

Non-preemptive multiprocessor static scheduling for systems with precedence and strict periodicity constraints

Omar Kermia¹, Liliana Cucu², Yves Sorel¹

(1)INRIA Rocquencourt,
BP 105 - 78153 Le Chesnay Cedex, France
Phone: +33 1 39 63 52 60 - Fax: +33 1 39 63 51 93
omar.kermia@inria.fr, yves.sorel@inria.fr
(2)IPP Hurray Research Group,
Polytechnic Institute of Porto, Portugal
Phone: +351 22 8340502 - Fax:+351 228340529
lcucu@dei.isep.ipp.pt

Abstract

In this paper we propose a greedy heuristic to solve the non preemptive multiprocessor static scheduling problem with precedence and strict periodicity constraints. The system of periodic tasks is described by a graph where dependent tasks are connected by precedence constraints. First, each task is repeated within the LCM of all periods of tasks (hyper-period) allowing to unroll the graph over the hyper-period. We propose an algorithm adding some missing edges. Then, the tasks of the new graph are classified by their periods and assigned to the processors. Finally, we propose an extension of the SynDEx heuristic which already allocates and schedules tasks graph with precedence constraints, but tasks have the same period equal to the execution time of the complete graph. We apply this extended heuristic on the unrolled graph in order to allocate and schedule the periodic tasks onto the multiprocessor. We compute the complexity of the proposed heuristic, and perform a performance comparison which shows its effectiveness.

1 Introduction

Distributed real-time embedded applications found in domains such as avionic, automobile, autonomous robotics, telecommunication lead to non preemptive multiprocessor static scheduling problem with precedence and strict periodicity constraints. These applications are of crucial importance because if all the constraints are not satisfied this may have disastrous consequences.

The multiprocessor periodic scheduling problem was discussed by Liu in 1969 [1]. Given a set of n tasks and m resources, where each task x has rational weight $x:w = x:e=x:p$ ($x:e$ and $x:p$ referred to execution requirement and period), $0 < x:w < 1$, a periodic schedule according to him is one that allocates a resource to a task x exactly $x:e$ time unit of slots in each interval $[x:p:k, x:p:(k+1)]$ for all $k \in \mathbb{N}$. Scheduling decisions may be made only at integral times and a task may use either zero or one resources at a time. To solve this problem Balaiah [2] proposed the fair scheduling that considers a relaxed version of the Liu's problem in which tasks are not restricted to using zero or one resource at a time (resource sharing). But this approach can not be applied to our problem because, contrary to Liu's problem, we have a machines instead of resources and in this case resource sharing would mean that one machine can execute more than one task on the same time. Xu [3] considers task sets where individual tasks are further divided into subtasks with precedence relation in order to give more parallelism within tasks, and in this idea Shi-azi [4] developed an algorithm of static multiprocessor scheduling of periodic tasks which are independent and their subtasks are related. [5] gives a lot of known results to solve the periodic schedule but he considers that all the tasks are at the same period. In this paper we deal with a non preemptive multiprocessor scheduling with precedence and multiple strict periodicities constraint and we propose an algorithm which take into account the interprocessor communications.

This problem is solved here by a partitioning method and the partitioning problem is analogous to the bin-packing problem known NP-hard [6]. Thus, there are no optimal algorithms which solve it, whereas heuristics techniques are very fast despite their sub-optimal results. Because embedded distributed real-time systems involve rapid prototyping of useful solutions, we choose greedy heuristics which are very fast and produce good results.

2 Model and heuristic principles

[7] gives a model based on graph theory o

schedules operations graph with precedence constraints in the case where operations have the same period equal to the execution time of the entire graph.

At each main step of the heuristic, one operation is selected from the set of schedulable operations (an operation becomes schedulable when all its predecessors in the graph are already scheduled). In order to select an operation, we first seek, for each schedulable operation, its best processor, i.e. which minimizes the cost function. Then, among these operations, we select the most critical one, i.e. which maximizes the cost function, and schedule it onto its best processor where it is finally allocated to. This operation is removed from the set of schedulable operations and all its successors that become schedulable are added to this set. The next main step is ready to begin with this new set, and the heuristic stops when it is empty.

Here, we extend the SynDx heuristic in order to satisfy, in addition to precedences, several strict periodicities constraints. The extended heuristic is applied onto the unrolled graph instead of the initial one. At each main step, like in the original heuristic, one operation is selected from the set of schedulable operations. This set contains two kinds of operations: those belonging to an assigned class and those belonging to a non-assigned class. If the operation belongs to an assigned class, it is considered as allocated to this processor, and the cost function is computed relatively to this processor. If the operation belongs to a non-assigned class, the cost function is computed relatively to all the processors in order to find its best processor. Once every operation (assigned and non-assigned) has found its best processor, we select the most critical one, which becomes scheduled and allocated. Then, the heuristic proceeds like the original one. Moreover, in order to guarantee that the execution of the repeated operations satisfy the periodicity constraints, we verify at each main step that a schedulable operation fulfills two conditions given below.

Let B a schedulable operation. Θ is the set of the periodic operations already scheduled whose successors, corresponding to the same repetition, was not yet scheduled. We denote by O_1 the last operation scheduled on the processor P_1 and $S(X)$ the start time of operation X .

► First condition: if the scheduling of the operation B on the processor P_1 does not need data transfer, it is necessary that: $\forall A \in \Theta$ and $T(A)$ its period, $S(O_1) + C(O_1) + C(B) \leq S(A) + T(A)$.

► Second condition: if the scheduling of the operation B on the processor P_1 needs data transfer (from the operation O_2 scheduled on the processor P_2 with Com the execution time of this inter-processor communication), it is necessary that: $\forall A \in \Theta$ and $T(A)$ is its period, $S(O_1) + C(O_1) + C(B) \leq S(A) + T(A)$ and $S(O_2) + C(O_2) + Com + C(B) \leq S(A) + T(A)$.

3.4 Example

In order to illustrate how the proposed heuristic progresses we apply it to the graph of figure 1 with the same multiprocessor (two processors connected by a communication medium). The periods of the operations are: $T(A)=4$, $T(B)=T(C)=5$ and $T(D)=20$. The execution times are: $C(A)=C(B)=C(C)=2$ and $C(D)=5$ of the operations and $Com=2$ of the inter-processor communication.

• Graph unrolling: the LCM of the periods is $P=20$, $n_1=20/4=5$, $n_2=20/5=4$, $n_3=20/5=4$, $n_4=20/20=1$ are respectively the repetition numbers of the graph operations A , B , C and D . The edge weights between the repetitions of A and those of B , the same number between the repetitions of A and those of C , one edge between B and D , and one edge between C and D . The resulting graph is depicted on figure 3.

• Classification and assignment: the operations are classified according to their periods in three classes: $c_1=\{A\}$, $c_2=\{B;C\}$, $c_3=\{D\}$. Because the multiprocessor architecture is composed of only two processors, we assign the first class to one processor and the second class to the other one, and the third class is not assigned.

• Execution of the extended SynDx heuristic: by applying this heuristic on the unrolled graph (figure 3) we obtain the allocation and scheduling given on figure 4.

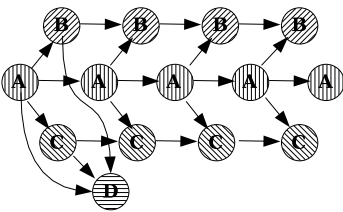


Figure 3: The unrolled dependence graph

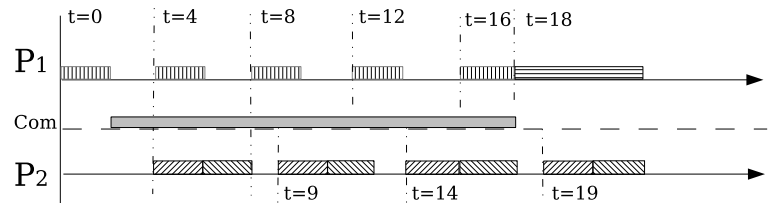


Figure 4: Allocation and scheduling

4 Heuristic complexity and performance

In order to compute the complexity of the proposed heuristic, we start by reminding the complexity of the SynDx heuristic which is equal to $N^2 \times M=2$ with M the processors number, N the number of operations in the dependence graph, and compute the number of unrolled graph operations which is equal to $N_{unrolling} = \sum_{i=1}^N P= T(O_i)$ with O_i an operation of the graph and $T(O_i)$ its period and P the hyper-period. Consequently, the complexity of the proposed heuristic is $N_{unrolling}^2 \times M=2$.

We compared the execution times of our heuristic to the one given in [8] on a similar problem in terms of number of operations and processors. The results are summarized in the table given below.

	Number of operations	Number of processors	Execution time (sec)
Simulated annealing heuristic	624	9	5572
Proposed heuristic	681	9	32.61

5 Conclusion and further research

The paper proposes an heuristic to allocate and schedule systems of operations with precedence and periodicity constraints onto multiprocessors, and gives some complexity and performance results showing that it can be used in realistic embedded real-time applications. We plan to investigate other possibilities of adding edges in the unrolled graph, and propose a flexible method of adding these edges which deals with a large kind of applications and adapts to the users needs.

References

- [1] C. L. Liu. Scheduling algorithms of multiprocessors in a hard real-time environment. In *JPL Space Programs Summary 37-60*, volume 2, pages 28–37, November 1969.
- [2] S. K. Bahuah. Fairness in periodic real-time scheduling. In *RTSS '95: Proceedings of the 16th IEEE Real-Time Systems Symposium (RTSS '95)*, page 200, Washington, DC, USA, 1995. IEEE Computer Society.
- [3] J. Xu. Multiprocessor scheduling of processes with release times, deadlines, precedence, and exclusion relations. *IEEE Transactions on Parallel Distributed Systems*, 19(2), February 1993.
- [4] S. Rungten and A. Shizazi. Static multiprocessor scheduling of periodic real-time tasks with precedence constraints and communication costs. In *Proceeding of the 28th Annual Hawaii International conference on system Sciences*, Department of Computer Science and Engineering, University of Texas, 1995.
- [5] J. Leung. Handbook of scheduling. pages 7–4.7–5. Published by CRC Press, Boca Raton, FL, USA, May 2004.
- [6] M. R. Garey and D. S. Johnson. *Computers and Intractability. A guide to the theory of NP-completeness*. W. H. Freeman, 1979.
- [7] L. Cucu and Y. Soel. Schedulability condition of systems with precedence and periodicity constraints without preemption. In *Proceedings of 11th Real-Time Systems Conference, RTS'03*, Paris, March 2003.
- [8] S-T. Cheng and A. Agarwala. Allocation and scheduling of real-time periodic tasks with relative timing constraints. In *Technical Report CS-TR-3402*, Computer science, University of Maryland, College Park, January 1995.
- [9] K. Ramamitham. Allocation and scheduling of precedence-related periodic tasks. In *IEEE Transactions on Parallel and Distributed Systems*, volume 6, University of Massachusetts, April 1995.
- [10] T. Gandpie, C. Lavaenne, and Y. Soel. Optimized rapid prototyping of real-time embedded heterogeneous multiprocessors. In *Proceedings of 7th International Workshop on Hardware/Software Co-Design*, May 1999.