processors that is needed to find a feasible assignment of a task set to a multiprocessor system. Obviously, in the worst case, only one task can be assigned to each processor. Therefore, we have to be concerned with the question: When can a set of tasks be scheduled on less than one processor for *each* task?

The answer to this question is given in Theorem 2. In short, Theorem 2 says, that if K processors are needed to schedule a set of K tasks, then the load on each processor cannot be much less than $1/2$.

THEOREM 2. *If the total load of a set of K real-time tasks satisfies*

$$U \leq \frac{K}{2^{1/K} + 1} \quad (52)$$

then the task set can be scheduled with the RM algorithm on less than K processors. The condition is tight.

The following example shows that Theorem 2 is a true multiprocessor result, and cannot be obtained as a corollary of Theorem 1. We select an integer $K \geq 3$, and set

$$C_i = 2^{i/K} \quad T_i = 2^{i/K} (2^{1/K} + 1) \quad \text{for } i = 1, \dots, K \quad (53)$$

Since $U = \frac{K}{2^{1/K} + 1}$ the task set can be scheduled on less than K processors by Theorem 2. However, since $1/(2^{1/K} + 1) > \sqrt{2} - 1$ the schedulability conditions in Theorem 1 fail for any pair of tasks. As a result, assignment schemes based on Theorem 2 can achieve a better processor utilization than schemes based on uniprocessor results alone.

The following corollary provides a good approximation of Theorem 2. The given bound is asymptotically exact for $K \rightarrow \infty$, and never differs by more than 0.07 from the exact value.

COROLLARY 2. *If the total load U of a set of K real-time tasks ($K \geq 2$) satisfies*

$$U \leq \frac{K}{2} - \frac{\ln 2}{4} \quad (54)$$

then the task set can be scheduled with the RM algorithm on less than K processors.

PROOF. The claim follows from the following inequalities ($K \geq 2$):

$$\frac{K}{2} - \frac{\ln 2}{4} < \frac{K}{2^{1/K} + 1} \leq \frac{K}{2} - \frac{1}{6} \quad (55)$$

□

Next we give the proof of Theorem 2. We will need a technical lemma, given in Lemma 4. With Lemma 4, we can show that for the task set which satisfies the inequality of Theorem 2 with the lowest total load, the execution times C_i satisfy the same ordering relations $C_1 \leq \dots \leq C_k \leq 2C_1$ as the periods. A proof of the lemma is given in Appendix A.

LEMMA 4. Assume that a set of K arbitrary nonnegative numbers y_1, \dots, y_K has the following property. For every $1 \leq j \leq K$ there exists at least one index i such that

$$y_j = m_i := \min\{2y_1, \dots, 2y_{i-1}, y_{i+1}, \dots, y_K\} \text{ or } 2y_j = m_i \quad (56)$$

For such a set, either

$$y_1 \leq \dots \leq y_K \leq 2y_1 \quad (57)$$

or there exists an index $1 \leq j \leq K$ and numbers $a > b \geq 0$ such that

$$y_i = \begin{cases} a & \text{if } i < j \\ b & \text{if } i = j \\ 2a & \text{if } i > j \end{cases} \quad (58)$$

PROOF OF THEOREM 2. We want to find the smallest value of U for a set of K tasks that cannot be scheduled on less than K processors. Since for $K = 1$ there is nothing to show, we will assume $K \geq 2$. The proof will proceed in four steps.

- 1) Formulate Theorem 2 in terms of a problem of minimizing U as a function of its variables (C_i, T_i) .
 - 2) Fix the execution times $\underline{C} = (C_1, C_2, \dots, C_K)$ and minimize U over the periods $\underline{T} = (T_1, T_2, \dots, T_K)$. Use the result to express the periods as functions of the execution times.
 - 3) Transform the minimization problem into a convex minimization problem.
 - 4) Solve the minimization problem with standard methods.
- 1) Assume that the tasks $(C_1, T_1), \dots, (C_K, T_K)$ cannot be scheduled on less than K processors. We want to show that the total load U violates (52). With Lemma 3 we can assume that

$$T_1 \leq \dots \leq T_K \leq 2T_1 \quad (59)$$

The assumption that the K tasks cannot be scheduled on less than K processors is equivalent to the statement that no two tasks can be scheduled on one processor. By (13), this is equivalent to

$$\begin{cases} T_i < C_i + C_j & \text{if } i < j \\ T_i < C_i + 2C_j & \text{if } i > j \end{cases} \quad (60)$$

So we have to solve the following minimization problem.

$$\text{minimize } U(\underline{C}, \underline{T}) = \sum_{i=1}^K \frac{C_i}{T_i} \quad (61)$$

subject to

$$T_i < C_i + C_j \quad i < j \quad (62)$$

$$T_i < C_i + 2C_j \quad i > j \quad (63)$$

$$0 \leq C_i \leq T_i \quad i = 1, \dots, K \quad (64)$$

$$T_1 \leq \dots \leq T_k \leq 2T_1 \quad (65)$$

As in the proof of Theorem 1, we replace “<” by “≤” to enforce that the minimum is attained at some point.

2) Since the partial derivatives of $U(\underline{C}, \underline{T})$

$$\frac{\partial U}{\partial T_i}(\underline{C}, \underline{T}) = -\frac{C_i}{T_i^2} < 0 \quad (66)$$

are negative, $U(\underline{C}, \underline{T})$ is minimal if we choose the T_i as large as possible. We enforce the first two side conditions, (62) and (63), by setting

$$T_i = C_i + \min\{2C_1, \dots, 2C_{i-1}, C_{i+1}, \dots, C_k\} \quad (67)$$

With (67), the objective function can be phrased exclusively in terms of \underline{C} . To simplify (67), we will show that in the absolute minimum of $U(\underline{C}, \underline{T})$, conclusion (57) of Lemma 4 holds. To this end, we first have to show that the assumptions of Lemma 4 hold. Assume that (56) is not satisfied, that is, for some index j there exists neither $i < j$ such that $m_i = C_j$, nor $i > j$ such that $m_i = 2C_j$. Since the total derivative

$$\frac{dU}{dC_j}(\underline{C}) = \frac{m_j}{(C_j + m_j)^2} > 0 \quad (68)$$

is positive, we can lower the value of the functional $U(\underline{C})$ by lowering the values for C_j . Thus, condition (56) of Lemma 4 is satisfied in any critical point of the functional.

Next we show that the second outcome of Lemma 4 is not feasible in a minimum of the functional. If the task execution times satisfy (58) then

$$C_i = \begin{cases} a & \text{if } i < j \\ b & \text{if } i = j \\ 2a & \text{if } i > j \end{cases} \quad (69)$$

with $0 \leq b < a$. At such a point the periods of the tasks are given by

$$T_i = \begin{cases} a+b & \text{if } i < j \\ 2a+b & \text{if } i = j \\ 2a+2b & \text{if } i > j \end{cases} \quad (70)$$

We obtain for the total load factor

$$U(\underline{C}, \underline{T}) = (K-1) \frac{a}{a+b} + \frac{b}{2a+b} \quad (71)$$

Setting

$$t = \frac{b}{a} \quad (72)$$

equation (71) reads

$$U(t) = (K-1) \frac{1}{1+t} + \frac{t}{2+t} \quad (73)$$

which is nonincreasing with t . Consequently, for $0 \leq t \leq 1$, U assumes its minimum at $t = 1$, that is $a = b$. Hence, condition (57) holds.

Since we showed that in a minimum the C_i satisfy condition (57), we obtain with (67) that

$$T_i = C_i + C_{i+1} \quad i = 1, \dots, K-1 \quad (74)$$

$$T_K = C_K + 2C_1 \quad (75)$$

Note that the side condition in (65) is always satisfied. Summarizing, we have reduced the problem to

$$\text{minimize } U(\underline{C}) = \sum_{i=1}^{K-1} \frac{C_i}{C_i + C_{i+1}} + \frac{C_K}{C_K + 2C_1} \quad (76)$$

subject to

$$C_i \geq 0 \quad i = 1, \dots, K \quad (77)$$

3) We perform a transformation of variables. Define

$$x_i := \log_2 \frac{C_{i+1}}{C_i} \quad \text{if } i < K \quad (78)$$

$$x_K := \log_2 \frac{2C_1}{C_K}$$

Then the optimization problem of (76) and (77) reads

$$\text{minimize } U(\underline{x}) = \sum_{i=1}^K \frac{1}{1+2^{x_i}} \quad (79)$$

subject to

$$x_i \geq 0 \quad i = 1, \dots, K \quad (80)$$

$$\sum_{i=1}^K x_i = 1 \quad (81)$$

where (81) is a consequence of definition (78).

4) Since the optimization problem in (79)-(81) is strictly convex, and since both the functional and the side conditions are symmetric under permutation of the indices, the unique minimum is also symmetric under permutations of indices. From the side condition in (81), we directly obtain that the solution must be

$$x_i^* = 1/K \quad i = 1, \dots, K \quad (82)$$

Transforming back to the original variables we obtain the solution in terms of execution times and periods:

$$C_i^* = a2^{i/K} \quad (83)$$

$$T_i^* = a2^{i/K}(2^{1/K} + 1) \quad (84)$$

where $a > 0$ can be any number. For this task set, U has the minimal value

$$U^* = K/(2^{1/K} + 1) \quad (85)$$

To show that the bound of Theorem 2 is tight, we verify that the task set with parameters as in (83) and (84) can be scheduled on less than K processors. But if all C_i^* are replaced by $\tilde{C}_i > C_i^*$, then the resulting task can only be scheduled on K processors. This completes the proof. \square

IV. ASSIGNMENT SCHEMES FOR MULTIPROCESSOR SYSTEMS

In this section, we design a new class of assignment schemes for distributing a set of real-time tasks to a set of processors. The new assignment schemes apply our schedulability conditions from Section III. To evaluate the performance of our assignment schemes, we determine the number of processors needed by our schemes for processing a task set. Ideally, one would like to compare the number of processors needed by our heuristic assignment schemes with the minimum number processors N_{opt} that is required by an optimal task assignment scheme. However, since it is not feasible in practice to analytically determine N_{opt} for an arbitrary task set, we will use the total load U as a lower bound for N_{opt} .

We propose two assignment schemes. The first scheme, referred to as *rate-monotonic small-tasks* or *RMST* scheme, is intended for task sets where the load factor U_i of each real-time task is small compared to the processing speed of each processor. The second scheme, referred to as *rate-monotonic general-tasks* or *RMGT* scheme, applies to general task sets.

Previously proposed assignment schemes only consider the load factors of the tasks [2], [3], [5], [10]. Our schemes gain superiority by additionally taking into account the task periods. Before we present the schemes, let us review the maximum performance that can be achieved with a task assignment scheme using information on the load factors only. In this case, the sufficient schedulability conditions given in [9] and its variants [5], [10] are the best available schedulability conditions. With these conditions, the load assigned to any pair of processors is guaranteed to exceed $\ln 2/2$ per processor. So, if N processors are used, then

$$U > \frac{\ln 2}{2} N \quad (86)$$

This bound cannot be improved beyond $(\sqrt{2}-1)N$. For example, a set with K real-time tasks where $U_i = \sqrt{2}-1+\epsilon$ for each task cannot be scheduled onto K processors with this condition. Thus, all task assignment schemes that use the scheduling condition of [9] are strictly limited in the performance they can achieve.

Similar arguments show that if the load factor of every task is less than α , then a next-fit task assignment scheme based on [9] will ensure that the load on all but one processor is at least $\ln 2 - \alpha$. This shows that

$$U > (\ln 2 - \alpha)N \quad (87)$$

For $\alpha < 1/2$ this bound is best possible. More sophisticated bin-packing heuristics, such as first-fit or best-fit, can improve the average performance, but one can construct task sets such that $U \leq N \ln 2$. Consequently, in the worst case, $N/N_{opt} \geq 1/\ln 2$.

A moment's consideration shows, that the bounds given in inequalities in (86) and (87) are far below the bounds that can be achieved with an optimal assignment scheme. By Theorem 2, any set of K real-time tasks with $U_i \leq 1/2 - \epsilon$ for all tasks can be scheduled on less than K processors if the number of tasks is sufficiently large. And we can certainly assign task sets with

TABLE II
EXAMPLE PARAMETER SET

$i =$	1	2	3	4	5	6	7	8	9	10
$T_i =$	65	280	36	150	20	45	400	7	230	60
$C_i =$	16	27	11	31	3	14	113	2	70	19
$U_i =$	0.2461	0.964	0.3055	0.2067	0.15	0.3111	0.2825	0.2857	0.3043	0.3167
$S_i =$	0.0223	0.1292	0.1699	0.2288	0.3219	0.4918	0.6438	0.8073	0.8454	0.9068

small load factors to processors in such a way that the load on all but one processor exceeds $1/2$. This argument suggests that the best bound for an optimal assignment scheme is of the form

$$U > \frac{N}{2} - const. \tag{88}$$

Similarly, if the load factor of every task is bounded above by α , one may hope to prove that

$$U > (1 - \alpha)N - const. \tag{89}$$

Note that these inequalities have the same form as corresponding inequalities for periodic tasks *without* deadline constraints. In particular, the leading terms $N/2$ and $(1 - \alpha)N$ are best possible.

The RMST and RMGT assignment schemes proposed here create task assignments that satisfy bounds of the form (89) and (88), respectively. In the following we discuss the schemes in detail and prove properties of their worst-case performance.

A. RMST—An Assignment Scheme for Small Tasks

We first consider the problem of scheduling a set of tasks with small load factors. Denote by

$$\alpha := \max_{i=1, \dots, K} U_i \tag{90}$$

the maximal load factor of any single task. For all practical purposes, we may assume that a task set contains only “small” tasks if $\alpha \leq 1/2$.

Recall that by Corollary 1 the minimal achievable load on a single processor is larger than $1 - \beta \ln 2$ where β is defined as in (2). The main idea of RMST is to partition the tasks in such a way that on each processor, β has a small value.

It is convenient to visualize the values of S_i as defined in (1) for a given task set as points on a circle with circumference one. Starting at any point on the circle and proceeding clockwise, we assign tasks to processors, using the schedulability condition of Theorem 1. Then, the value of β at each processor is given by the length of the arc spanned by the tasks that are assigned to that processor.

The RMST scheme is summarized in Algorithm 1. It can be easily verified that the computational complexity of Algorithm 1 is determined by the sorting of tasks in Step 1. Thus, the complexity of RMST is $O(K \log K)$ where K is the size of the task set.

- (1) Order the task set such that $0 \leq S_1 \leq \dots \leq S_K < 1$ and set $S_{K+1} := S_1 + 1$. Set the task index to $i := 1$, and the processor index to $n = 0$.
- (2) Select an empty processor with index $n := n + 1$. Assign task τ_i to processor n , that is $\rho_n := U_i$. Set $S := S_i$, and $\tilde{\beta}_n := 0$.
- (3) Increase the task index, $i := i + 1$, and set $\tilde{\beta}_n := S_i - S$. If the schedulability condition from Corollary 2, i.e.,

$$U_i + \rho_n \leq \max\{\ln 2, 1 - \tilde{\beta} \ln 2\}$$
 is satisfied, assign task τ_i to processor n by setting $\rho_n := \rho_n + U_i$, and continue with step (3). Otherwise, continue with step (2).
- (4) When all tasks have been assigned set $\tilde{\beta}_n := S_{K+1} - S$ and terminate.

Algorithm 1. Rate-monotonic small tasks (RMSTs).

We will illustrate the RMST scheme with an example. In Table II we show the parameters for a set of 10 tasks, also including the values of S_i . The task set is already ordered according to S_i . For the parameter set in Table I we obtain $\alpha = 0.3167$. In Fig. 1, we depict the values for S_i on a circle with circumference one. The shaded areas indicate the clockwise assignment of tasks to processors as obtained by RMST. The areas are labeled with the value of β for each processor.⁴ Thus, three processors are required to schedule the set of tasks of Table I, where tasks $\tau_1, \tau_2, \tau_3, \tau_4$ are assigned to the first processor, τ_5, τ_6, τ_7 to the second processor, and $\tau_8, \tau_9, \tau_{10}$ to the third processor. Note that any assignment scheme proposed in the literature would utilize four processors.




Figure 1: Processor Assignment for Algorithm 1.

4. Note that β is slightly different from $\tilde{\beta}_n$ in Algorithm 1; $\tilde{\beta}_n$ also contains the value of S_i of the first task τ_i not assigned to processor n .

Next we will show that the RMST scheme satisfies performance bounds of the form given in (88).

THEOREM 3. *For any given task set τ_1, \dots, τ_K , the RMST scheme arrives at a feasible assignment of tasks to processors. If the maximal load factor of any single task is*

$$\alpha := \max_{i=1, \dots, K} \frac{C_i}{T_i} < 1 \quad (91)$$

then the number N of processors needed by RMST satisfies the inequality

$$N < \frac{U}{1-\alpha} + 2 - \frac{1-\ln 2}{1-\alpha}. \quad (92)$$

The inequality in (92) is not quite tight; however, the inequality is asymptotically tight since one can construct arbitrarily large task sets for which RMST requires at least $U/(1-\alpha)$ processors. For $\alpha \leq 1/2$, the bounds obtained with the RMST scheme improve upon any previously proposed assignment scheme. For $\alpha > 1/2$, the RMGT scheme presented in the next subsection will give better bounds than RMST. Since in most practical real-time systems the load imposed by individual real-time tasks is small compared to the power of the processors, the result of Theorem 3 shows to be highly relevant for practical real-time applications.

PROOF OF THEOREM 3. Denote by K_n the number of tasks assigned to the n th processor, and by ρ_n the load on the processor due to these tasks. Define $\tilde{\beta}_n$ as $\tilde{\beta}_n = S_{i_{n+1}} - S_{i_n}$, where i_n is the index of the first task assigned to processor n . We also set $S_{i_{N+1}} = 1 + S_1$. Note that these definitions are identical to the values of $\tilde{\beta}_n$ after Algorithm 1 has terminated.

The task assignment resulting from Algorithm 1 is feasible by Corollary 1. We now estimate N , the number of processors needed by Algorithm 1 to schedule a given task set. By our schedulability condition in (6), the loads on the processors satisfy

$$\rho_n > 1 - \tilde{\beta}_n \ln 2 - \alpha \quad n = 1, \dots, N-1 \quad (93)$$

$$\rho_{N-1} + \rho_N > 1 - \tilde{\beta}_{N-1} \ln 2 \quad (94)$$

Adding (93) and (94), and using that by construction

$$\sum_{n=1}^N \tilde{\beta}_n = 1 \quad (95)$$

we obtain

$$U = \sum_{n=1}^N \rho_n \quad (96)$$

$$> \sum_{n=1}^{N-2} (1 - \tilde{\beta}_n \ln 2 - \alpha) + 1 - \tilde{\beta}_{N-1} \ln 2 \quad (97)$$

$$\geq (N-2)(1-\alpha) + 1 - \ln 2 \quad (98)$$

Solving for N yields assertion (92). \square

In the following corollaries we compare the number of processors required by RMST to schedule a task set with the

number of processors required by an optimal assignment scheme. Corollary 3 gives the bound of the performance parameter N/N_{opt} for any number N of processors. In Corollary 4 we present the asymptotic limit which is included in Table I.

COROLLARY 3. *For any task set, the number N of processors needed by RMST satisfies*

$$\frac{N}{N_{opt}} < \frac{1}{1-\alpha} + \frac{2}{N_{opt}} \quad (99)$$

N_{opt} denotes the number of processors utilized by an optimal assignment scheme.

PROOF. For any schedule, we must have $N_{opt} \geq U$. With (92) we obtain

$$N < \frac{N_{opt}}{1-\alpha} + 2 - \frac{1-\ln 2}{1-\alpha} \quad (100)$$

Dividing by N_{opt} gives the claim. \square

COROLLARY 4. *Let $\{\tau_i \mid i = 1, 2, \dots\}$ be a given infinite task set. Assume that the load factor of any single task does not exceed α . Denote by $U(k)$ the total load factor of tasks τ_1, \dots, τ_k . Denote by $N(k)$ the number of processors used by RMST to schedule τ_1, \dots, τ_k , and by $N_{opt}(k)$ the number of processors needed in an optimal assignment. If*

$$\lim_{k \rightarrow \infty} U(k) = \infty \quad (101)$$

then

$$\liminf_{k \rightarrow \infty} \frac{U(k)}{N(k)} \geq 1 - \alpha \quad (102)$$

$$\limsup_{k \rightarrow \infty} \frac{N(k)}{N_{opt}(k)} \leq \frac{1}{1-\alpha} \quad (103)$$

The bounds are tight.

PROOF. We obtain (102) by dividing by N in (92) of Theorem 3 and passing to the limit. For (103) we pass to the limit in inequality (99) of Corollary 3. \square

B. RMGT—An Assignment Scheme for General Task Sets

The RMST scheme from the previous subsection provides excellent bounds for task sets where the maximal load factor of any task τ_i is limited by $U_i \leq 1/2$. Next we propose a task assignment scheme, referred to as *rate-monotonic general-tasks (RMGT)* scheme which applies to unrestricted task sets. We show that RMGT is able to find a feasible task assignment with less than $N < 2U + 2$ processors.

For the RMGT scheme we partition the task set into two groups, such that the load factors of tasks in the first and second group, respectively, satisfy $U_i \leq 1/3$ and $U_i > 1/3$. Tasks in the first group are assigned to processors with the RMST scheme in Algorithm 1. Tasks from the second group are assigned to processors with a first-fit heuristic. The heuristic assigns at most two tasks to one processor using the exact schedulability condition from (13). The complete RMGT scheme is summarized in Algorithm 2.

Partitioning the task set in Algorithm 2 involves a computational complexity of $O(K)$ where K is the total number of tasks.

Recall from Subsection IV.A that Algorithm 1 has a complexity of $O(K_1 \log K_1)$, where K_1 is the number of tasks in group \mathcal{G}_1 . Since a first fit bin-packing algorithm for tasks in \mathcal{G}_2 can be implemented with a complexity of $O(K_2 \log K_2)$, where K_2 is the number of tasks in the second group [6], the worst case computational complexity of the RMGT scheme is given by $O(K \log K)$.

In Theorem 4, we show that the number of processors needed by RMGT for scheduling an arbitrary set of real-time tasks satisfies a bound similar to (88). In Theorem 5 and Corollary 5 we prove bounds for the performance of the RMGT scheme.

(1) Partition the set of tasks into two groups:

$$\mathcal{G}_1 = \{\tau_i \mid U_i \leq 1/3\} \quad \mathcal{G}_2 = \{\tau_i \mid U_i > 1/3\}$$

(2) Use the RMST scheme in Algorithm 1 to assign the tasks in set \mathcal{G}_1 .

(3) Assign tasks in \mathcal{G}_2 as follows:

(3.1) Set the task index to $i := 1$ and assign τ_1 to an empty processor with index 1. Set $p_1 := U_1$.

(3.2) Increase the task index to $i := i + 1$ and consider task τ_i . Use a first-fit heuristic to find a processor n that contains only one task, say τ_j , and satisfies the following condition:

$$\begin{cases} \lfloor \frac{p_i}{T_i} \rfloor (T_i - C_i) \geq C_j \text{ or } T_j \geq \lfloor \frac{p_i}{T_i} \rfloor C_i + C_j & \text{if } T_i < T_j \\ \lfloor \frac{p_i}{T_j} \rfloor (T_j - C_j) \geq C_i \text{ or } T_i \geq \lfloor \frac{p_i}{T_j} \rfloor C_j + C_i & \text{if } T_i \geq T_j \end{cases}$$

If such a processor exists, assign task τ_i to processor n by setting $p_n := p_n + U_i$, and mark processor n as full. Otherwise, assign task τ_i to an empty processor with index m , and set $p_m := U_i$. Continue with step (3.2).

(3.3) Terminate when all tasks in \mathcal{G}_2 have been assigned.

Algorithm 2. Rate-monotonic general tasks (RMGTs).

THEOREM 4. For any task set τ_1, \dots, τ_K , the RMGT algorithm arrives at a feasible processor assignment. The number N of processors utilized by RMGT to schedule the task set satisfies

$$N < 2U + \frac{5}{2} \ln 2 - \frac{1}{3} \approx 2U + 1.42 \quad (104)$$

PROOF. Given a task set $\tau_1, \tau_2, \dots, \tau_K$. By Theorem 3 and Lemma 1, RMGT arrives at a feasible processor assignment. To estimate the number of processors used by RMGT, assume that Algorithm 2 assigns to n_1 processors each two tasks from \mathcal{G}_2 , and to n_2 processors each one task from \mathcal{G}_2 . Also assume that Algorithm 1 utilizes n_3 processors for tasks in \mathcal{G}_1 . Denote the total load factors assigned to the three groups of processors by $U^{(1)}$, $U^{(2)}$, and $U^{(3)}$, respectively. Then we obtain:

$$U^{(1)} > \frac{2n_1}{3} \geq \frac{n_1}{2} \quad (105)$$

$$U^{(2)} > \frac{n_2}{2} - \frac{\ln 2}{4} \quad (106)$$

$$U^{(3)} > \frac{2n_3}{3} - \frac{1}{3} - \ln 2 \geq \frac{n_3}{2} - \ln 2 + \frac{1}{6} \quad (107)$$

Inequality (105) holds, because each processor in this group has been assigned two tasks, each with a load larger than $1/3$. Inequality (106) follows from Corollary 2. Inequality (107) follows from Theorem 3 provided that $n_3 \geq 3$. Because the joint load on any two processors exceeds $\ln 2$ by

construction, (107) holds also if $n_3 < 3$. Adding the three inequalities and solving for N yields the claim in (104). \square

THEOREM 5. The number of processors utilized by RMGT to assign a task set satisfies

$$\frac{N}{N_{opt}} \leq \frac{7}{4} + \frac{15 \ln 2 - 4}{8N_{opt}} \approx \frac{7}{4} + \frac{0.79}{N_{opt}} \quad (108)$$

where N_{opt} is the number of processors utilized by an optimal assignment. The inequality becomes false, if the right hand side is replaced by anything less than $7/4$.

PROOF. As in the proof of Theorem 4, we denote by n_2 the number of processors that have been assigned one task from \mathcal{G}_2 . By adding inequalities (105), (106), and (107), we obtain

$$N_{opt} \geq U > \frac{2}{3}(N - n_2) + \frac{n_2}{2} - \frac{1}{3} - \ln 2 - \frac{\ln 2}{4} \quad (109)$$

$$= \frac{2}{3}N - \frac{n_2}{6} - \frac{1}{3} - \frac{5}{4} \ln 2 \quad (110)$$

Moreover, since we used exact schedulability conditions to schedule class \mathcal{G}_2 , we have that

$$N_{opt} \geq n_2 \quad (111)$$

We obtain

$$N_{opt} \geq \min_{0 \leq n \leq N} \max \{U, n_2\} \quad (112)$$

$$\geq \frac{4}{7}N + \frac{2}{7} - \frac{15}{14} \ln 2 \quad (113)$$

which proves (108).

To prove that the second claim of the theorem holds we will construct a class of arbitrarily large task sets such that $N/N_{opt} = 7/4$. The task set consists of $13m$ ($m > 0$) tasks. For the construction of the tasks we select two small positive numbers ε and δ with

$$\varepsilon \ln 2 \leq \delta < \frac{3}{2} \varepsilon \ln 2 \quad (114)$$

We label the tasks by τ_{ij} , where $1 \leq i \leq m$, and $1 \leq j \leq 13$. The tasks are given by

$$C_{ij} = U_{ij} T_{ij}, \quad T_{ij} = 2^{S_{ij}} \quad (115)$$

The values for U_{ij} and S_{ij} are given by

$$U_{ij} = \begin{cases} 1/2 - \delta & \text{if } j = 2, 5, 8, 11 \\ 1/3 & \text{if } j = 1, 6, 10 \\ 1/6 - \delta & \text{if } j = 3, 4, 7, 9, 12, 13 \end{cases} \quad (116)$$

and

$$S_{ij} = \begin{cases} (12i + j)\varepsilon & \text{if } 1 \leq j \leq 12 \\ (12i + 11)\varepsilon & \text{if } j = 13 \end{cases} \quad (117)$$

An optimal assignment will require exactly $4m$ processors to schedule the given task set. In particular, the optimal assignment distributes the following groups of tasks to one processor each.

$$\{\tau_{11}, \tau_{12}, \tau_{13}\}, \{\tau_{14}, \tau_{15}, \tau_{16}\}, \{\tau_{17}, \tau_{18}, \tau_{19}, \tau_{113}\}, \{\tau_{110}, \tau_{111}, \tau_{112}\}$$

$$\text{for } i = 1, \dots, m \quad (118)$$

Since the total load of the task set is given by $U \geq (1 - 4\delta)(4m)$, the given assignment is optimal, if ε and δ are chosen sufficiently small. The assignment is feasible by Corollary 1.

RMGT as given in Algorithm 2 first divides the task set into two groups. In the first group we have for each i

$$\mathcal{G}_1 = \{\tau_{ij} \mid j = 1, 3, 4, 6, 7, 9, 10, 12, 13; i = 1, \dots, m\} \quad (119)$$

These tasks are assigned with Algorithm 1, resulting in the following distribution of task groups to one processor each.

$$\{\tau_{i1}, \tau_{i3}, \tau_{i4}\}, \{\tau_{i6}, \tau_{i7}, \tau_{i9}\}, \{\tau_{i10}, \tau_{i12}, \tau_{i13}\} \text{ for } i = 1, \dots, m \quad (120)$$

The second group is given by

$$\mathcal{G}_2 = \{\tau_{ij} \mid j = 2, 5, 8, 11; i = 1, \dots, m\} \quad (121)$$

The load factor of each task in \mathcal{G}_2 is given by $1/2 - \delta$. Thus, according to (13), no two tasks from \mathcal{G}_2 can be scheduled on the same processor. So, RMGT needs $7m$ processors for the same task set. \square

The following corollary provides the asymptotic limit of the performance parameter N/N_{opt} which is included in Table I.

COROLLARY 5. Let $\{\tau_i \mid i = 1, 2, \dots\}$ be a given infinite task set. Denote by $U(k)$ the sum of the load factors of tasks τ_1, \dots, τ_k . Denote by $N(k)$ the number of processors used by RMGT to schedule τ_1, \dots, τ_k , and by $N_{opt}(k)$ the number needed in an optimal assignment. If

$$\lim_{k \rightarrow \infty} U(k) = \infty \quad (122)$$

then

$$\liminf_{k \rightarrow \infty} \frac{N(k)}{N_{opt}(k)} \geq \frac{1}{2} \quad (123)$$

$$\limsup_{k \rightarrow \infty} \frac{N(k)}{N_{opt}(k)} \leq \frac{7}{4} \quad (124)$$

Both bounds are tight.

PROOF. Inequality (123) is obtained by dividing (104) of Theorem 4 by $2N$, and passing to the limit. Inequality (124) follows by passing to the limit in (108) of Theorem 5. \square

V. AVERAGE-CASE PERFORMANCE EVALUATION OF RMST AND RMGT

So far, the performance bounds of the new assignment schemes, RMST and RMGT, were derived under worst-case assumptions. While a worst-case analysis assures that the performance bounds are satisfied for any task set, it does not provide insight into the average behavior of the assignment schemes. To obtain the average-case performance of the RMST and RMGT schemes, one can analyze the schemes with probabilistic assumptions, or conduct simulation experiments to empirically study the average performance. Since a probabilistic analysis of our algorithms is beyond the scope of this study, we resort to simulation to gain insight into the average-case behavior of RMST and RMGT.

We present simulation experiments for large task sets with $100 \leq K \leq 1,000$ tasks. In each experiment, we vary the value of parameter $\alpha = \max_{i=1, \dots, K} U_i$, the maximal load factor occurring in the task set. The task periods are selected to be uniformly distributed with values $1 \leq T_i \leq 500$. The execution times of the tasks are also taken from a uniform distribution with range $0 < C_i \leq \alpha T_i$. The performance metric in all experiments is the number of processors required to assign a given task set. We compare the RMST and RMGT schemes with several existing assignment schemes. All assignment schemes are executed on identical task sets. The following assignment schemes are considered:

- rate-monotonic small-tasks (RMST) (Section IV.A),
- rate-monotonic general-tasks (RMGT) (Section IV.B),
- rate-monotonic-next-fit (RMNF) [5],
- rate-monotonic-first-fit (RMFF) [5],
- rate-monotonic-best-fit (RMBF) [10].⁵

Since an optimal task assignment cannot be calculated for large task sets, we use the total load ($U = \sum_{i=1}^K U_i$) to obtain a lower bound for the number of processors required.

The outcome of the simulation experiments is shown in Figs. 2-4. The maximum load of a task is set to $\alpha = 0.2$ in Fig. 2, to $\alpha = 0.5$ in Fig. 3, and to $\alpha = 0.8$ in Fig. 4. In the figures, each data point depicts the average value of 15 independently generated task sets with identical parameters. Note that the RMGT scheme gives the best performance in all experiments. In Fig. 2, the performance of RMST and RMGT cannot be distinguished since RMST and RMGT are identical if $\alpha < 0.3$. As we increase the value of α , we observe that the performance of RMST decreases. Note from Figs. 2-4, that in all schemes the rate of increase for the number of processors required is approximately proportional to the total load. In Table III we show the minimum and maximum values of

$$N/U := \frac{\text{Number of Processors Used}}{\text{Total Load}}$$

as obtained for the different assignment schemes in all task sets. The data in Table III was collected from the same simulation experiments used for Figs. 2-4.






Fig. 2. Task sets with $\alpha = 0.2$.

5. In all simulations, the results obtained with RMFF and RMBF were identical.

Fig. 3. Task sets with $\alpha = 0.5$.Fig. 4. Task sets with $\alpha = 0.8$.TABLE III
SIMULATION RESULTS FOR N/U

	RMNF		RMFF		RMST		RMGT	
	min.	max.	min.	max.	min.	max.	min.	max.
$\alpha = 0.2$	1.42	1.55	1.30	1.45	1.06	1.20	1.06	1.20
$\alpha = 0.5$	1.50	1.64	1.29	1.36	1.15	1.26	1.14	1.22
$\alpha = 0.8$	1.50	1.67	1.27	1.38	1.30	1.44	1.18	1.33

VI. CONCLUSIONS

We derived new schedulability conditions for scheduling periodic real-time tasks on uniprocessor and multiprocessor systems. Each processor was assumed to use the rate-monotonic scheduling algorithm. Based on the schedulability conditions we developed two task assignment schemes for multiprocessor systems, called *RMST* and *RMGT*.

For each scheme we obtained upper bounds for the performance parameter N/N_{opt} , where N is the number of processors required to schedule a task set with *RMST* or *RMGT*, and N_{opt} is the number of processors needed in an optimal assignment. We also obtained lower bounds for the average processor utilization. We provided asymptotic limits of these bounds. For *RMST*, the asymptotic bound for N/N_{opt} was proven to be $1/(1 - \alpha)$ where α is the maximal load factor of the tasks in the given task set. For the *RMGT* scheme we proved an asymptotic bound of $7/4$. We compared the average-case performance of the *RMST* and *RMGT* schemes with existing schemes via simulations, and showed that our new schemes are also superior if worst case conditions do not hold.

Note that the performance improvement of the *RMST* and *RMGT* schemes, compared to previously existing results, were achieved with rather simple assignment algorithms. The strength of the presented schemes resulted from the novel schedulability conditions.

APPENDIX PROOF OF LEMMA 4

For $K \leq 2$ the statement is trivially correct. For $K > 2$, consider a task set that satisfies the assumption but violates (57). We have to distinguish two cases. Either, $y_{j-1} > y_j$ for some index $1 < j \leq K$, or $y_K > 2y_1$. We will only prove the first case, since the proof for the second case is completely analogous.

Assume that there exists an index $1 < j \leq K$ with $y_{j-1} > y_j$. Then we obtain by definition of m_i , that for all $i < j - 1$,

$$m_i \leq y_j < y_{j-1} \quad (125)$$

and for all $i > j$,

$$m_i \leq 2y_j < 2y_{j-1} \quad (126)$$

But with the assumption in (56) we must have that $m_j = 2y_{j-1}$. By definition of m_j , this means

$$2y_j < 2y_{j-1} \leq \begin{cases} 2y_i & \text{if } i < j-1 \\ y_i & \text{if } i > j \end{cases} \quad (127)$$

Using (56), we have

$$m_i = \begin{cases} 2y_j & \text{if } i < j \\ 2y_{j-1} & \text{if } i = j \\ y_j & \text{if } i > j \end{cases} \quad (128)$$

It follows that (58) holds with $a = y_{j-1}$ and $b = y_j$. \square

REFERENCES


- [1] A. Burchard, Y. Oh, J. Liebeherr, and S.H. Son, "A linear time online task assignment scheme for multiprocessor systems," *Proc. 11th IEEE Workshop Real-Time Operating Systems and Software*, pp. 28-31, May 1994.
- [2] S. Davari and S.K. Dhall, "An on line algorithm for real-time allocation," *19th Ann. Hawaii Int'l Conf. System Sciences*, pp. 133-141, 1986.
- [3] S. Davari and S.K. Dhall, "An on line algorithm for real-time allocation," *IEEE Real-Time Systems Symp.*, pp. 194-200, 1986.
- [4] S.K. Dhall, "Scheduling periodic-time-critical jobs on single processor and multiprocessor computing systems," PhD thesis, Univ. of Illinois at Urbana-Champaign, 1977.
- [5] S.K. Dhall and C.L. Liu, "On a real-time scheduling problem," *Operations Research*, vol. 26, no. 1, pp. 127-140, Jan./Feb. 1978.
- [6] D.S. Johnson, A. Demers, J.D. Ullman, M.R. Garey, and R.L. Graham, "Worst case performance bounds for simple one-dimensional packing algorithms," *SIAM J. Computing*, vol. 3, pp. 299-325, 1974.
- [7] J.P. Lehoczky, L. Sha, and Y. Ding, "The rate-monotonic scheduling algorithm: Exact characterization and average behavior," *IEEE Real-Time Systems Symp.*, pp. 166-171, 1989.
- [8] J.Y.-T. Leung and J. Whitehead, "On the complexity of fixed-priority scheduling of periodic real-time tasks," *Performance Evaluation*, vol. 2, pp. 237-250, 1982.
- [9] C.L. Liu and J.W. Layland, "Scheduling algorithms for multiprogramming in a hard real time environment," *J. ACM*, vol. 20, no. 1, pp. 46-61, Jan. 1973.

- [10] Y. Oh and S.H. Son, "Tight performance bounds of heuristics for a real-time scheduling problem." Technical Report CS-93-24, Univ. of Virginia, Dept. of Computer Science, May 1993.
- [11] K. Ramamritham, J.A. Stankovic, and W. Zhao, "Distributed scheduling of tasks with deadlines and resource requirements," *IEEE Trans. Computers*, vol. 38, no. 8, pp. 1,110-1,123, Aug. 1989.
- [12] L. Sha, J.P. Lehoczky, and R. Rajkumar, "Solutions for some practical problems in prioritized preemptive scheduling," *IEEE Real-Time Systems Symp.*, pp. 181-191, 1986.
- [13] *Hard Real-Time Systems*, J.A. Stankovic and K. Ramamritham, eds. Los Alamitos, Calif.: IEEE CS Press, 1988.



Almut Burchard received the Diplom-Mathematiker degree from the University of Heidelberg, Germany, in 1989, and the PhD in mathematics from the Georgia Institute of Technology in 1994. She has been an instructor of mathematics at Princeton University since 1994. Her principal field of interest is functional analysis.

Dr. Burchard is a member of the American Mathematical Society and the International Association of Mathematical Physics.



Jörg Liebeherr (S'88-M'92) received the Diplom-Informatiker degree from the University of Erlangen-Nürnberg, Germany, in 1988, and the PhD degree in computer science from the Georgia Institute of Technology in 1991. From 1990-1991, he did his dissertation work partly at the IBM T.J. Watson Research Center. In 1992, he was a postdoctoral fellow in the Computer Science Division of the University of California, Berkeley. He has been an assistant professor in the Department of Computer Science at the University of Virginia since September 1992.

Dr. Liebeherr's research interests are computer networks, distributed multimedia systems, real-time systems, and performance evaluation. He is a member of the IEEE and ACM (SIGCOMM and SIGMETRICS).



Yingfeng Oh received a BS degree from Zhenjiang ShipBuilding College, Jiangsu Province, China, in 1985, an MS degree from the Institute of Computing Technology, Academia Sinica, Beijing, China, in 1988, and a PhD degree from the University of Virginia, Charlottesville, in 1994. The degrees are all in computer science.

Dr. Oh is currently a research associate in the Department of Computer Science at the University of Virginia. His research interests include real-time scheduling theory, databases, multimedia systems, real-time networking, and parallel systems.



Sang Hyuk Son (S'84-M'86) received the BS degree in electronics engineering from Seoul National University in 1976, and the MS degree in electrical engineering from the Korea Advanced Institute of Science and Technology in 1978. He earned the MS and PhD degrees in computer science from the University of Maryland, College Park, in 1984 and 1986, respectively. In 1986, he joined the Department of Computer Science at the University of Virginia, where he is currently an associate professor.

His current research interests include real-time computing, database systems, distributed systems, and fault-tolerant computing. He has authored or coauthored more than 80 papers and several book chapters in these areas.

Dr. Son served as a guest editor for the *ACM SIGMOD Record* special issue on real-time database systems and as the program chair of the 10th IEEE Workshop on Real-Time Operating Systems and Software (RTOS) held in May 1993. He was the general chair of the 11th RTOS, held in May 1994. He is the editor of the book *Advances in Real-Time Systems* (Prentice Hall, 1995). He also served as an ACM lecturer in 1991-1993. He is a member of the ACM and the IEEE Computer Society.