Mälardalen University Press Licentiate Theses
No. 119

PARTITIONED SCHEDULING OF REAL-TIME TASKS ON MULTI-CORE PLATFORMS

Farhang Nemati

2010

MÄLARDALEN UNIVERSITY
SWEDEN

School of Innovation, Design and Engineering

ISSN 1651-9256
ISBN 978-91-86135-74-4
Printed by Mälardalen University, Västerås, Sweden

Populärvetenskaplig sammanfattning

Klassiska programvarusystem som exempelvis ordbehandlare, bildbehandlare och webbläsare har typiskt en förväntad funktion att uppfylla, till exempel, en användare ska kunna producera typsatt skrift under relativt smärtfria former. Man kan generalisera och säga att korrekt funktion är av yttersta vikt för hur populär och användbar en viss programvara är medans exakt hur en viss funktion realiseras är av underordnad betydelse. Tittar man istället på så kallade realtidssystem så är, utöver korrekt funktionalitet hos programvaran, också det tidsmässiga utförandet av funktionen av yttersta vikt. Med andra ord så bör, eller måste, de funktionella resultaten produceras inom vissa specificerade tidsramar. Ett exempel är en airbag som inte får utlösas för tidigt eller för sent. Detta kan tyckas relativt okomplicerat, men tittar man närmare på hur realtidssystem är konstruerade så finner man att ett system vanligtvis är uppdelat i ett antal delar som körs (exekveras) parallellt. Dessa delar kallas för tasks och varje task är en sekvens (del) av funktionalitet, eller instruktioner, som genomförs samtidigt med andra tasks. Dessa tasks exekverar på en processor, själva hjärnan i en dator. Realtidsanalyser har tagits fram för att förutsäga hur sekvenser av taskexekveringar kommer att ske givet att antal tasks och deras karakteristik.

Utvecklingen och modernisering av processorer har tvingat fram så kallade multicoreprocessorer - processorer med multipla hjärnor (cores). Tasks kan nu, jämfört med hur det var förr, köras parallellt med varandra på olika cores, vilket samtidigt förbättrar effektiviteten hos en processor med avseende på hur mycket som kan exekveras, men även komplicerar både analys och förutsägbarhet med avseende på hur dessa tasks körs. Analys behövs för att kunna förutsäga korrekt tidsmässigt beteende hos programvaran i ett realtidssystem.

I denna licentiatavhandling har vi föreslagit en metod att fördela ett realtidssystems tasks på ett antal processorer givet en multicorearkitektur. Denna metod ökar avsevärt både prestation, förutsägbarhet och resursutnyttjandet hos multicorebaserade realtidsystemet genom att garantera tidsmässigt korrekt exekvering av programvarusystem med komplexa beroenden vilka har direkt påverkan på hur lång tid ett task kräver för att exekvera.

Abstract

In recent years multiprocessor architectures have become mainstream, and multi-core processors are found in products ranging from small portable cell phones to large computer servers. In parallel, research on real-time systems has mainly focused on traditional single-core processors. Hence, in order for real-time systems to fully leverage on the extra capacity offered by new multi-core processors, new design techniques, scheduling approaches, and real-time analysis methods have to be developed.

In the multi-core and multiprocessor domain there are mainly two scheduling approaches, global and partitioned scheduling. Under global scheduling each task can execute on any processor at any time while under partitioned scheduling tasks are statically allocated to processors and migration of tasks among processors is not allowed. Besides simplicity and efficiency of partitioned scheduling protocols, existing scheduling and synchronization methods developed for single-core processor platforms can more easily be extended to partitioned scheduling. This also simplifies migration of existing systems to multi-cores. An important issue related to partitioned scheduling is distribution of tasks among processors which is a bin-packing problem.

In this thesis we propose a partitioning framework for distributing tasks on the processors of multi-core platforms. Depending on the type of performance we desire to achieve, the framework may distribute a task set differently, e.g., in an application in which tasks process huge amounts of data the goal of the framework may be to decrease cache misses. Furthermore, we propose a blocking-aware partitioning heuristic algorithm to distribute tasks onto the processors of a multi-core architecture. The objective of the proposed algorithm is to decrease blocking overhead of tasks which reduces the total utilization and has the potential to reduce the number of required processors. Finally, we have implemented a tool to facilitate evaluation and comparison of different multiprocessor scheduling and synchronization approaches, as well as

different partitioning heuristics. We have applied the tool in the evaluation of several partitioning heuristic algorithms, and the tool is flexible to which any new scheduling or synchronization protocol as well as any new partitioning heuristic can easily be added.

Acknowledgments

First, I want to thank my supervisors, Thomas Nolte, Christer Norström, Anders Wall for guiding and helping me during my studies. I specially thank Nolte for all his support and encouragement.

I would like to give many thanks to the people who, with their support, have made PROGRESS to progress; Hans Hansson, Ivica Crnkovic, Paul Pettersson, Sasikumar Punnekkat, Björn Lisper, Mikael Sjödin, Kristina Lundkvist, Jan Gustafsson, Cristina Seceleanu, Frank Lüders, Jan Carlson, Dag Nyström, Andreas Ermedahl, Radu Dobrin, Daniel Sundmark, Rikard Land and Jukka Mäki-Turja.

I also thank people at IDT; Gunnar, Malin, Åsa, Harriet, Monica, Jenny, Monika, Else-Maj, Susanne, Maria and Carola for making many things easier.

During my studies, trips, coffee breaks and parties I have had a lot of fun and I wish to give many thanks to Aida, Aneta, Séverine, Hongyu, Pasqualina, Rafia, Kathrin, Ana, Sara, Eun-Young, Adnan, Andreas H., Moris, Hüseyin, Marcelo, Bob (Stefan), Luis (Yue), Mikael, Jagadish, Nikola, Rui, Holger, Federico, Saad, Mehrdad, Johan K., Johan F., Juraj, Luka, Leo, Josip, Antonio, Tibi, Lars, Rikard Li., Etienne, Thomas Le., Amine, Adam, Andreas G., Batu, Fredrik, Jörgen, Giacomo, and others for all the fun and memories.

I want to give my gratitude to my parents for their support and love in my life.

Last but not least, my special thanks goes to my wife, Samal, for all the support, love and fun.

This work has been supported by the Swedish Foundation for Strategic Research (SSF), via the research programme PROGRESS.

Farhang Nemati
Västerås, May, 2010

List of Publications

Papers Included in the Licentiate Thesis[1]

Paper A *Efficiently Migrating Real-Time Systems to Multi-Cores*. Farhang Nemati, Moris Behnam, Thomas Nolte. In 14th IEEE International Conference on Emerging Techonologies and Factory (ETFA'09), pages 1205-1212, September, 2009.

Paper B *Blocking-Aware Partitioning for Multiprocessors*. Farhang Nemati, Thomas Nolte, Moris Behnam. Technical Report, MRTC (Mälardalen Real-Time Research Centre), Mälardalen University, March, 2010.

Paper C *Partitioning Real-Time Systems on Multiprocessors with Shared Resources*. Farhang Nemati, Thomas Nolte, Moris Behnam. In submission.

Paper D *A Flexible Tool for Evaluating Scheduling, Synchronization and Partitioning Algorithms on Multiprocessors*. Farhang Nemati, Thomas Nolte. In submission.

---

[1] The included articles have been reformatted to comply with the licentiate layout

Additional Papers, not Included in the Licentiate Thesis

Conferences and Workshops

- *Multiprocessor Synchronization and Hierarchical Scheduling*. Farhang Nemati, Moris Behnam, Thomas Nolte. In 38th International Conference on Parallel Processing (ICPP'09) Workshops, pages 58-64, September, 2009.

- *Investigation of Implementing a Synchronization Protocol under Multiprocessors Hierarchical Scheduling*. Farhang Nemati, Moris Behnam, Thomas Nolte, Reinder J. Bril (Eindhoven University of Technology, The Netherlands). In 14th IEEE International Conference on Emerging Technologies and Factory (ETFA'09), pages 1670-1673, September, 2009.

- *Towards Hierarchical Scheduling in AUTOSAR*. Mikael Åsberg, Moris Behnam, Farhang Nemati, Thomas Nolte. In 14th IEEE International Conference on Emerging Techonologies and Factory (ETFA'09), pages 1181-1188, September, 2009.

- *An Investigation of Synchronization under Multiprocessors Hierarchical Scheduling*. Farhang Nemati, Moris Behnam, Thomas Nolte. In Work-In-Progress (WIP) Proceedings of the 21st Euromicro Conference on Real-Time Systems (ECRTS'09), pages 49-52, July, 2009.

- *Towards Migrating Legacy Real-Time Systems to Multi-Core Platforms*. Farhang Nemati, Johan Kraft, Thomas Nolte. In Work-In-Progress (WIP) track of the 13th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'08), pages 717-720, September, 2008.

- *Validation of Temporal Simulation Models of Complex Real-Time Systems*. Farhang Nemati, Johan Kraft, Christer Norström. In 32nd IEEE International Computer Software and Application Conference (COMPSAC'08), pages 1335-1340, July, 2008.

MRTC reports

- *A Framework for Real-Time Systems Migration to Multi-Cores*. Farhang Nemati, Johan Kraft, Thomas Nolte. MRTC report ISSN 1404-3041 ISRN MDH-MRTC-235/2009-1-SE, Mälardalen Real-Time Research Centre, Mälardalen University, 2009.

Contents

I

Thesis

Chapter 1

Introduction

Inherent in problems with power consumption and related thermal problems, multi-core platforms seem to be the way towards increasing performance of processors, and single-chip multiprocessors (multi-cores) are today the dominating technology for desktop computing. The performance achieved by multi-core architectures was previously only provided by High Performance Computing (HPC) systems. The HPC programmers are required to have a deep understanding of the respective hardware architecture in order to adjust the program explicitly for that hardware. This is not a suitable approach in embedded systems development, due to requirements on productivity, portability, maintainability, and short time to market.

The performance improvements of using multi-core processors depend on the nature of the applications as well as the implementation of the software. To take advantage of the concurrency offered by a multi-core architecture, appropriate algorithms have to be used to divide the software into tasks (threads) and distribute tasks fairly on processors to increase the overall performance. Real-time systems are typically multi threaded, hence they are easier to adapt to multi-core platforms than single-threaded, sequential programs. If the tasks are independent of eachother, they can run concurrently to improve performance. Looking at real-time systems, from a practical point of view, a static and manual assignment of processors is often preferred for predictability reasons. Real-time systems can highly benefit from multi-core architectures, as critical functionality can have dedicated cores and independent tasks can run concurrently. Moreover, since the processors are located on the same chip and typically have shared memory, communication between them is very fast.

Many of todays existing *legacy* real-time systems are very large and complex, typically consisting of millions of lines of code which have been developed and maintained for many years. Due to the huge development investments made in these legacy systems, it is normally not an option to throw them away and to develop a new system from scratch. A significant challenge when migrating legacy real-time systems to multi-core platforms is that they have been developed for uniprocessor (single-core) platforms where the execution model is actually sequential. Thus the software may need adjustments where assumptions of uniprocessor have impact.

Mainly, two approaches for scheduling real-time systems on multiprocessors exist [1, 2, 3, 4]; global and partitioned scheduling. Under global scheduling protocols, e.g., Global Earliest Deadline First (G-EDF), tasks are scheduled by a single scheduler and each task can be executed on any processor. A single global queue is used for storing tasks. A task can be preempted on a processor and resumed on another processor, i.e., migration of tasks among cores is permitted. Under a partitioned scheduling protocol, tasks are statically assigned to processors and the tasks within each processor are scheduled by a uniprocessor scheduling protocol, e.g., Rate Monotonic (RM) and EDF. Each processor is associated with a separate ready queue for scheduling task jobs. There are systems in which some tasks cannot migrate among cores while other tasks can migrate. For such systems neither global or partitioned scheduling methods can be used. A two-level hybrid scheduling approach [4], which is a mix of global and partitioned scheduling methods, is used for those systems.

In the multiprocessor research community, considerable work has been done on scheduling algorithms where it is assumed that tasks are independent. However in practice a typical real-time system includes tasks that share resources. On the other hand, synchronization in the multiprocessor context has not received enough attention. Under partitioned scheduling, if all tasks that share the same resource can be allocated on the same processor the uniprocessor synchronization protocols can be used [5]. This is not always possible, and some adjustments have to be done to the protocols to support synchronization of tasks across processors. The uniprocessor lock-based synchronization protocols have been extended to support inter processor synchronization among tasks [6, 7, 8]. However, under global scheduling methods, the uniprocessor synchronization protocols [9, 1] can not be reused without modification. Instead, new lock-based synchronization protocols have been developed to support resource sharing under global scheduling methods [10, 11].

Partitioned scheduling protocols have been used more often and are supported by commercial real-time operating systems [12], because of their simplicity, efficiency and predictability. However, they suffer from the problem of allocating tasks to processors (partitioning), which is a bin-packing problem [13] and is known to be a NP-hard problem in the strong sense. Thus, to take advantage of performance offered by multi-cores, partitioned scheduling protocols should be coordinated with appropriate partitioning (allocating tasks on processors) algorithms. Heuristic approaches and sufficient feasibility tests for bin-packing algorithms have been studied to find a near-optimal partitioning [2, 3]. However, the existing partitioning algorithms for multiprocessors (multi-cores) mostly assume independent tasks while in real applications, tasks often share resources.

1.1 Contributions

The main contributions of this thesis are as follows.

1. **Partitioning Framework**
   We have proposed a framework that coordinates partitioned scheduling with allocation of tasks (partitioning) on a multi-core platform. Depending on the application the coordination may be different, e.g., in an application in which tasks process huge amounts of data the goal of coordination may be decreasing cache misses, or in an application in which tasks heavily share resources, the coordination will be towards decreasing blocking overhead by allocating tasks sharing the same resources to the same processor as far as possible. Paper A directs this contribution.

2. **Partitioning Heuristic**
   We have proposed a partitioning algorithm, based on bin-packing, for allocating tasks onto processors of a multi-core platform (Chapter 3). Tasks can access mutually exclusive resources and the goal of the algorithm is to decrease the overall blocking overhead in the system. This may consequently increase the schedulability of a task set and reduce the number of processors. We proposed the the partitioning algorithm in Paper B. In Paper C we have further evaluated our algorithm and compared it to a similar algorithm originally proposed in [12].

3. **Implementation**
   We have implemented a tool to facilitate evaluation and comparison of different multiprocessor scheduling and synchronization approaches as well as different partitioning heuristics. We have implemented our partitioning algorithm together with a similar existing algorithm and added

them to the tool. By using the tool, we have performed experiments to evaluate the performance of our heuristic. This tool has been made extensible to allow easy addition of future protocols and algorithms. This contribution is directed by Paper D.

1.2 Thesis Outline

The outline of the thesis is as follows. In Chapter 2 we give a background describing of real-time systems, scheduling, multiprocessors, multi-core architectures, the problems and the existing solutions, e.g., scheduling and synchronization protocols. Chapter 3 gives an overview of our proposed partitioning framework, heuristic algorithm, and the evaluation tool. In Chapter 4 we present our conclusion and future work. We present the technical overview of the papers that are included in this thesis in Chapter 5, and we present these papers in Chapters 6 - 9.

Chapter 2

Background

2.1 Real-Time Systems

In a real-time system, besides the functional correctness of the system, the output should satisfy timing attributes as well [14], e.g., the outputs should be within deadlines. A real-time system is typically developed following a concurrent programming approach in which a system may be divided into several parts, called *tasks*, and each task, which is is a sequence of operations, executes in parallel with other tasks. A task may issue an infinite number of instances called *jobs* during run-time.

Each task has timing attributes, e.g., *deadline* before which the task should finish its execution, *Worst Case Execution Time* (WCET) which is the maximum time that a task needs to perform and complete its execution when executing without interference from other tasks. The execution of a task can be periodic or aperiodic; a periodic task is triggered with a constant time, denoted as *period*, in between instances, and an aperiodic task may be triggered at any arbitrary time instant.

Real-time systems are generally categorized into two categories; *hard real-time systems* and *soft real-time systems*. In a hard real-time system tasks are not allowed to miss their deadlines, while in a soft real-time system some tasks may miss their deadlines. A safety-critical system is a type of hard-real time system in which missing deadlines of tasks may lead to catastrophic incidents, hence in such a system missing deadlines are not tolerable.

2.2 Multi-core Platforms

A multi-core (single-chip multiprocessor) processor is a combination of two or more independent processors (cores) on a single chip. The cores are connected to a single shared memory via a shared bus. The cores typically have independent L1 caches and may share an on-chip L2 cache.

Multi-core architectures are today the dominating technology for desktop computing and are becoming the defacto processors. The performance of using multiprocessors, however, depends on the nature of the applications as well as the implementation of the software. To take advantage of the concurrency offered by a multi-core architecture, appropriate algorithms have to be used to divide the software into tasks (threads) and to distribute tasks on cores to increase the system performance. If an application is not (or can not) be fairly divided into tasks, e.g., one task does all the heavy work, a multi-core will not help improving the performance significantly. Real-time systems can highly benefit from multi-core processors, as they are typically multi-threaded, hence making it easier to adapt them to multi-cores than single-threaded, sequential programs, e.g., critical functionality can have dedicated cores and independent tasks can run concurrently to improve performance. Moreover, since the cores are located on the same chip and typically have shared memory, communication between cores is very fast.

Multi-core platforms introduce significant challenges, and existing software systems need adjustments to be adapted on multi-cores. Many existing legacy real-time systems are very large and complex, typically consisting of huge amount of code. It is normally not an option to throw them away and to develop a new system from scratch. A significant challenge is to adapt them to work efficiently on multi-core platforms. If the system contains independent tasks, it is a matter of deciding on which processors each task should be executed. In this case scheduling protocols from single-processor platforms can easily be reused. However, tasks are usually not independent and they may share resources. This means that, to be able to adapt the existing systems, synchronization protocols are required to be changed or new protocols have to be developed.

For hard real-time systems, from a practical point of view, a static assignment of processors, i.e., partitioned scheduling (Section 2.3.1), is often the more common approach [2], often inherent in reasons of predictability and simplicity. On the other hand, the well-studied and verified scheduling analysis methods from the single-processor domain has the potential to be reused. However, fairly allocating tasks onto processors (partitioning) is a challenge, which is a bin-packing problem.

Finally, the processors on a multi-core can be identical, which means that all processors have the same performance, this type of multi-core architectures are called *homogenous*. However, the architecture may suffer from heat and power consumption problems. Thus, processor architects have developed multi-core architectures consisting of processors with different performance in which tasks can run on appropriate processors, i.e., the tasks that do not need higher performance can run on processors with lower performance, decreasing energy consumption.

2.3 Real-Time Scheduling on Multiprocessors

The major approaches for scheduling real-time systems on multiprocessors are *partitioned scheduling*, *global scheduling*, and the combination of these two called *hybrid scheduling* [1, 2, 3, 4].

2.3.1 Partitioned Scheduling

Under partitioned scheduling tasks are statically assigned to processors, and the tasks within each processor are scheduled by a single-processor scheduling protocol, e.g., RM and EDF [15]. Each task is allocated to a processor on which its jobs will run. Each processor is associated with a separate ready queue for scheduling its tasks' jobs.

A significant advantage of partitioned scheduling is that well-understood and verified scheduling analysis from the uniprocessor domain can be reused. Another advantage is the run-time efficiency of these protocols as the tasks and jobs do not suffer from migration overhead. A disadvantage of partitioned scheduling is that it is a bin-packing problem which is known to be NP-hard in the strong sense, and finding an optimal distribution of tasks among processors (cores) in polynomial time is not generally realistic. Another disadvantage of partitioned scheduling algorithms is that prohibiting migration of tasks among processors decreases the utilization bound, i.e., it has been shown [3] that task sets exist that are only schedulable if migration among processors is allowed. Non-optimal heuristic algorithms have been used for partitioning a task set on a multiprocessor platform. An example of a partitioned scheduling algorithm is Partitioned EDF (P-EDF) [2].

2.3.2 Global Scheduling

Under global scheduling algorithms tasks are scheduled by a single system-level scheduler, and each task or job can be executed on any processor. A single global queue is used for storing ready jobs. At any time instant, at most m ready jobs with highest priority among all ready jobs are chosen to run on a multiprocessor consisting of m processors. A task or its jobs can be preempted on one processor and resumed on another processor, i.e., migration of tasks (or its corresponding jobs) among cores is permitted. An example of a global scheduling algorithm is Global EDF (G-EDF) [2]. The global scheduling algorithms are not necessarily optimal either, although in the research community new multiprocessor scheduling algorithms have been developed that are optimal. Proportionate fair (Pfair) scheduling approaches are examples of such algorithms [16, 17]. However, this particular class of scheduling algorithms suffer from high run-time overhead as they may have to increase the number of preemptions and migrations significantly.

2.3.3 Hybrid Scheduling

There are systems that cannot be scheduled by either pure partitioned or pure global scheduling; for example some tasks cannot migrate among cores while other tasks are allowed to migrate. An example approach for those systems is the two-level hybrid scheduling approach [4], which is based on a mix of global and partitioned scheduling methods. In such protocols, at the first level a global scheduler assigns jobs to processors and at the second level each processor schedules the assigned jobs by a local scheduler.

Recently more general approaches, such as cluster based scheduling [18, 19], have been proposed which can be categorized as a generalization of partitioned and global scheduling protocols. Using such an approach, tasks are statically assigned to clusters and tasks within each cluster are globally scheduled. In turn, clusters are transformed into tasks and are globally scheduled on a multiprocessor. Cluster-based scheduling can be physical or virtual. In physical cluster-based scheduling the processors of each cluster are statically mapped to a subset of processors of the multiprocessor [18]. In virtual cluster-based scheduling the processors of each cluster are dynamically mapped (one-to-many) onto processors of the multiprocessor. Virtual clustering is more general and less sensitive to task-cluster mapping compared to physical clustering.

2.4 Resource Sharing on Multiprocessors

In the multiprocessor domain, considerable work has been done on scheduling protocols, but usually under the assumption that tasks are independent. However in practice a typical real-time system must allow for resource sharing among tasks. Generally there are two classes of resource sharing, i.e., lock-based and lock-free synchronization protocols. In the lock-free approach [20], operations on simple software objects, e.g., stacks, linked lists, are performed by retry loops, i.e., operations are retried until the object is accessed successfully. The advantages of lock-free algorithms is that they do not require kernel support and as there is no need to lock, priority inversion does not occur. The disadvantage of these approaches is that it is not easy to apply them to hard real-time systems as the worst case number of retries is not easily predictable. In this thesis we have focused on a lock-based approach, thus in this section we present an overview of the existing lock-based synchronization methods.

On a multiprocessor platform a job, besides lower priority jobs, can be blocked by higher priority jobs (those that are assigned to different processors) as well. This does not rise any problem on uniprocessor platforms. Another issue, which is not the case in uniprocessor synchronization, is that on a uniprocessor, a job J_i can not be blocked by a lower priority job J_j arriving after J_i. However, on a multiprocessor, assuming jobs J_i and J_j are assigned on different processors, the lower priority job J_i can arrive later than the higher priority job J_i and block J_i. Those cases introduce more complexity and pessimism into schedulability analysis.

For multiprocessor systems, Rajkumar present MPCP (Multiprocessor Priority Ceiling Protocol) [6], which extends PCP [9] to multiprocessors allowing for synchronization of tasks sharing mutually exclusive resources using the partitioned Fixed Priority Scheduling (FPS) protocol.

Gai et al. [7, 8] present the MSRP (Multiprocessor Stack Resource Policy), which extends SRP [1] to multiprocessor platforms and works under the P-EDF scheduling protocol.

Lopez et al. [5] present an implementation of SRP under P-EDF. In this work they propose a solution in which all tasks that directly or indirectly share resources are allocated to the same processor and a uniprocessor synchronization protocol, i.e., SRP, is used to manage resource sharing within each processor. However, if all tasks that directly or indirectly share resources can not be allocated to the same processor the solution can not be used.

Block et al. [10] present FMLP (Flexible Multiprocessor Locking Protocol), which is the first synchronization protocol for multiprocessors that can be applied to both partitioned and global scheduling algorithms, i.e., P-EDF and G-EDF. An implementation of FMLP has been described in [21] and a comparison between FMLP and MPCP has been presented in [22].

Recently, Easwaran and Andersson have proposed a synchronization protocol [11] under global fixed priority scheduling protocol. In this paper they have derived schedulability analysis of the priority inheritance protocol under global scheduling algorithms, for the first time.

2.4.1 The Multiprocessor Priority Ceiling Protocol (MPCP)

Definition The MPCP is used for synchronizing a set of tasks sharing lock-based resources under a partitioned FPS protocol, i.e., RM. Under MPCP, resources are divided into *local* and *global* resources. Local resources are shared only among tasks from the same processor and global resources are shared by tasks assigned to different processors. The local resources are protected using a uniprocessor synchronization protocol, i.e., PCP. A task blocked on a global resource suspends making the processor available for the local tasks. A critical section in which a task performs a request for a global resource is called a *global critical section* (gcs). Similarly a critical section where a task requests for a local resource is denoted as a *local critical section* (lcs).

Under MPCP, the blocking time of a task, in addition to local blocking, has to include remote blocking terms where a task is blocked by tasks (with any priority) executing on other processors. However, the maximum remote blocking time of a job is bounded and is a function of the duration of critical sections of other jobs. This is a consequence of assigning any gcs a ceiling greater than the priority of any other task, hence a gcs can only be blocked by another gcs and not by any non-critical section. Assume ρ_H is the highest priority among all tasks. The priority of a job J_i executing within a gcs in which it requests R_k is called remote ceiling of gcs and equals to $\rho_H + 1 + \max\{\rho_j | \tau_j \textit{ requests } R_k \textit{ and } \tau_j \textit{ is not on } J_i\textit{'s processor}\}$.

Global critical sections cannot be nested in local critical sections and vice versa. Global resources potentially lead to high blocking times, thus tasks sharing the same resources are preferred to be assigned to the same processor as far as possible. We have proposed an algorithm that attempts to reduce the blocking times by assigning tasks to appropriate processors (Chapter 3).

To determine the schedulability of each processor under RM scheduling the following test is performed:

$$\forall k\, 1 \leq i \leq n, \sum_{k=1}^{i} C_k/T_k + B_i/T_i \leq i(2^{1/i} - 1) \tag{2.1}$$

where n is the number of tasks assigned to the processor, and B_i is the maximum blocking time of task τ_i which includes remote blocking factors as well as local blocking time. However this condition is sufficient but not necessary. Thus for more precise schedulability, a test of task response time [23] can be performed.

Blocking times under MPCP Before explaining the blocking factors of the blocking time of a job, the following terminology has to be explained:

- n_i^G: The number of global critical sections of task τ_i.
- $\{J'_{i,r}\}$: The set of jobs on processor P_r (other than J_i's processor) with global critical sections having priority higher than the global critical sections of jobs that can directly block J_i.
- $NH_{i,r,k}$: The number of global critical sections of job $J_k \in \{J'_{i,r}\}$ having priority higher than a global critical section on processor P_r that can directly block J_i.
- $\{GR_{i,k}\}$: The set of global resources that will be locked by both J_i and J_k.
- $NC_{i,k}$: The number of global critical sections of J_k in which it request a global resource in $\{GR_{i,k}\}$.
- β_i^{local}: The longest local critical section among jobs with a priority lower than that of job J_i executing on the same processor as J_i which can block J_i.
- $\beta L_i^{\text{global}}$: The longest global critical section of any job J_k with a priority lower than that of job J_i executing on a different processor than J_i's processor in which J_k requests a resource in $\{GR_{i,k}\}$.
- $\beta H_{i,k}^{\text{global}}$: The longest global critical section of job J_k with a priority higher than that of job J_i executing on a different processor than J_i's processor. In this global critical section, J_k requests a resource in $\{GR_{i,k}\}$.

- ${\beta'_{i,k}}^{\text{global}}$: The longest global critical section of job $J_k \in \{J'_{i,r}\}$ having priority higher than a global critical section on processor P_r that can directly block J_i.
- $\beta^{\text{lg}}_{i,k}$: The longest global critical section of a lower priority job J_k on the J_i's host processor.

The maximum blocking time B_i of task τ_i is a summation of five blocking factors:

$$B_i = B_{i,1} + B_{i,2} + B_{i,3} + B_{i,4} + B_{i,5}$$

where:

1. $B_{i,1} = n_i^G \beta_i^{\text{local}}$ each time job J_i is blocked on a global resource and suspends, the local lower priority jobs may execute and lock local resources and block J_i when it resumes.
2. $B_{i,2} = n_i^G \beta L_i^{\text{global}}$ when a job J_i is blocked on a global resource which is locked by a lower priority job executing on another processor.
3. $B_{i,3} = \sum_{\substack{\rho_i \leq \rho_k \\ J_k \textit{ is not on } J_i\textit{'s processor}}} NC_{i,k} \lceil T_i/T_k \rceil \beta H^{\text{global}}_{i,k}$ when higher priority jobs on processors other than J_i's processor block J_i.
4. $B_{i,4} = \sum_{\substack{J_k \in \{J'_{i,r}\} \\ P_r \neq J_i\textit{'s processor}}} NH_{i,r,k} \lceil T_i/T_k \rceil {\beta'_{i,k}}^{\text{global}}$ when the gcs's of lower priority jobs on processor P_r (different from J_i's processor) are preempted by higher priority gcs's of $J_k \in \{J'_{i,r}\}$.
5. $B_{i,5} = \sum_{\substack{\rho_i \leq \rho_k \\ J_k \textit{ is on } J_i\textit{'s processor}}} \min n_i^G + 1, n_k^G \beta^{\text{lg}}_{i,k}$ when J_i is blocked on global resources and suspends a local job J_k can execute and enter a global section which can preempt J_i when it executes in non-gcs sections.

2.4.2 The Multiprocessor Stack Resource Policy (MSRP)

Definition The MSRP is used for synchronizing a set of tasks sharing lock-based resources under a partitioned EDF (P-EDF). The shared resources are classified as either (i) local resources that are shared among tasks assigned to the same processor, or (ii) global resources that are shared by tasks assigned to different processors. Under MSRP, tasks synchronize local resources using SRP, and access to global resources is guaranteed a bounded blocking time.

Further, under MSRP, when a task is blocked on a global resource it performs *busy wait* (spin lock). This means that the processor is kept busy without doing any work, hence the duration of spin lock should be as short as possible which means locking a global resource should be reduced as far as possible. To achieve this goal under MSRP, the tasks executing in global critical sections become non-preemptive. The tasks blocked on a global resource are added to a FCFS (First Come First Served) queue. Global critical sections are not allowed to be nested under MSRP.

Gai et al. in [8] compare their implementation of MSRP to MPCP. They point out that the complexity of implementation as a disadvantage of MPCP and that wasting more local processor time (due to busy wait) as a disadvantage of MSRP. They have performed two case studies for the comparison. The results show that MPCP works better when the duration of global critical sections are increased while MSRP outperforms MPCP when critical sections become shorter. Also in applications where tasks access many resources, and resources are accessed by many tasks, which lead to more pessimism in MPCP, MSRP has a significant advantage compared to MPCP.

Blocking times under MSRP Under MSRP, if a task's job, J_i, attempts to access a global resource, R_q, it becomes non-preemptive. If the resource R_q is free it locks the resource, but if R_q is already locked by another job running on a different processor, J_i performs busy wait. The upper bound of busy wait time that any job executing on processor P_k can experience to access a global resource R_q is as follows.

$$\text{spin}(P_k, R_q) = \sum_{\forall P_l \neq P_k} \max_{\forall J_j \text{ on } P_l} (|Cs_{j,q}|) \tag{2.2}$$

where $|Cs_{j,q}|$ refers to the length of any critical section $Cs_{j,q}$ of J_j accessing R_q.

As a job performs busy wait its global critical sections become longer and consequently its Worst Case Execution Time (WCET) is increased. Thus, the worst case time any job, J_i executing on processor P_k, busy waits can be added to its WCET:

$$C'_i = C_i + \sum_{\forall \text{ global } R_q \text{ accessed by } J_i} \text{spin}(P_k, R_q) \tag{2.3}$$

where C'_i is the actual worst case execution time of J_i.

According to MSRP (similar to SRP), a job can be blocked only once, and as it starts executing it cannot be blocked. The worst case blocking time of any job J_i executing on processor P_k, is calculated as follows:

$$B_i = \max(B_i^{\text{local}}, B_i^{\text{global}}) \tag{2.4}$$

where B_i^{local} and B_i^{local} are the worst case blocking overhead from local resources and global resources respectively, and being defined as follows:

$$\begin{aligned} B_i^{\text{local}} = \max\{|Cs_{j,q}| \mid \ & (J_j \text{ is on } P_k) \ \wedge \ (R_q \text{ is local}) \ \wedge \ (\lambda_i > \lambda_j) \wedge \\ & (\lambda_i \leq \text{ceil}(R_q))\} \end{aligned} \tag{2.5}$$

where λ_i is the *static preemption level* of J_i [1].

$$\begin{aligned} B_i^{\text{global}} = \max\{|Cs_{j,q}| + \text{spin}(P_k, R_q) \mid \ & J_j \text{is not on } P_k \wedge \\ & R_q \text{ is global}\} \end{aligned} \tag{2.6}$$

2.4.3 The Flexible Multiprocessor Locking Protocol (FMLP)

Definition In FMLP, resources are categorized into *short* and *long* resources, and wether a resource is short or long is user specified. There is no limitation on nesting resource accesses, except that requests for long resources cannot be nested in requests for short resources.

Under FMLP, deadlock is prevented by grouping resources. A group includes either global or local resources, and two resources are in the same group if a request for one is nested in a request for the other one. A group lock is assigned to each group and only one task from the group at any time can hold the lock.

The jobs that are blocked on short resources perform busy-wait and are added to a FIFO queue. Jobs that access short resources hold the group lock and execute non-preemptively. A job accessing a long resource under G-EDF holds the group lock and executes preemptively using priority inheritance, i.e., it inherits the maximum priority of any higher priority job blocked on any resource within the same group. Tasks blocked on a long resource are added to a FIFO queue.

Under global scheduling, FMLP actually works under a variant of G-EDF for Suspendable and Non-preemptable jobs (GSN-EDF) [10] which guarantees that a job J_i can only be blocked (with a constraint duration) by another non-preemptable job when J_i is released or resumed.

Blocking times under FMLP In FMLP, any job J_i can face three types of blocking overhead:

- *Busy-wait blocking* of job J_i, specified by BW_i, is the maximum duration of time that the job can busy-wait on a short resource.
- *Non-preemptive blocking* occurs when a preemptable job J_i is one of the m highest priority jobs but it is not scheduled because a lower priority job is non-preemptively executing instead. Non-preemptive blocking of J_i denoted by NPB_i is the maximum duration of time that J_i is non-preemptively blocked.
- *Direct blocking* occurs when job J_i is one of the m highest priority jobs but it is suspended because it issues a request for an outermost long resource from group G but another job holds a resource from the same group (holds the group's lock). Direct blocking of job J_i specified by DB_i is the maximum duration of time that J_i can be direct blocked.

The worst case blocking time of any job J_i is the summation of the three sources of blocking times:

$$B_i = \mathrm{BW}_i + \mathrm{NPB}_i + \mathrm{DB}_i \tag{2.7}$$

The detailed calculations of the three sources of blocking times are presented in the appendix to the online version of [10][1].

2.5 Assumptions of the Thesis

With respect to the above presented background material, the work presented in this thesis has been developed under the following limitations:

Real-Time Systems:
We assume hard real-time systems.

Multi-core Architecture:
We assume identical multi-core architectures. However, as future work we believe that this assumption can be relaxed.

---

[1] Available at http://www.cs.unc.edu/˜anderson/papers/rtcsa07along.pdf

Scheduling Protocol:

The focus of this thesis is partitioned scheduling approaches. In the future we will extend our research to global and hybrid scheduling protocols as well.

Synchronization Protocol:

We have focused on MPCP as the synchronization protocol under which our heuristic attempts to decrease blocking overhead, and extending the heuristic to other protocols remains a future work.

Chapter 3

Heuristic Methods for Partitioning Task Sets on Multiprocessors

In this chapter we present a partitioning framework in which a task set is attempted to be efficiently allocated onto a single-chip shared memory multiprocessor (multi-core) platform with identical processors.

A scheduling framework for multi-core processors is presented by Rajagopalan et al. [24]. The framework tries to balance between the abstraction level of the system and the performance of the underlying hardware. The framework groups dependant tasks, which, for example, share data, to improve the performance. The paper presents Related Thread ID (RTID) as a mechanism to help the programmers to identify groups of tasks.

An approach for migration to multi-core is presented by Lindhult in [25]. The author presents the parallelization of sequential programs as a way to achieve performance on multi-core processors. The targeted language is PLEX, Ericsson's in-house developed event-driven real-time programming language used for Ericsson's telephone exchange system.

The *grey-box* modeling approach for designing real-time embedded systems is presented in [26]. In the grey-box task model the focus is on task-level abstraction and it targets performance of the processors as well as timing constraints of the system.

Furthermore,we have proposed a heuristic blocking-aware algorithm to allocate a task set on a multi-core platform to reduce the blocking overhead of tasks.

Partitioning (allocation tasks on processors) of a task set on a multiprocessor platform is a bin-packing problem which is known to be a NP-hard problem in the strong sense; therefore finding an optimal solution in polynomial time is not realistic in the general case [13]. Heuristic algorithms have been developed to find near-optimal solutions.

A study of bin-packing algorithms for designing distributed real-time systems is presented in [27]. The presented method partitions a software into modules to be allocated on hardware nodes. In their approach they use two graphs; a graph which models software modules and a graph that represents the hardware architecture. The authors extend the bin-packing algorithm with heuristics to minimize the number of required bins (processors) and the required bandwidth for the communication between nodes.

Liu et al. [28] present a heuristic algorithm for allocating tasks in multi-core-based massively parallel systems. Their algorithm has two rounds; in the first round processes (groups of threads - partitions in this thesis) are assigned to processing nodes, and the second round allocates tasks in a process to the cores of a processor. However, the algorithm does not consider synchronization between tasks.

Baruah and Fisher have presented a bin-packing partitioning algorithm, the *first-fit decreasing* (FFD) algorithm, in [29] for a set of independent sporadic tasks on multiprocessors. The tasks are indexed in non-decreasing order based on their relative deadlines, and the algorithm assigns the tasks to the processors in first-fit order. The tasks on each processor are scheduled under uniprocessor EDF.

Lakshmanan et al. [12] investigate and analyze two alternatives of execution control policies (suspend-based and spin-based remote blocking) under MPCP. They have developed a blocking-aware task allocation algorithm, an extension to the *best-fit decreasing* (BFD) algorithm, and evaluated it under both execution control policies. Their blocking-aware algorithm is of great relevance to our proposed algorithm, hence we have presented their algorithm in more detail in Section 3.3. Together with our algorithm we have also implemented and evaluated their blocking-aware algorithm and compared the performances of both algorithms.

3.1 Task and Platform Model

Our target system is a task set that consists of n sporadic tasks, $\tau_i(T_i, C_i, \rho_i, \{c_{i,p,q}\})$ where T_i is the minimum inter-arrival time between two successive jobs of task τ_i with worst-case execution time C_i and ρ_i as its priority. The tasks share a set of resources, R, which are protected using semaphores. The set of critical sections, in which task τ_i requests resources in R, is denoted by $\{c_{i,p,q}\}$, where $c_{i,p,q}$ indicates the maximum execution time of the p^{th} critical section of task τ_i, in which the task locks resource $R_q \in R$. Critical sections of tasks should be sequential or properly nested. The deadline of each job is equal to T_i. A job of task τ_i, is specified by J_i. The utilization factor of task τ_i is denoted by u_i where $u_i = C_i/T_i$.

We have also assumed that the multiprocessor (multi-core) platform is composed of identical, unit-capacity processors (cores) with shared memory. The task set is partitioned into partitions $\{P_1, \ldots, P_m\}$, where m represent the number of required processors and each partition is allocated onto one processor (core).

3.2 Partitioning Framework for Multi-cores

In this section we present our framework in which tasks are grouped into partitions and each partition is allocated onto one core (processor). At each step when a task is assigned to a partition the following requirements should be satisfied:

1. All partitions are schedulable.
2. The best partition for assigning the task is chosen in a way that the cost is minimized.

Cost function In the framework a cost function is used to calculate the cost of assigning a task to a partition. The cost function can be derived from constraints and preferences which are extracted from the system as well as those offered by the system experts. In the proposed framework, we use a set of constraints and preferences to derive the cost function and to test the schedulability of each partition (processor). The following constraints are used:

1. **Timing constraints**
   Specify timing attributes, e.g., deadline, Worst Case Execution Time

(WCET). Those constraints are used for schedulability test of each partition.

2. **Resource sharing constraints**
   In a typical real-time system, as tasks share resources the corresponding constraints should be considered for schedulability analysis. These constraints together with the timing constraints may be used for deriving the cost function as well.

3. **Task preferences**
   This may include more than one category of preferences. Each category consists (as a matrix) of cost values for each pair of tasks, when they are co-allocated on the same partition. These preferences facilitate allocation of tasks on the processors (partitions) by attracting dependent tasks on the same processor and forcing independent tasks to be allocated on different processors as far as possible. Each category (matrix) of cost values represents an aspect of system performance, e.g., an aspect can be increasing cash hits, or reducing blocking times (Section 3.3). The importance of each category is indicated by a coefficient. The number of categories as well as the value of their coefficient (their importance) depend on the *partitioning strategy*.

Partitioning strategy A partitioning strategy indicates the importance of the type of system performance (e.g., increasing cache hits, decreasing blocking overhead, etc.) we wish to achieve and gives a coefficient parameter to each matrix. The value of each coefficient depends on the importance of the performance that the matrix represents. For example in a system that processes large amounts of data the partitioning strategy can be that the tasks which share data heavily are assigned to the same partition to increase cache hits. Similarly in a system in which tasks share mutually exclusive resources, the target partitioning strategy can be assigning tasks sharing the same resources to the same processor as far as possible. This is the concrete partitioning strategy of our blocking-aware algorithm presented in Section 3.3.

Task weight Generally looking at bin-packing algorithms, e.g., the best-fit decreasing (BFD), objects are allocated into bins in the order of their size, e.g., the heavier objects are packed first. In the context of allocation of tasks onto processors, with independent tasks the utilization of the tasks are considered as their size. However, with dependent tasks other parameters (depending on the partitioning strategy) should be considered in their size (weight). The weight of a task indicates the importance of the task according to the partitioning strategy. For example in a partitioning strategy for reducing inter-core communication, the weight of a task may include the total number of messages it sends or receives during its execution time, or in a partitioning strategy for reducing blocking times of tasks their weight (size) should include blocking parameters.

3.3 Heuristic Partitioning Algorithms with Resource Sharing

In this section we present our proposed blocking-aware heuristic algorithm to allocate tasks onto the processors of a single chip multiprocessor (multi-core) platform. The algorithm extends a bin-packing algorithm with synchronization parameters. The results of our experimental evaluation [30] shows significant performance increment compared to the existing similar algorithm [12] and a reference *blocking-agnostic* bin-packing algorithm. The blocking-agnostic algorithm, in the context of this thesis, refers to a bin-packing algorithm that does not consider blocking parameters to increase the performance of partitioning, although blocking times are included in the schedulability test.

In our algorithm task constraints are identified, e.g., dependencies between tasks, timing attributes, and resource sharing preferences, and we extend the best-fit decreasing (BFD) bin-packing algorithm with blocking time parameters. The objective of the heuristic is (based on the constraints and preferences) to decrease blocking overheads by assigning tasks to appropriate processors (partitions).

In a blocking-agnostic BFD algorithm, bins (processors) are ordered in non-increasing order of their utilization and tasks are ordered in non-increasing order of their size (utilization). The algorithm attempts to allocate the task from the top of the ordered task set onto the first processor that fits it (i.e., the first processor on which the task can be allocated while all processors are schedulable), beginning from the top of the ordered processor list. If none of the processors can fit the task, a new processor is added to the processor list. At each step the schedulability of all processors should be tested, because allocating a task to a processor can increase the remote blocking time of tasks previously allocated to other processors and may make the other processors unschedulable. This means, it is possible that some of the previous processors become unschedulable even if a task is allocated to a new processor, which makes the algorithm fail.

The algorithm proposed in [12] was called Synchronization-Aware Partitioning Algorithm, and we call our algorithm Blocking-Aware Partitioning Algorithm. However, to ease refereing them, from now on we refer them as SPA and BPA respectively. In practice, industrial systems mostly use Fixed Priority Scheduling (FPS) protocols. To our knowledge the only synchronization protocol under fixed priority partitioned scheduling, for multiprocessor platforms, is MPCP. Both our algorithm (BPA) and the existing one (SPA) assume that MPCP is used for lock-based synchronization. Thus, we derive heuristics based on the blocking parameters in MPCP. However, our algorithm can be easily extended to other synchronization protocols, e.g., MSRP.

3.3.1 Blocking-Aware Algorithm (BPA)

The algorithm attempts to allocate a task set onto processors in two rounds. The output of the round with better partitioning results will be chosen as the output of the algorithm. In each round the tasks are allocated to the processors (partitions) in a different way. When a bin-packing algorithm allocates an object (task) to a bin (processor), it usually attempts to put the object in a bin that fits it better, and it does not consider the unallocated objects. The rationale behind the two rounds is that the heuristic tries to consider both past and future by looking at tasks allocated in the past and those that are not allocated yet. In the first round the algorithm considers the tasks that are not allocated to any processor yet; and attempts to take as many as possible of the best related tasks (based on remote blocking parameters) with the current task. In the second round it considers the already allocated tasks and tries to allocate the current task onto the processor that contains best related tasks to the current task. In the second round, the algorithm performs more like the usual bin packing algorithms (i.e., attempts to find the best bin for the current object). Briefly, the algorithm in the first round looks at the future and in the second round it considers the past.

Before starting the two rounds the algorithm performs some basic steps:

- A heuristic weight is assigned to each task which is a function of task's utilization as well as the blocking parameters that lead to potential remote blocking time interfered by other tasks:

$$w_i = u_i + \left\lceil \left(\sum_{\rho_i < \rho_k} \mathrm{NC}_{i,k} \beta_{i,k} \left\lceil \frac{T_i}{T_k} \right\rceil + \mathrm{NC}_i \max_{\rho_i \geq \rho_k} \beta_{i,k} \right) / T_i \right\rceil \tag{3.1}$$

where, $\mathrm{NC}_{i,k}$ is the number of critical sections of τ_k in which it shares a resource with τ_i and $\beta_{i,k}$ is the longest critical section among them, and NC_i is the total number of critical sections of τ_i.

Considering the remote blocking terms of MPCP (Section 2.4.1), the rationale behind the definition of weight is that the tasks that can be punished more by remote blocking become heavier. Thus, they can be allocated earlier and attract as many as possible of the tasks with which they share resources.

- Next, the *macrotasks* are generated. A macrotask includes tasks that directly or indirectly share resources, e.g., if tasks τ_i and τ_j share resource R_p and tasks τ_j and τ_k share resource R_q, all three tasks belong to the same macrotask. A macrotask has two alternatives; it can either be broken or unbroken. A macrotask is set as broken if it cannot fit in one processor (i.e., it can not be scheduled by a single processor even if no other task is allocated onto the processor), otherwise it is set as unbroken. If a macrotask is unbroken, the partitioning algorithm always allocate all tasks in the macrotask to the same partition (processor). Thus, all resources shared by tasks within the macrotask will be local. However, tasks within a broken macrotask have to be distributed into more than one partition. Similar to tasks, a weight is assigned to each macrotask, which equals to the sum of weights of its tasks.

- After generationg the macrotasks, the unbroken macrotasks along with the tasks not belonging to any unbroken macrotasks (i.e., the tasks that either do not share any resource or they belong to a broken macrotask) are ordered in a single list in non-increasing order of their weights. We denote this list the *mixed list*.

In the both rounds the strategy of task allocation depends on attraction between tasks. In the partitioning framework in Section 3.2 co-allocation of tasks is based on a cost function. In our blocking-aware algorithm we denote the function *attraction function* which has the same role in partitioning tasks. The attraction of task τ_k to a task τ_i is defined based on the potential remote blocking overhead that task τ_k can introduce to task τ_i if they are allocated onto different processors. We represent the attraction of task τ_k to task τ_i as $v_{i,k}$:

$$v_{i,k} = \begin{cases} \mathrm{NC}_{i,k}\beta_{i,k}\left\lceil \frac{T_i}{T_k} \right\rceil & \rho_i < \rho_k; \\ \mathrm{NC}_i\beta_{i,k} & \rho_i \geq \rho_k \end{cases} \tag{3.2}$$

The rationale behind the attraction function is to allocate the tasks which may remotely block a task, τ_i, to the same processor as of τ_i (in order of the amount of remote blocking overhead) as far as possible.

The definition of weight (Equation 3.1) and attraction function (Equation 3.2) are heuristics to guide the algorithm under MPCP. These function may differ under other synchronization protocols, e.g., MSRP, which have different remote blocking terms.

After the basic steps the algorithm continues with the rounds:

First Round The following steps are repeated within the first round until all tasks are allocated to processors (partitions):

- All processors are ordered in non-increasing order of their size (utilization).
- The object (a task or an unbroken macrotask) at the top of the mixed list is picked to be allocated.

  (i) If the object is a task and it does not belong to any broken macrotask it will be allocated onto the first processor that fits it (all processors are schedulable), beginning from the top of the ordered processor list. If none of the processors can fit the task a new processor is added to the list and the task is allocated onto it.

  (ii) If the object is an unbroken macrotask, all its tasks will be allocated onto the first processor that fits them (all processors can successfully be scheduled). If none of the processors can fit the tasks (at least one processor becomes unschedulable), they will be allocated onto a new processor.

  (iii) If the object is a task that belongs to a broken macrotask, the algorithm orders the tasks (those that are not allocated yet) within the macrotask in non-increasing order of attraction to the task based on Equation 3.2. We denote this list as *attraction list* of the task. The task itself will be on the top of its attraction list. Although creation of a attraction list begins from a task, in continuation tasks are added to the list that are most attracted to all of the tasks in the list, i.e., the sum of its attraction to the tasks in the list is maximized. The best processor for allocation which is the processor that fits the most tasks from the attraction list is selected, beginning from the top of the list. If none of the existing processors can fit any of the tasks, a new processor is added and as many tasks as possible from the attraction list are allocated to the processor. However, if the new processor cannot fit any task from the attraction list, i.e., at least one of the processors become unschedulable, the first round fails and the algorithm moves to the second round.

Second Round The following steps are repeated until all tasks are allocated to processors:

- The object at the top of the mixed list is picked.

  (i) If the object is a task and it does not belong to any broken macrotask, this step is performed the same way as in the first round.

  (ii) If the object is an unbroken macrotask, in this step the algorithm performs the same way as in the first round.

  (iii) If the object is a task that belongs to a broken macrotask, the ordered list of processors is a concatenation of two ordered lists of processors. The top list contains the processors that include some tasks from the macrotask of the picked task; this list is ordered in non-increasing order of processors' attraction to the task based on Equation 3.2, i.e., the processor which has the greatest sum of attractions of its tasks to the picked task is the most attracted processor to the task. The second list of processors is the list of the processors that do not contain any task from the macrotask of the picked task and are ordered in non-increasing order of their utilization. The picked task will be allocated onto the first processor from the processor list that will fit it. The task will be allocated to a new processor if none of the existing ones can fit it. The second round of the algorithm fails if allocating the task to the new processor makes at least one of the processors unschedulabe.

If both rounds fail to schedule a task set the algorithm fails. If one of the rounds fails the result will be the output of the other round. Finally, if both rounds succeed to schedule the task set, the one with less partitions (processors) will be the output of the algorithm.

3.3.2 Synchronization-Aware Algorithm (SPA)

In this section we present the partitioning algorithm originally proposed by Lakshmanan et al. in [12].

- Similar to BPA, the macrotasks are generated (in [12], macrotasks are denoted as bundles). A number of processors (enough processors that fit the total utilization of the task set, i.e., $\lceil u_i \rceil$) are added.
- The utilization of macrotasks and tasks are considered as their size and all the macrotasks together with all other tasks are ordered in a list in non-increasing order of their utilization. The algorithm attempts to allocate each macrotask onto a processor. Without adding any new processor, all macrotasks and tasks that fit are allocated onto the processors and the macrotasks that cannot fit are put aside. After each allocation, the processors are ordered in their non-increasing order of utilization.
- The remaining macrotasks are ordered in the order of the cost of breaking them. The cost of breaking a macrotask is defined based on the estimated cost (blocking time) introduced into the tasks by transforming a local resource into a global resource (i.e., the tasks sharing the resource are allocated to different processors). The estimated cost of transforming a local resource R_q into a global resource is defined as follows.

  $$\mathrm{Cost}(R_q) = \text{Global Overhead} - \text{Local Discount} \tag{3.3}$$

  The Global Overhead is calculated as follows.

  $$\text{Global Overhead} = \max(|Cs_q|)/\min_{\forall \tau_i}\{\rho_i\} \tag{3.4}$$

  where $\max(|Cs_q|)$ is the length of longest critical section accessing R_q.

  And the Local Discount is defined as follows.

  $$\text{Local Discount} = \max_{\forall \tau_i \text{ accessing } R_q} (\max(|Cs_{i,q}|)/\rho_i) \tag{3.5}$$

  where $\max(|Cs_{i,q}|)$ is the length of longest critical section of τ_i accessing R_q.

  The cost of breaking any macrotask, mTask_k, is calculated as the maximum of blocking overhead caused by transforming its accessed resources into global resources.

$$\mathrm{Cost}(\mathrm{mTask}_k) = \sum_{\forall R_q \text{ accessed by } \mathrm{mTask}_k} \mathrm{Cost}(R_q) \tag{3.6}$$

- The macrotask with minimum breaking cost is picked and is broken in two pieces such that the size of one piece is as close as the largest utilization available among processors. This means, tasks within the selected macrotask are ordered in decreasing order of their size (utilization) and the tasks from the ordered list are added to the processor with the largest available utilization as far as possible. In this way, the macrotask has been broken in two pieces; (1) the one including the tasks allocated to the processor and (2) the tasks that could not fit in the processor. If the fitting is not possible a new processor is added and the whole algorithm is repeated again.

The SPA algorithm does not consider any blocking parameters while it allocates the current task to a processor, but only its utilization, i.e., the tasks are ordered in order of their utilization only. The BPA, on the other hand, assigns a heuristic weight (Equation 3.1) which besides the utilization includes the blocking parameters as well. Another issue is that no relationship (e.g., based on blocking parameters) among individual tasks within a macrotask is considered which (as in the BPA) could help to allocate tasks from a broken bundle to appropriate processors to decreases the blocking times. The attraction function (Equation 3.2) facilitates the BPA, to allocate the most attracted tasks from the current task's broken macrotask, on a processor. As the experimental results in [30] show, considering these issues can improve the partitioning significantly.

3.3.3 Implementation

As the scheduling and synchronization protocols together with partitioning heuristics are being developed in the research community, the industry needs to evaluate the different methods to choose appropriate methods for migrating to multi-core platforms. This arises the need for developing tools to facilitate investigation and evaluation of different approaches and compare them to each other according to different factors.

We have developed a tool for evaluation of different scheduling and synchronization protocols as well as different partitioning algorithms. The output of the tool includes different information and graphs to facilitate evaluation and comparison of different approaches. Furthermore the tool can assist practitioners (given a scheduling and synchronization protocol as well as a partitioning heuristic) to find an appropriate solution for distributing a task set onto the processors of a multi-core platform.

We have implemented our blocking-aware partitioning algorithm together with the algorithm proposed in [12] and added both approaches to the tool. The tool has been developed in a modular manner to which any new scheduling and synchronization algorithms as well as any new partitioning heuristic can easily be added. However, the focus of the tool is partitioned scheduling and synchronization approaches as well as partitioning heuristics while extending the tool to global scheduling methods remains as a future work.

3.4 Summary

In this chapter we presented an overview of the contributions of the thesis, which is the development of a partitioning framework, a partitioning heuristic, and a tool to evaluate different scheduling, synchronization and partitioning algorithms for multi-cores. The target multi-core architecture is identical, unit capacity, shared memory multi-core platforms. Regarding the scheduling algorithm, the focus in this thesis is partitioned scheduling, and finally concerning the synchronization protocol, we have focused on MPCP.

In the framework, a task set is partitioned on a multi-core in a way to maximize the desired system performance. Tasks are grouped together based on timing constraints, resource sharing constraints, and task preferences. In the partitioning process, the framework uses a cost function to calculate the cost of allocating a task to a partition.

The heuristic algorithm for partitioning specifically focuses on decreasing the blocking overhead when allocating a task set on a multi-core platform. We have developed and implemented the blocking-aware heuristic algorithm together with a similar existing algorithm and a reference blocking-agnostic algorithm.

We have implemented a tool to evaluate different techniques of partitioned scheduling and synchronization, as well as task allocation algorithms. The tool has been developed in a modular manner to which new protocols and algorithms can easily be add in the future.

Chapter 4

Conclusions

4.1 Summary

In this thesis we have pointed out the increasing interest in multi-core architectures, and we have explained some of the challenges regarding migrating to these platforms. We have briefly discussed the existing scheduling approaches, e.g., partitioned and global scheduling. We have also presented an overview of the existing synchronization protocols for lock-based resource sharing on multiprocessor platforms, e.g., MPCP and FMLP.

We have proposed a general partitioning framework for distribution of a task set on a multi-core platform, which includes a heuristic algorithm that extends a bin-packing algorithm with a cost function based on task constraints and preferences.

We have also proposed a heuristic blocking-aware partitioning algorithm which extends a bin-packing algorithm with synchronization factors. The algorithm allocates a task set onto the processors of an identical multi-core platform. The objective of the algorithm is to decrease blocking times of tasks by means of allocating the tasks that directly or indirectly share resources onto the same processors as far as possible. This generally increases shedulability of a task set and can lead to fewer required processors compared to a blocking-agnostic bin-packing algorithm. Since in practice most systems use fixed priority scheduling protocols, we have developed our algorithm under MPCP, the only existing synchronization protocol for multiprocessors (multi-cores) which works under fixed priority scheduling [6]. However, this protocol introduces large amounts of blocking time overheads especially when the critical sections (in which a task accesses to resources) are relatively long and the accessing ratio to the resources is high.

Finally, we have implemented a tool to facilitate evaluation and comparison of different multiprocessor scheduling and synchronization approaches as well as different partitioning heuristics.

4.2 Future Work

In the future we plan to work further on the resource sharing issue in the multicore domains and we will investigate the possibility of improvement of the existing protocols as well as development of new approaches.

One future work will be to extend our partitioning algorithm to other synchronization protocols, e.g., MSRP and FMLP for partitioned scheduling. Another interesting future work is to apply our approach to real systems, and to study the performance gained by the algorithm on these systems.

In the domain of multiprocessor scheduling and synchronization our future work also includes investigating global and hierarchical scheduling protocols and appropriate synchronization protocols.

Looking at the tool that we have developed, the focus of the tool is currently partitioned scheduling approaches, and extending the tool to global scheduling protocols is another interesting future work.

Chapter 5

Overview of Papers

5.1 Paper A

Farhang Nemati, Moris Behnam, Thomas Nolte. *Efficiently Migrating Real-Time Systems to Multi-Cores*. In 14th IEEE International Conference on Emerging Techonologies and Factory (ETFA'09), pages 1205-1212, September, 2009.

Summary Power consumption and thermal problems limit a further increase of speed in single-core processors. Multi-core architectures have therefore received significant interest. However, a shift to multi-core processors is a big challenge for developers of embedded real-time systems, especially considering existing ”legacy” systems which have been developed with uniprocessor assumptions. These systems have been developed and maintained by many developers over many years, and cannot easily be replaced due to the huge development investments they represent. An important issue while migrating to multi-cores is how to distribute tasks among cores to increase performance offered by the multi-core platform. In this paper we propose a partitioning algorithm to efficiently distribute legacy system tasks along with newly developed ones onto different cores. The target of the partitioning is increasing system performance while ensuring correctness.

My contribution The basic idea of this paper was suggested by Farhang Nemati. Farhang was the main driver in writing and finalization of the paper.

5.2 Paper B

Farhang Nemati, Thomas Nolte and Moris Behnam. *Blocking-Aware Partitioning for Multiprocessors*. Technical Report, MRTC (Mälardalen Real-Time Research Centre), Mälardalen University, March, 2010.

Summary In the multi-core and multiprocessor domain there are two scheduling approaches, global and partitioned scheduling. Under global scheduling each task can execute on any processor while under partitioned scheduling tasks are allocated to processors and migration of tasks among processors is not allowed. Under global scheduling the higher utilization bound can be achieved, but in practice the overheads of migrating tasks is high. On the other hand, besides simplicity and efficiency of partitioned scheduling protocols, existing scheduling and synchronization methods developed for uniprocessor platforms can more easily be extended to partitioned scheduling. This also simplifies migration of existing systems to multi-cores. An important issue related to partitioned scheduling is how to distribute tasks among processors/cores to increase performance offered by the platform. However, existing methods mostly assume independent tasks while in practice a typical real-time system contains tasks that share resources and they may block each other. In this paper we propose a blocking-aware partitioning algorithm to distribute tasks onto different processors. The proposed algorithm allocates a task set onto processors in a way that blocking times of tasks are decreased. This reduces the total utilization which has the potential to decrease the total number of needed processors/cores.

My contribution The idea of this paper was suggested by Farhang Nemati. Farhang was the main driver in writing the paper and he was responsible for implementation and evaluation of the algorithm proposed in the paper.

5.3 Paper C

Farhang Nemati, Thomas Nolte and Moris Behnam. *Partitioning Real-Time Systems on Multiprocessors with Shared Resources*. In submission.

Summary There are two main approaches to task scheduling on multiprocessor/multi-core platforms; 1) global scheduling, under which migration of tasks among processors is allowed, and 2) partitioned scheduling under which tasks are allocated onto processors and task migration is not allowed. Under global scheduling a higher utilization bound can be achieved, but in practice the overheads of migrating tasks is high. On the other hand under partitioned scheduling, besides simplicity and efficiency, existing scheduling and synchronization methods developed for uniprocessor platforms can more easily be extended to partitioned scheduling. However the partitioned scheduling protocols suffer from the problem of partitioning tasks among processors/cores which is a bin-packing problem. Therefore, several heuristic algorithms have been developed for partitioning a task set on multiprocessor platforms. However, these algorithms typically assume independent tasks while in practice real-time systems often contain tasks that share resources and hence may block each other.

In this paper we propose a blocking-aware partitioning algorithm which allocates a task set onto processors in a way that the overall amount of blocking times of tasks are decreased. The algorithm reduces the total utilization which, in turn, has the potential to decrease the total number of required processors (cores). In this paper we evaluate our algorithm and compare it with an existing similar algorithm. The comparison criteria includes both number of schedulable systems as well as processor reduction performance.

My contribution Farhang Nemati was the main driver in writing the paper and he was responsible for further evaluation of the algorithm. He was also responsible for implementing an algorithm similar to the algorithm proposed in Paper B, and comparing the two algorithms.

5.4 Paper D

Farhang Nemati, Thomas Nolte. *A Flexible Tool for Evaluating Scheduling, Synchronization and Partitioning Algorithms on Multiprocessors*. In submission.

Summary Multi-core platforms seem to be the way towards increasing performance of processors. Single-chip multiprocessors (multi-cores) are today the dominating technology for desktop computing. As the multi-cores are becoming the defacto processors, the need for new scheduling and resource sharing protocols has arisen. There are two major types of scheduling under multiprocessor/multi-core platforms. Global scheduling, under which migration of tasks among processors is allowed, and partitioned scheduling under which tasks are allocated onto processors and task migration is not allowed. The partitioned scheduling protocols suffer from the problem of partitioning tasks among processors/cores, which is a bin-packing problem. Heuristic algorithms have been developed for partitioning a task set on multiprocessor platforms. However, taking such technology to an industrial setting, it needs to be evaluated such that appropriate scheduling, synchronization and partitioning algorithms are selected.

In this paper we present our work on a tool for investigation and evaluation of different approaches to scheduling, synchronization and partitioning on multi-core platforms. Our tool allows for comparison of different approaches with respect to a number of parameters such as number of schedulable systems and number of processors required for scheduling. The output of the tool includes a set of information and graphs to facilitate evaluation and comparison of different approaches.

My contribution Farhang Nemati was the main driver in writing the paper and he was responsible for implementation of the tool.

Bibliography

[1] T. Baker. Stack-based scheduling of real-time processes. *Journal of Real-Time Systems*, 3(1):67–99, 1991.

[2] T. Baker. A comparison of global and partitioned EDF schedulability test for multiprocessors. Technical report, 2005.

[3] J. Carpenter, S. Funk, P. Holman, A. Srinivasan, J. Anderson, and S. Baruah. A categorization of real-time multiprocessor scheduling problems and algorithms. In *Handbook on Scheduling Algorithms, Methods, and Models*. Chapman Hall/CRC, Boca, 2004.

[4] U. Devi. Soft real-time scheduling on multiprocessors. In *PhD thesis, available at www.cs.unc.edu/˜anderson/diss/devidiss.pdf*, 2006.

[5] J. M. López, J. L. Díaz, and D. F. García. Utilization bounds for EDF scheduling on real-time multiprocessor systems. *Journal of Real-Time Systems*, 28(1):39–68, 2004.

[6] R. Rajkumar. *Synchronization in Real-Time Systems: A Priority Inheritance Approach*. Kluwer Academic Publishers, 1991.

[7] P. Gai, G. Lipari, and M. D. Natale. Minimizing memory utilization of real-time task sets in single and multi-processor systems-on-a-chip. In *proceedings of 22nd IEEE Real-Time Systems Symposium (RTSS'01)*, pages 73–83, 2001.

[8] P. Gai, M. Di Natale, G. Lipari, A. Ferrari, C. Gabellini, and P. Marceca. A comparison of MPCP and MSRP when sharing resources in the janus multiple processor on a chip platform. In *proceedings of 9th IEEE Real-Time And Embedded Technology Application Symposium (RTAS'03)*, pages 189–198, 2003.

[9] L. Sha, R. Rajkumar, and J. P. Lehoczky. Priority inheritance protocols: An approach to real-time synchronization. *Journal of IEEE Transactions on Computers*, 39(9):1175–1185, 1990.

[10] A. Block, H. Leontyev, B. Brandenburg, and J. Anderson. A flexible real-time locking protocol for multiprocessors. In *proceedings of 13th IEEE Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA'07)*, pages 47–56, 2007.

[11] A. Easwaran and B. Andersson. Resource sharing in global fixed-priority preemptive multiprocessor scheduling. In *proceedings of 30th IEEE Real-Time Systems Symposium (RTSS'09)*, pages 377–386, 2009.

[12] K. Lakshmanan, D. de Niz, and R. Rajkumar. Coordinated task scheduling, allocation and synchronization on multiprocessors. In *proceedings of 30th IEEE Real-Time Systems Symposium (RTSS'09)*, pages 469–478, 2009.

[13] D. S. Johnson. *Near-optimal bin packing algorithms*. Massachusetts Institute of Technology, 1973.

[14] J. A. Stankovic and K. Ramamritham, editors. *Tutorial: hard real-time systems*. 1989.

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

[16] S. K. Baruah, N. K. Cohen, C. G. Plaxton, and D. A. Varvel. Proportionate progress: A notion of fairness in resource allocation. *Journal of Algorithmica*, 15(6):600–625, 1996.

[17] J. Anderson, P. Holman, and A. Srinivasan. Fair scheduling of real-time tasks on multiprocessors. In *Handbook of Scheduling*, 2005.

[18] J. M. Calandrino, J. H. Anderson, and D. P. Baumberger. A hybrid real-time scheduling approach for large-scale multicore platforms. In *proceedings of 19th IEEE Euromicro Conference on Real-time Systems (ECRTS'07)*, pages 247–258, 2007.

[19] I. Shin, A. Easwaran, and I. Lee. Hierarchical scheduling framework for virtual clustering of multiprocessors. In *proceedings of 20th IEEE*

Euromicro Conference on Real-time Systems (ECRTS'08), pages 181–190, 2008.

[20] U. Devi, H. Leontyev, and J. Anderson. Efficient synchronization under global EDF scheduling on multiprocessors. In *proceedings of 18th IEEE Euromicro Conference on Real-time Systems (ECRTS'06)*, pages 75–84, 2006.

[21] B. Brandenburg, J. Calandrino, A. Block, H. Leontyev, and J. Anderson. Synchronization on multiprocessors: To block or not to block, to suspend or spin? In *proceedings of 14th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'08)*, pages 342–353, 2008.

[22] B. B. Brandenburg and J. H. Anderson. A comparison of the M-PCP , D-PCP , and FMLP on LITMUSRT. In *proceedings of the 12th International Conference on Principles of Distributed Systems (OPODIS'08)*, pages 105–124, 2008.

[23] A. Burns. Preemptive priority based scheduling: An appropriate engineering approach. In *Principles of Real-Time Systems*, pages 225–248. Prentice Hall, 1994.

[24] M. Rajagopalan, B. T. Lewis, and T. A. Anderson. Thread scheduling for multi-core platforms. In *proceedings of the 11th Workshop on Hot Topics in Operating Systems (HotOS'07)*, 2007.

[25] J. Lindhult. Operational semantics for plex a basis for safe parallelization. In *Licentiate thesis, Mälardalen University Press*, 2008.

[26] A. Prayati, C. Wong, P. Marchal, F. Catthoor, H. de Man, N. Cossement, R. Lauwereins, D. Verkest, and A. Birbas. Task concurrency management experiment for power-efficient speed-up of embedded mpeg4 im1 player. In *proceedings of International Conference on Parallel Processing Workshops (ICPPW'00)*, pages 453–460, 2000.

[27] D. de Niz and R. Rajkumar. Partitioning bin-packing algorithms for distributed real-time systems. *Journal of Embedded Systems*, 2(3-4):196–208, 2006.

[28] Y. Liu, X. Zhang, H. Li, and D. Qian. Allocating tasks in multi-core processor based parallel systems. In *proceedings of Network and Parallel Computing Workshops, in conjunction with IFIP'07*, pages 748–753, 2007.

[29] S. Baruah and N. Fisher. The partitioned multiprocessor scheduling of sporadic task systems. In *proceedings of 26th IEEE Real-Time Systems Symposium (RTSS'05)*, pages 321–329, 2005.

[30] F. Nemati, T. Nolte, and M. Behnam. Partitioning real-time systems on multiprocessors with shared resources. In *In submition*, 2010.