

Efficient Method for Periodic Task Scheduling with Storage Requirement Minimization

Karine DESCHINKEL, Sid-Ahmed-Ali TOUATI

University of Versailles Saint-Quentin-en-Yvelines, France

Abstract. In this paper, we study the general problem of one-dimensional periodic task scheduling under storage requirement, irrespective of machine constraints. We have already presented in [9] a theoretical framework that allows an optimal optimization of periodic storage requirement in a periodic schedule. This problem is used to optimize processor register usage in embedded systems. Our storage optimization problem being NP-complete [8], solving an exact integer linear programming formulation is too expensive in practice. In this article, we propose an efficient two-steps heuristic using model's properties that allows fast resolution times while providing nearly optimal results. This method includes the resolution of a integer linear program with a totally unimodular constraints matrix in first step, then the resolution of a linear assignment problem. Our solution has been implemented and included inside a compiler for embedded processors.

1 Introduction

This article addresses the problem of storage optimization in cyclic data dependence graphs (DDG), which is for instance applied in the practical problem of periodic register allocation for innermost loops on modern Instruction Level Parallelism (ILP) processors[10]. The massive introduction of ILP processors since the last two decades makes us re-think new ways of optimizing register/storage requirement in assembly codes before starting the instruction scheduling process under resource constraints. In such processors, instructions are executed in parallel thanks to the existence of multiple small computation units (adders, multipliers, load-store units, etc.). The exploitation of this new fine grain parallelism (at the assembly code level) asks to completely revisit the old classical problem of register allocation initially designed for sequential processors. Nowadays, register allocation has not only to minimize the storage requirement, but has also to take care of parallelism and total schedule time. In this research article, we do not assume any resource constraints (except storage requirement); Our aim is to analyze the trade-off between memory (register pressure) and parallelism in a periodic task scheduling problem. Note that this problem is abstract enough to be considered in other scheduling disciplines that worry about conjoint storage and time optimization in repetitive tasks (manufacturing, transport, networking, etc.).

Existing techniques in this field usually apply a periodic instruction scheduling under resource constraints that is sensitive to register/storage requirement. Therefore a great amount of work tries to schedule the instructions of a loop (under resource and time constraints) such that the resulting code does not use more than R values simultaneously alive. Usually they look for a schedule that minimizes the storage requirement

under a fixed scheduling period while considering resource constraints [3, 4, 6, 1]. In this paper, we satisfy register/storage constraints early before instruction scheduling under resource constraints: we directly handle and modify the DDG in order to fix the storage requirement of any further subsequent periodic scheduling pass while taking care of not altering parallelism exploitation if possible. This idea uses the concept of reuse vector used for multi-dimensional scheduling [11, 12].

This article is on continuation on our previous work on register allocation [10]. In that paper, we showed that register allocation implies a loop unrolling. However, the general problem of storage optimization does not require such loop unrolling. So the current paper is an abstraction of our previous results on register optimization. Furthermore, it extends it with a new heuristic and experimental results.

Our article is organized as follows. Section 2 recalls our task model and notations already presented in [9]. Section 3 recalls the exact problem of optimal periodic scheduling under storage constraints with integer linear programming: our detailed results on the optimal resolution of this problem have been presented in [9]. Since the exact model is not practical (too expensive in terms of resolution time), our current article provides a new look by writing an efficient approximate method in Section 4, that we call *SIR-ALINA*. Before concluding, Section 5 presents the results of our experimental evaluation of *SIRALINA*, providing practical evidence of its efficiency.

2 Tasks Model

Our task model is similar to [2]. We consider a set of l generic tasks (instructions inside a program loop) T_0, \dots, T_{l-1} . Each task T_i should be executed n times, where n is the number of loop iterations. n is an unknown, unbounded, but finite integer. This means that each task T_i has n instances. The k^{th} occurrence of task T_i is noted $T\langle i, k \rangle$, which corresponds to task i executed at the k^{th} iteration of the loop, with $0 \leq k < n$.

The tasks (instructions) may be executed in parallel. Each task may produce a result that is read/consumed by other tasks. The considered loop contains some data dependencies represented with a graph G such that:

- V is the set of the generic tasks of the loop body, $V = \{T_0, \dots, T_{l-1}\}$.
- E is the set of edges representing precedence constraints (flow dependences or other serialization constraints). Any edge $e = (T_i, T_j) \in E$ has a latency $\delta(e) \in \mathbb{N}$ in terms of processor clock cycles and a distance $\lambda(e) \in \mathbb{N}$ in terms of number of loop iterations. The distance $\lambda(e)$ means that the edge $e = (T_i, T_j)$ is a dependence between the task $T\langle i, k \rangle$ and $T\langle j, k + \lambda(e) \rangle$ for any $k = 0, \dots, n - 1 - \lambda(e)$.

We make a difference between tasks and precedence constraints depending whether they refer to data to be stored into registers or not

1. V_R is the set of tasks producing data to be stored into registers.
2. E_R is the set of flow dependence edges through registers. An edge $e = (T_i, T_j) \in E_R$ means that the task $T\langle i, k \rangle$ produces a result stored into a register and read/consumed by $T\langle j, k + \lambda(e) \rangle$. The set of consumers (readers) of a generic task T_i is then the set:

$$Cons(T_i) = \{T_j \in V \mid e = (T_i, T_j) \in E_R\}$$

Figure 1 is an example of a data dependence graph (DDG) where bold circles represent V_R the set of generic tasks producing data to be stored into registers. Bold edges represent flow dependences (each sink of such edge reads/consumes the data produced by the source and stored in a register). Tasks that are not in bold circles are instructions that do not write into registers (write the data into memory or simply do not produce any data). Non-bold edges are other data or precedence constraints different from flow dependences. Every edge e in the DDG is labeled by the pair $(\delta(e), \lambda(e))$.

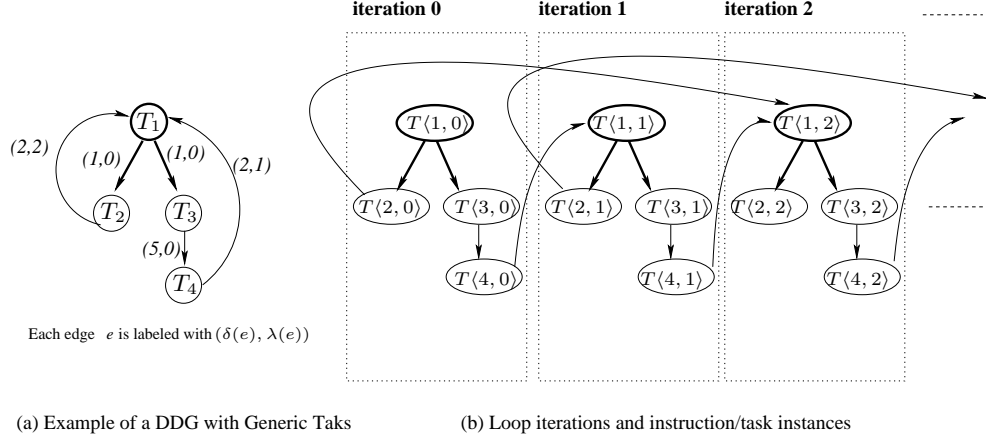


Fig. 1. Example of Data Dependence Graphs with Recurrent Tasks

In our generic processor model, we assume that the reading and writing from/into registers may be delayed from the starting time of task execution. Let assume $\sigma(T\langle i, k \rangle) \in \mathbb{N}$ as the starting execution time of task $T\langle i, k \rangle$. We thus define two delay functions δ_r and δ_w in which

$$\begin{aligned} \delta_w : V_R &\rightarrow \mathbb{N} \\ T_i &\mapsto \delta_w(T_i) \mid 0 \leq \delta_w(T_i) \\ &\text{the writing time of data produced by } T\langle i, k \rangle \text{ is } \sigma(T\langle i, k \rangle) + \delta_w(T_i) \\ \delta_r : V &\rightarrow \mathbb{N} \\ T_i &\mapsto \delta_r(T_i) \mid 0 \leq \delta_r(T_i) \\ &\text{the reading time of the data consumed by } T\langle i, k \rangle \text{ is } \sigma(T\langle i, k \rangle) + \delta_r(T_i) \end{aligned}$$

These two delays functions depend on the target processor and model almost all regular hardware architectures (VLIW, EPIC/IA64 and superscalar processors).

3 Exact Problem Formulation

This section recalls the exact integer linear model for solving the problem of Periodic Scheduling with Storage Minimisation (PSSM). It is built for a fixed desired period $p \in \mathbb{N}$. For more details on this problem, please refer to [9].

3.1 Basic Variables

- a schedule variable $\sigma_i \geq 0$ for each task $T_i \in V$, including σ_{K_i} for each killing node K_i . We assume a finite upper bound L for such schedule variables (L sufficiently large, $L = \sum_{e \in E} \delta(e)$); The schedule variables are integer. As our scheduling is periodic, we only consider the integer execution date of the first task occurrence $\sigma_i = T\langle i, 0 \rangle$ and the execution date of any other occurrence $T\langle i, k \rangle$ becomes equal to $\sigma(T\langle i, k \rangle) = \sigma_i + k * p$.
- a binary variables $\theta_{i,j}$ for each $(T_i, T_j) \in V_R^2$. It is set to 1 iff (T_i, T_j) is a reuse edge;
- a reuse distance $\mu_{i,j}$ for all $(T_i, T_j) \in V_R^2$; The reuse distance are nonnegative integer variables.

3.2 Linear Constraints

– Data dependences

The schedule must at least satisfy the precedence constraints defined by the DDG.

$$\forall e = (T_i, T_j) \in E : \sigma_j - \sigma_i \geq \delta(e) - p \times \lambda(e) \quad (1)$$

– Flow dependences

Each flow dependence $e = (T_i, T_j) \in E_R$ means that the task occurrence $T\langle j, k + \lambda(e) \rangle$ reads the data produced by $T\langle i, k \rangle$ at time $\sigma_j + \delta_r(T_j) + (\lambda(e) + k) \times p$. Then, we should schedule the killing node K_i of the task T_i after all T_i 's consumers.

$$\forall T_i \in V_R, \forall T_j \in Cons(T_i) | e = (T_i, T_j) \in E_R : \sigma_{K_i} \geq \sigma_j + \delta_r(T_j) + p \times \lambda(e) \quad (2)$$

– Storage dependences

There is a storage dependence between K_i and T_j if (T_i, T_j) is a reuse edge.

$$\forall (T_i, T_j) \in V_R^2 : \theta_{i,j} = 1 \implies \sigma_{K_i} - \delta_w(T_j) \leq \sigma_j + p \times \mu_{i,j}$$

This involvement can result in the following inequality :

$$\forall (T_i, T_j) \in V_R^2 : \sigma_j - \sigma_{K_i} + p \times \mu_{i,j} + M_1(1 - \theta_{i,j}) \geq -\delta_w(T_j) \quad (3)$$

where M_1 is an arbitrarily large constant.

If there is no register reuse between two tasks T_i and T_j , then $\theta_{i,j} = 0$ and the storage dependence distance $\mu_{i,j}$ must be set to 0.

$$\forall (T_i, T_j) \in V_R^2 : \mu_{i,j} \leq M_2 \theta_{i,j} \quad (4)$$

where M_2 is an arbitrarily large constant.

– Reuse Relations

The reuse relation must be a bijection from V_R to V_R . A register can be reused by one task and a task can reuse one released register:

$$\forall T_i \in V_R : \sum_{T_j \in V_R} \theta_{i,j} = 1 \quad (5)$$

$$\forall T_j \in V_R : \sum_{T_i \in V_R} \theta_{i,j} = 1 \quad (6)$$

3.3 Objective Function

As proved in [10], the storage requirement is equal to $\sum \mu_{i,j}$. In our periodic scheduling problem, we want to minimize the storage requirement: Minimize $z = \sum_{(T_i, T_j) \in V_R^2} \mu_{i,j}$

Using the above integer linear program to solve an NP-problem as PSSM is not efficient in practice. With a classical Branch and Bound method, we are only able to solve small instances (DDG sizes), in practice around 12 nodes. For this reason, we suggest to make use of the problem structure to propose an efficient heuristic as follows.

4 SIRALINA : A Two Steps Approximate Resolution Method

Our resolution strategy is based on the analysis of the model constraints. As the problem involves scheduling constraints and assignment constraints, and the reuse distances are the link between these two sets of constraints, we attempt to decompose the problem into two subproblems :

- A scheduling problem : to find a scheduling for which the potential reuse distances are as small as possible.
- An assignment problem : to select which pairs of tasks will share the same register.

4.1 Preliminaries

If edge $e = (T_i, T_j) \in V_R^2$ is a reuse edge, its reuse distance should satisfy the inequality given by 3, where $\theta_{ij} = 1$. This inequality gives a lower bound for each reuse distance. If $(T_i, T_j) \in V_R^2$ is a reuse edge (E_r denotes the set of reuse edge) then :

$$\forall (T_i, T_j) \in E_r : \mu_{i,j} \geq \frac{1}{p}(\sigma_{K_i} - \delta_w(T_j) - \sigma_j) \quad (7)$$

If $(T_i, T_j) \in V_R^2$ is not a reuse edge then $\mu_{ij} = 0$ according to the inequality 4.

$$\forall (T_i, T_j) \notin E_r : \mu_{i,j} = 0$$

The aggregation of constraint 7 for each reuse edge provides a lower bound of the objective function value. $z = \sum_{(T_i, T_j) \in V_R^2} \mu_{i,j} \geq \frac{1}{p}(\sum_{(T_i, T_j) \in E_r} \sigma_{K_i} - \delta_w(T_j) - \sigma_j)$ As the reuse relation is a bijection from V_R to V_R , the left sum of the inequality can be separated into two parts.

$$\begin{aligned} \sum_{(T_i, T_j) \in E_r} \sigma_{K_i} - \delta_w(T_j) - \sigma_j &= \sum_{i \in V_R} \sigma_{K_i} - \sum_{j \in V_R} (\delta_w(T_j) + \sigma_j) \\ &= \sum_{i \in V_R} \sigma_{K_i} - \sum_{j \in V_R} \sigma_j - \sum_{j \in V_R} \delta_w(T_j) \end{aligned}$$

We deduce from this inequality a lower bound for the number of required registers. In this context, it may be useful to find an appropriate scheduling for which this value is minimal. As $\sum_{j \in V_R} \delta_w(T_j)$ is a constant for the problem, we could ignore it in the

following optimization problem.

We consider *the scheduling problem (P)*:

$$\begin{cases} \min \sum_{i \in V_R} \sigma_{K_i} - \sum_{j \in V_R} \sigma_j \\ \text{subject to :} \\ \sigma_j - \sigma_i \geq \delta(e) - p \times \lambda(e), \quad \forall e = (T_i, T_j) \in E \\ \sigma_{K_i} - \sigma_j \geq \delta_r(T_j) + p \times \lambda(e), \quad \forall T_i \in V_R, \forall T_j \in \text{Cons}(T_i) | e = (T_i, T_j) \in E_R \end{cases} \quad (8)$$

As the constraints matrix of the integer linear program of System 8 is totally unimodular, *i.e.*, the determinant of each square sub-matrix is equal to 0 or to ± 1 , we can use polynomial algorithms to solve this problem [7]. This would allow us to consider huge DDG. The resolution of problem (P) by a simplex method will provide optimal values σ_i^* for each task $T_i \in V_R$ and the optimal values $\sigma_{K_i}^*$ for each killing node K_i .

Once the scheduling variables have been fixed, the minimal value of each potential reuse distance would be equal to $\overline{\mu_{ij}} = \lceil \frac{\sigma_{K_i}^* - \delta_w(T_j) - \sigma_j^*}{p} \rceil$ according to 7. Knowing the reuse distance values $\overline{\mu_{ij}}$ if T_j reuses the register freed by T_i , the storage allocation which consists of choosing which task reuses which released register can be modeled as a linear assignment problem.

We consider *the linear assignment problem (A)*:

$$\begin{cases} \min \sum_{(T_i, T_j) \in V_R^2} \overline{\mu_{i,j}} \theta_{ij} \\ \text{Subject to} \\ \sum_{T_j \in V_R} \theta_{i,j} = 1, \quad \forall T_i \in V_R \\ \sum_{T_i \in V_R} \theta_{i,j} = 1, \quad \forall T_j \in V_R \\ \theta_{ij} \in \{0, 1\} \end{cases} \quad (9)$$

where $\overline{\mu_{i,j}}$ is a fixed value for each edge $e = (T_i, T_j) \in V_R^2$.

4.2 Heuristic

We suggest to solve the problem with the following heuristic :

- Solve the problem (P) to deduce the optimal values σ_i^* for each task $T_i \in V_R$ and the optimal values $\sigma_{K_i}^*$ for each killing node K_i ,
- Compute the cost $\overline{\mu_{ij}} = \max \left(\lceil \frac{\sigma_{K_i}^* - \delta_w(T_j) - \sigma_j^*}{p} \rceil \right)$ for each edge $e = (T_i, T_j) \in V_R^2$,
- Solve the linear assignment problem (A) with the Hungarian algorithm [5] which solves assignment problems in polynomial time ($O(n^3)$) to deduce the optimal values $\theta_{i,j}^*$,
- If $\theta_{i,j}^* = 1$ for the edge $e = (T_i, T_j) \in V_R^2$, then (T_i, T_j) is a reuse edge and the reuse distance is equal to $\overline{\mu_{ij}}$.

5 Experiments

We now present the results obtained on several DDG extracted from many well known benchmarks (Spec95, whetstone, livermore, lin-ddot, DSP filters, etc.). The data dependence graphs of all these loops are present in [8]. The small test instances have 2 nodes and 2 edges, the large instances have multiples hundreds of nodes and edges. We use the ILOG CPLEX 10.2 to solve the integer linear program. The experiment was run on PC under linux, equipped with a Pentium IV 2.13 Ghz processor, and 2 Giga bytes of memory. In practice, the optimal method [9] can solve small instances, around 10 nodes. As far as we know about our problem, it is still NP-complete even for DDG chains and trees [8]: we do not have simple DDG instances larger than 10 nodes. We are able to check the efficiency of our heuristic on small DDG instances by comparing its results against the optimal ones. The theoretical computation of the approximation ratio constitutes an additional problem which is not studied here.

Table 1 presents the results of SIRALINA against optimal method using common benchmarks. This table presents the results for the minimal period of each benchmark. Note that every benchmark has its own minimal period, defined as the critical circuit of the DDG, which is inherent to the data dependences [2]. The first column represents the name of the benchmark. The second and third column represent the instance size (numbers of DDG nodes and edges). Columns number 4 and 5 give the storage requirement (objective function values) computed by the optimal and SIRALINA methods (some instances could not be solved). The two last columns give the resolution times in seconds. As can be seen in this table, SIRALINA is fast and nearly optimal. Sometimes SIRALINA is slightly longer than the optimal method for two reasons: 1) the timer is too precise (milliseconds) and the interactions with operating system disturbs our timing measurements, and 2) SIRALINA performs in two steps while the optimal method performs in one step (resolving a unique integer linear program). In our context, we consider that a time difference which is less than 0.1 seconds is negligible. Another interesting remark is that the processing time of SIRALINA is relatively constant compared to the processing time of the optimal method. Another improvement of SIRALINA compared to the optimal method is that SIRALINA performs in relatively a constant resolution time (irrespective of the considered period p). Figure 2 illustrates some examples, where we can see that the optimal method performs in a high variable processing time in function of the period p while SIRALINA is more stable. Figure 3 shows that SIRALINA is still nearly optimal with various period values. This remark has been checked for all other benchmarks and periods: in almost all benchmarks, SIRALINA computes nearly optimal results for all periods in a satisfactory (fast) processing time.

For large instances we compare our results against the ones obtained by two heuristics (f3 and f5) proposed in [10]. Results are reported in table 2, where S_{f_3} and S_{f_5} design the storage requirement (objective function values) computed by heuristics f3 and f5 and the last column gives the resolution time in seconds with SIRALINA. Benchmarks presented in table 2 correspond to instances where loop bodies are duplicated ten times (number of DDG nodes multiplied by 10) and for which results with heuristic f3 and f5 have been presented in [10].

Benchmark	$ V $	$ E $	S_{opt}	$S_{siralina}$	T_{opt}	$T_{siralina}$
lin-ddot	4	4	7	7	0.007	0.066
liv-loop1	9	11	5	5	0.364	0.067
liv-loop5	5	5	3	3	0.005	0.066
liv-loop23	20	26	10	12	605.548	0.069
spec-dod-loop1	13	15	5	6	198.472	0.067
spec-dod-loop2	10	10	3	3	0.084	0.067
spec-dod-loop3	11	11	3	4	0.257	0.067
spec-dod-loop7	4	4	35	35	0.004	0.066
spec-fp-loop1	5	6	2	2	0.006	0.067
spec-spice-loop1	2	2	3	3	0.004	0.067
spec-spice-loop2	9	10	15	15	2.757	0.067
spec-spice-loop3	4	5	2	2	0.005	0.067
spec-spice-loop4	12	51	8	8	0.088	0.068
spec-spice-loop5	2	2	1	1	0.003	0.067
spec-spice-loop6	6	7	14	14	0.016	0.067
spec-spice-loop7	5	5	40	40	0.005	0.067
spec-spice-loop8	4	4	7	7	0.005	0.067
spec-spice-loop9	11	17	7	7	26.242	0.067
spec-spice-loop10	4	4	2	2	0.005	0.069
spec-tom-loop1	15	18	5	7	604.278	0.068
test-christine	18	17	230	230	600.847	0.068
Elliptic	36	59	NA	10	NA	0.074
whet-cycle4-1	4	4	1	1	0.005	0.066
whet-cycle4-2	4	4	2	2	0.006	0.067
whet-cycle4-4	4	4	4	4	0.01	0.067
whet-cycle4-8	4	4	8	8	0.013	0.069
whet-loop1	16	28	5	6	0.2	0.068
whet-loop2	7	10	5	5	0.006	0.067
whet-loop3	5	16	4	4	0.006	0.067

Table 1. SIRALINA and optimal Results

Benchmark	S_{f3}	S_{fs}	$S_{siralina}$	$T_{siralina}$
lin-ddot-10	40	27	7	0.072
liv-loop1-10	80	60	5	0.094
spec-dod-loop3-10	100	75	4	0.107
spec-spice-loop1-10	20	19	3	0.069
spec-spice-loop3-10	30	29	2	0.071
spec-spice-loop6-10	60	48	14	0.081

Table 2. SIRALINA and other heuristics

6 Conclusion

This article presents an efficient heuristic for the periodic task scheduling problem under storage constraints. Our model is based on the theoretical approach of reuse graphs




Fig. 2. Processing Times of some Benchmarks vs. Period

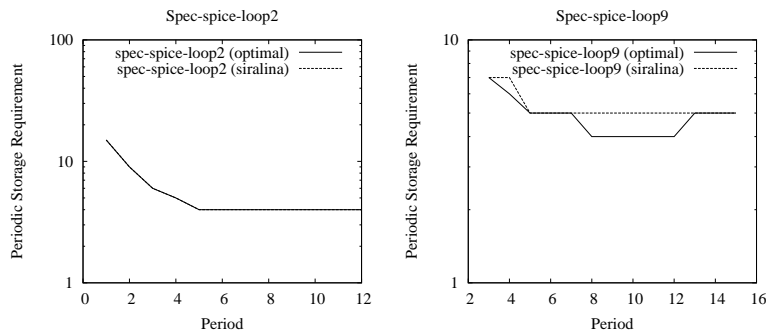


Fig. 3. Storage Requirement of some Benchmarks vs. Period

studied in [9]. Storage allocation is expressed in terms of reuse edges and reuse distances to model the fact that two tasks use the same storage location.

Since computing an optimal periodic storage allocation is intractable in large data dependence graphs (larger than 12 nodes for instance), we have identified a two steps resolution method. We call this simplified method as SIRALINA. A first optimal step provides scheduling variables and allows to compute the potential reuse distances if the corresponding reuse edge is added. Then a second step solves a linear assignment problem using the Hungarian method in order to select the appropriate reuse edges.

Our practical experiments on many DDGs show that SIRALINA provides satisfactory solutions with fast resolution times. Consequently, this method is included inside a compiler for embedded systems (in collaboration with STmicroelectronics).

Finally our future work will concentrate on the particular structure of the model constraints to consider the application of lagrangean relaxation to produce a bound stronger than the bound obtained by continuous relaxation and/or to find another heuristic. Furthermore it will be interesting to investigate how some hardware specificities could be taken into account. For instance, the use of a rotating register file (implemented inside some processors) implies the presence of a Hamiltonian reuse cycle. The handling of these kind of specificities is an additional challenge.

Acknowledgement

This work has been partially supported by the ANR MOPUCE project (ANR number 05-JCJC-0039).

References

1. Benoit Dupont de Dinechin. Parametric Computation of Margins and of Minimum Cumulative Register Lifetime Dates. In David C. Sehr and Utpal Banerjee and David Gelernter and Alexandru Nicolau and David A. Padua, editor, *LCPC*, volume 1239 of *Lecture Notes in Computer Science*, pages 231–245. Springer, 1996.
2. Claire Hanen and Alix Munier. A Study of the Cyclic Scheduling Problem on Parallel Processors. *Discrete Applied Mathematics*, 57(2-3):167–192, 1995.
3. A. E. Eichenberger, E. S. Davidson, and S. G. Abraham. Minimizing Register Requirements of a Modulo Schedule via Optimum Stage Scheduling. *International Journal of Parallel Programming*, 24(2):103–132, Apr. 1996.
4. D. Fimmel and J. Muller. Optimal Software Pipelining Under Resource Constraints. *International Journal of Foundations of Computer Science (IJFCS)*, 12(6):697–718, 2001.
5. Harold W. Kuhn. The Hungarian Method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
6. J. Janssen. *Compilers Strategies for Transport Triggered Architectures*. PhD thesis, Delft University, Netherlands, 2001.
7. A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley and Sons, New York, 1987.
8. Sid-Ahmed-Ali Touati. *Register Pressure in Instruction Level Parallelism*. PhD thesis, Université de Versailles, France, June 2002. <ftp.inria.fr/INRIA/Projects/a3/touati/thesis>.
9. Sid-Ahmed-Ali Touati. Periodic Task Scheduling under Storage Constraints. In *Proceedings of the Multidisciplinary International Scheduling Conference: Theory and Applications (MISTA'07)*, Aug. 2007.
10. Sid-Ahmed-Ali Touati and Christine Eisenbeis. Early Periodic Register Allocation on ILP Processors. *Parallel Processing Letters*, 14(2), June 2004. World Scientific.
11. M. M. Strout, L. Carter, J. Ferrante, and B. Simon. Schedule-Independent Storage Mapping for Loops. *ACM SIG-PLAN Notices*, 33(11):24–33, Nov. 1998.
12. W. Thies, F. Vivien, J. Sheldon, and S. Amarasinghe. A Unified Framework for Schedule and Storage Optimization. *ACM SIGPLAN Notices*, 36(5):232–242, May 2001.