



K. Pingali U. Banerjee D. Gelernter  
A. Nicolau D. Padua (Eds.)

# Languages and Compilers for Parallel Computing

7th International Workshop  
Ithaca, NY, USA, August 8-10, 1994  
Proceedings

Springer-Verlag  
Berlin Heidelberg New York  
London Paris Tokyo  
Hong Kong Barcelona  
Budapest

Series Editors

Gerhard Goos  
Universität Karlsruhe  
Vincenz-Priessnitz-Straße 3, D-76128 Karlsruhe, Germany

Juris Hartmanis  
Department of Computer Science, Cornell University  
4130 Upson Hall, Ithaca, NY 14853, USA

Jan van Leeuwen  
Department of Computer Science, Utrecht University  
Padualaan 14, 3584 CH Utrecht, The Netherlands

Volume Editors

Keshav Pingali  
Department of Computer Science, Cornell University  
Ithaca, NY 14853, USA

Utpal Banerjee  
Intel Corporation  
2200 Mission College Blvd. P.O. Box 58119, RN6-18  
Santa Clara, CA 95052, USA

David Gelernter  
Department of Computer Science, Yale University  
51 Prospect St., New Haven, CT 06520, USA

Alex Nicolau  
Department of Information & Computer Science, University of California  
444 Computer Science Bldg., Irvine, CA 92717, USA

David Padua  
Center for Supercomputing Research and Development  
465 Computer and Systems Research Laboratory  
1308 West Main St., Urbana, IL 61801, USA

CR Subject Classification (1991): F.1.2, D.1.3, D.3.1, B.2.1, I.3.1

ISBN 3-540-58868-X Springer-Verlag Berlin Heidelberg New York

CIP data applied for

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1995  
Printed in Germany

Typesetting: Camera-ready by author  
SPIN: 10479285 45/3140-543210 - Printed on acid-free paper

## Foreword

The papers in this volume are revised versions of papers presented at the Seventh Annual Workshop on Languages and Compilers for Parallel Computers which was held in Ithaca, NY August 8th-10th, 1994. This workshop series has traditionally been a forum for the presentation of state-of-the-art research in languages, restructuring compilers and runtime systems by leading groups in the US, Europe and Japan. This year, we received about 45 submissions in response to the call for papers. Because of time constraints, we were unable to accept all papers, and these proceedings contain the 32 papers that were selected for presentation.

Thanks are due to many people for making the workshop a success. The members of the standing program committee — Utpal Banerjee, David Gelernter, Alex Nicolau and David Padua — were a great source of advice and information. Helene Croft put the entire database of names and addresses online, which should make it easier to run this workshop in future years. She was also relentless in tracking down tardy authors, referees, speakers and attendees, which made the workshop run like clockwork! Thanks are due to the referees: Utpal Banerjee, Carrie Brownhill, A. Capitanio, Sudeep Gupta, David Gelernter, Laurie Hendren, J. Hummel, Richard Johnson, Indu Kodukula, David Kolson, Vladimir Kotlyar, Wei Li, Mayan Moudgill, Alex Nicolau, Rishiyur Nikhil, Steve Novack, F. Onion, David Padua, Anne Rogers, Paul Stodghill, Thorsten von Eicken, H. Wang and Richard Zippel. Finally, the Cornell Department of Computer Science, and the Cornell Theory Center gave generous financial grants, which allowed us to reduce the registration fee for participants.

October, 1994

Keshav Pingali  
Program Chair

# Table of Contents

## Starting Small: Fine-Grain Parallelism

<b>Fine-Grain Scheduling under Resource Constraints . . . . .</b>	<b>1</b>
---	----------

Paul Feautrier  
*Université de Versailles, France*

<b>Mutation Scheduling: A Unified Approach to Compiling for Fine-Grain Parallelism . . . . .</b>	<b>16</b>
--	-----------

Steven Novack, Alexandru Nicolau  
*University of California at Irvine*

<b>Compiler Techniques for Fine-Grain Execution on Workstation Clusters Using PAPERS . . . . .</b>	<b>31</b>
--	-----------

H.G. Dietz, W.E. Cohen, T. Muhammad, T.I. Mattox  
*Purdue University, West Lafayette, Indiana*

## Getting Your Ducks in a Row: Alignment and Distribution

<b>Solving Alignment Using Elementary Linear Algebra . . . . .</b>	<b>46</b>
--	-----------

David Bau, Induprakas Kodukula, Vladimir Kotlyar, Keshav Pingali, Paul Stodghill  
*Cornell University, Ithaca, New York*

<b>Detecting and Using Affinity in an Automatic Data Distribution Tool . . . . .</b>	<b>61</b>
--	-----------

Eduard Ayguadé, Jordi Garcia, Mercè Gironés, Jesús Labarta, Jordi Torres, Mateo Valero  
*Polytechnic University of Catalunya, Barcelona, Spain*

<b>Array Distribution in Data-Parallel Programs . . . . .</b>	<b>76</b>
---	-----------

Siddhartha Chatterjee, Robert Schreiber, Thomas J. Sheffler  
*NASA Ames Research Center, Moffett Field, California*  
John R. Gilbert  
*Xerox Research Center, Palo Alto, California*

## Postlinear Loop Transformations

<b>Communication-Free Parallelization via Affine Transformations .</b>	<b>92</b>
--	-----------

Amy W. Lim, Monica S. Lam  
*Stanford University, California*

**Finding Legal Reordering Transformations Using Mappings . . . 107**

Wayne Kelly, William Pugh  
*University of Maryland, College Park*

**A New Algorithm for Global Optimization for Parallelism and Locality . . . . . 125**

Bill Appelbe, Srinivas Doddapaneni, Charles Hardnett  
*Georgia Institute of Technology, Atlanta*

**What Next?****Polaris: Improving the Effectiveness of Parallelizing Compilers . 141**

William Blume, Rudolf Eigenmann, Keith Faigin, John Grout,  
 Jay Hoeflinger, David Padua, Paul Petersen, William Pot-  
 tenger, Lawrence Rauchwerger, Peng Tu, Stephen Weatherford  
*University of Illinois at Urbana*

**A Formal Approach to the Compilation of Data-Parallel Languages . . . . . 155**

J.A. Trescher, L.C. Breebaart, P.F.G. Dechering,  
 A.B. Poelman, J.P.M. de Vreught, H.J. Sips  
*Delft University of Technology, The Netherlands*

**The Data Partitioning Graph: Extending Data and Control Dependencies for Data Partitioning . . . . . 170**

Tsuneo Nakanishi, Kazuki Joe, Akira Fukuda, Keijiro Araki  
*Nara Institute of Science and Technology, Japan*  
 Hideki Saito, Constantine D. Polychronopoulos  
*University of Illinois at Urbana*

**Back to Basics: Program Analysis****Detecting Value-Based Scalar Dependence . . . . . 186**

Eric Stoltz, Michael Wolfe  
*Oregon Graduate Institute of Science & Technology, Portland*

**Minimal Data Dependence Abstractions for Loop Transformations . . . . . 201**

Yi-Qing Yang, Corinne Ancourt, François Irigoin  
*Ecole des Mines de Paris, France*

<b>Differences in Algorithmic Parallelism in Control Flow and Call Multigraphs . . . . .</b>	<b>217</b>
Vincent Sgro, Barbara G. Ryder <i>Rutgers University, New Brunswick, New Jersey</i>	
<b>Flow-Insensitive Interprocedural Alias Analysis in the Presence of Pointers . . . . .</b>	<b>234</b>
Michael Burke, Paul Carini, Jong-Deok Choi <i>IBM Watson Research Center, Yorktown Heights, NY</i> Michael Hind <i>State University of New York at New Paltz, and</i> <i>IBM Watson Research Center, Yorktown Heights, NY</i>	
<b>How to Communicate Better</b>	
<b>Incremental Generation of Index Sets for Array Statement Execution on Distributed-Memory Machines . . .</b>	<b>251</b>
S.D. Kaushik, C.-H. Huang, P. Sadayappan <i>The Ohio State University, Columbus</i>	
<b>A Unified Data-Flow Framework for Optimizing Communication</b>	<b>266</b>
Manish Gupta, Edith Schonberg, Harini Srinivasan <i>IBM Watson Research Center, Yorktown Heights, NY</i>	
<b>Interprocedural Communication Optimizations for Distributed Memory Compilation . . . . .</b>	<b>283</b>
Gagan Agrawal, Joel Saltz <i>University of Maryland, College Park</i>	
<b>Automatic Parallelization Considered Unnecessary</b>	
<b>Analysis of Event Synchronization in Parallel Programs . . . . .</b>	<b>300</b>
J. Ramanujam, A. Mathew <i>Louisiana State University, Baton Rouge</i>	
<b>Computing Communication Sets for Control Parallel Programs .</b>	<b>316</b>
Jeanne Ferrante <i>IBM Watson Research Center, Yorktown Heights, NY</i> Dirk Grunwald, Harini Srinivasan <i>University of Colorado, Boulder</i>	
<b>Optimizing Parallel SPMD Programs . . . . .</b>	<b>331</b>
Arvind Krishnamurthy, Katherine Yelick <i>University of California at Berkeley</i>	

## **Languages for Parallelism: Something Borrow, Something New**

**An Overview of the Opus Language and Runtime System . . . . . 346**

Piyush Mehrotra, Matthew Haines  
*Institute for Computer Applications in Science and Engineer-  
ing, Hampton, Virginia*

**SIMPLE Performance Results in ZPL . . . . . 361**

Calvin Lin, Lawrence Snyder  
*University of Washington, Seattle*

**Cid: A Parallel, “Shared-Memory” C for Distributed-Memory  
Machines . . . . . 376**

Rishiyur S. Nikhil  
*Digital Equipment Corporation, Cambridge, Massachusetts*

**EQ: Overview of a New Language Approach for Prototyping  
Scientific Computation . . . . . 391**

Thomas Derby, Robert Schnabel, Benjamin Zorn  
*University of Colorado at Boulder*

## **And Now, for Something Completely Different**

**Reshaping Access Patterns for Generating Sparse Codes . . . . . 406**

Aart J.C. Bik, Peter M.W. Knijnenburg, Harry A.G. Wijshoff  
*Leiden University, The Netherlands*

**Evaluating Two Loop Transformations for Reducing Multiple-  
Writer False Sharing . . . . . 423**

François Bodin, Thierry Montaut  
*IRISA, Campus de Beaulieu, Rennes, France*  
Elana D. Granston  
*Rice University, Houston, Texas*

**Parallelizing Tree Algorithms: Overhead vs. Parallelism . . . . . 440**

Jon A. Solworth, Bryan B. Reagan  
*University of Illinois at Chicago*

## **When Your Program Runs (Finally)**

**Autoscheduling in a Distributed Shared-Memory Environment . 455**

José E. Moreira, Constantine D. Polychronopoulos  
*University of Illinois at Urbana-Champaign*



**Optimizing Array Distributions in Data-Parallel Programs . . . . 472**

Krishna Kunchithapadam, Barton P. Miller  
*University of Wisconsin at Madison*

**Automatic Reduction Tree Generation for Fine-Grain Parallel  
 Architectures when Iteration Count is Unknown . . . . . 487**

Kei Hiraki  
*University of Tokyo, Japan*  
 Satoshi Sekiguchi  
*Electrotechnical Laboratory, MITI, Ibaraki, Japan*

# Fine-grain Scheduling under Resource Constraints

Paul Feautrier

Laboratoire PRiSM,  
Université de Versailles, 45 Avenue des Etats-Unis,  
78035 VERSAILLES FRANCE

**Abstract.** Many present-day microprocessors have fine grain parallelism, be it in the form of a pipeline, of multiple functional units, or replicated processors. The efficient use of such architectures depends on the capability of the compiler to schedule the execution of the object code in such a way that most of the available hardware is in use while complying with data dependences. In the case of one simple loop, the schedule may be expressed as an affine form in the loop counter. The coefficient of the loop counter in the schedule is the initiation interval, and gives the mean rate at which loop bodies may be executed. The dependence constraints may be converted to linear inequalities in the coefficients of a closed form schedule, and then solved by classical linear programming algorithms. The resource constraints, however, translate to non-linear constraints. These constraints become linear if the initiation interval is known. This leads to a fast searching algorithm, in which the initiation interval is increased until a feasible solution is found.

## 1 Introduction

Thinking about parallel programs is a notoriously difficult task. One of the most successful techniques for dealing with this problem is *scheduling*, i.e. the construction of a timetable for the operations of the program. Scheduling is a difficult problem. Various special cases have been proved to be NP-hard or NP-complete. Most of the complexity of scheduling can be assigned to the conjunction of two type of constraints:

- dependence constraints, which express the fact that some computations must be executed in a specified order if the meaning of the original program is to be preserved; these constraints are usually expressed as a dependence graph.
- resource constraints, which express the fact that the number of simultaneous operation at any given time is limited by the available resources in the target computer.

While any one of these constraints can be handled easily, it is their simultaneous presence which is at the origin of the difficulty. Fortunately, in many cases of computer science interest, it is possible to handle the resource constraints in an approximate way. It is customary in this context, to distinguish between *coarse grain*, *medium grain*, and *fine grain* scheduling.

In coarse grain scheduling – e.g., job shop scheduling or macrotasking – the tasks and the resources are few. The schedule is represented in tabular form, and there are approximate techniques, like list scheduling, with precise bounds on the approximation.

In medium grain scheduling, there are many tasks – typically as many as there are operations in an execution of the source program – and many identical resources – the processors in a massively parallel computer. The schedule must be obtained in closed form. One may ignore the resource constraints in computing the schedule [Fea92a, Fea92b], and then fold the schedule on the available processors. One may prove [Fea89] that this solution is *asymptotically efficient*, provided that the source program has enough intrinsic parallelism.

The situation is different for fine grain scheduling. Here the number of tasks is large. The resources are few and discrete. At the most, resources may be partitioned into classes, each class having a small number of identical resources.

Fine grain scheduling started some thirty years ago when the first computers with multiple functional units – like the CDC 6600 – were put on the market. It is now a very important technique, due to the advent of many computers with instruction level parallelism, like pipelined computers, VLIW or superscalar processors.

In fine grain scheduling, it is impossible to ignore the resource constraints. Several techniques have been proposed for solving the problem, at least in an approximate way (see [RF93] for a comprehensive review of the subject). Trace scheduling [Fis84] applies list scheduling to basic blocks; it tries to detect critical paths in the program control graph and to enlarge basic blocks by moving code around test instructions.

Software pipelining [RG81] applies to simple loops and aims at executing several instances of the loop body in a staggered way so as to maximize resource usage and minimize the total execution time. A solution to the software pipelining problem for a given loop is characterized by its *initiation interval*, i.e. the time span between two successive iteration of the loop. It is easy to derive bounds for the initiation interval: an upper bound is given by the sequential execution time of the loop body. A lower bound is deduced from an analysis of resource usage, see section 2.2, and another one can be obtained by constructing an unconstrained schedule.

In many algorithms for software pipelining, one assume the iteration interval is given, and applies list scheduling, taking care that each resource allocation is folded modulo the initiation interval when constructing the reservation table (see e.g. [Lam88]). The interval of admissible initiation intervals is explored by binary search until the optimal value is found.

The algorithm of [GS92] applies only if there is only one resource class. The program is first scheduled as if there were no resource constraints. Analysis of the resulting schedule allows one to delete some dependences, and the resulting dependence graph is cycle free. The final schedule is obtained by applying list scheduling with resource constraints to this graph. The resulting schedule is not optimal, but the authors show that the usual bound on the list scheduling

approximation applies.

This paper is an attempt to extend the scheduling techniques of [Fea92a], which are based on linear programming, to fine grain scheduling. The next section is a review of these techniques. The main theme of Section 3 is how to convert the resource constraints into bilinear constraints. This is done in two cases. In the first one, there is one unique resource of each type; in the second case, there may be several copies of a resource. In the conclusion, I discuss the complexity of the algorithm and point to some direction for future work.

## 2 A Review of Scheduling Techniques

I will consider here the problem of scheduling a single loop:

```
do  $i = 1, \dots$ 
   $S_1; \dots S_n$ 
end do
```

where the  $S_k$  are scalar or array assignments. I have emphasized the fact that the upper bound of the loop is irrelevant for the present problem. The solution must be in the form of the repetition of a uniform pattern, the loop upper bound controlling only the repetition factor.

A schedule is defined by  $n$  functions from the iteration counter,  $i$ , to an integral time. I will suppose that an appropriate unit of time has been chosen – e.g., the clock cycle – and that all delays and durations are integral multiple of this unit. Schedules are supposed to be of the form:

$$\theta(S, i) = \lfloor ai + b_S \rfloor, \quad (1)$$

where  $a$  and the  $b_S$  are rational numbers.  $a$  is known as the initiation interval of the schedule. The main objective of software pipelining is its minimization.

There are several reasons for choosing such a form. Firstly, all known methods for computing schedules apply only to affine forms. It is true that a schedule whose values are not integral has no meaning, but it has been shown that the floor of a causal schedule is also causal, and that if the iteration domain is large enough, schedules of the above form are nearly optimal [Qui87].

Before embarking on the solution proper, let us observe that there is some leeway in the selection of  $b_S$  in (1).  $a$  is necessarily a rational number – if it were not so, the schedule would not be periodic. It is easy to prove the following lemma:

**Lemma 1.** *Let  $a = A/D$  be the representation of the initiation interval in lowest terms. Any schedule of the form (1) is equivalent to a schedule of the form:*

$$\theta(i) = \left\lfloor \frac{Ai + B_S}{D} \right\rfloor \quad (2)$$

where the  $B_S$  are integers.

Recent research on medium-grain scheduling [MQRS90, Fea92a] favors schedules in which each statement has its own initiation interval. In the case of fine grain parallelism, such a schedule generates very complicated code<sup>2</sup>, hence our insistence on the same initiation interval for all statements.

All schedules must satisfy the so-called *causality condition*: let us write  $(S_k, i) \perp (S_l, j)$  if  $(S_k, i)$  and  $(S_l, j)$  are in dependence, and  $(S_k, i) \prec (S_l, j)$  if  $(S_k, i)$  is executed before  $(S_l, j)$  in the original program. The schedule must verify:

$$(S_k, i) \perp (S_l, j) \wedge (S_k, i) \prec (S_l, j) \Rightarrow \theta(S_k, i) + \partial(S_k) \leq \theta(S_l, j), \quad (3)$$

where  $\partial(S_k)$  is the duration of  $S_k$ .

## 2.1 Dependences

The choice of the dependence relation in (3) is somewhat arbitrary. Ordinary dependences include both the effect of data flow from operation to operation and the constraints generated by the pattern of memory usage in the object program. Value based dependences are much less constraining and are easily computed by Array Dataflow Analysis [Fea91]. There is a value-based dependence between  $(S, i)$  and  $(R, j)$  iff  $(S, i)$  writes into some memory cell  $\mathbf{a}$ , if  $(R, j)$  reads  $\mathbf{a}$ ,  $(S, i) \prec (R, j)$ , and there is no write to  $\mathbf{a}$  between  $(S, i)$  and  $(R, j)$ . The result of Array Dataflow Analysis may be represented by a Dataflow Graph (DFG), whose vertices are associated to statements and edges to dependences. Each edge  $e$  from  $S$  to  $R$  is decorated with a polyhedron  $\mathcal{P}_e$  and a transformation  $h_e$  such that if  $i \in \mathcal{P}_e$  then there is actually a value-based dependence from  $(S, h_e(i))$  to  $(R, i)$ . One may say that after Array Dataflow Analysis, all values produced by the source code have been given distinct names, and the program has been rewritten using these names. Array Dataflow Analysis may thus be seen as a compile time counterpart of Tomasulo Algorithm.

The shape of the dependence is given by the function  $h_e$ . The simplest case is that of uniform dependences for which  $h_e$  is a translation:

$$h_e(i) = i - d_e$$

where  $d_e$  is known as the dependence distance. One may encounter more complicated cases, where  $h_e$  is an affine function, or even a sublinear function<sup>3</sup>. The scheduling technique of [Fea92a] works whenever the dependence is affine and is not limited to uniform dependences.

Value based dependences will be used throughout this paper. In this context, the causality condition (3) simplifies to:

$$\forall e \in DFG, \forall i \in \mathcal{P}_e : \theta(R, i) \geq \theta(S, h_e(i)) + \partial(S). \quad (4)$$

<sup>2</sup> The size of the code grows as the least common multiple of the initiation intervals.

<sup>3</sup> A sublinear function contains integer divisions by constants.

This condition expresses the fact that since operation  $(R, i)$  uses a value which is computed by  $(S, h_e(i))$ , it cannot start before this operation has terminated.

The solution method starts by substituting the form (1) into (4). In the case of uniform dependences, one may prove that:

**Lemma 2.** *The causality condition (4) is equivalent to:*

$$ad_e + b_R - b_S \geq \partial(S). \quad (5)$$

*Proof.* That (5) implies (4) is proved in [Fea92a] Theorem 6. To prove the reverse implication, choose for  $i$  a multiple of  $D$ . Notice that if  $n$  is an integer, we have the identity  $\lfloor n + x \rfloor = n + \lfloor x \rfloor$ . (4) simplifies to:

$$\lfloor B_R/D \rfloor \geq \lfloor (B_R - Ad_e)/D \rfloor - \partial(S).$$

Since the left hand side of this inequality is an integer, we have:

$$\lfloor B_R/D \rfloor \geq (B_R - Ad_e)/D - \partial(S).$$

Now, obviously,  $x \geq \lfloor x \rfloor$ , hence:

$$B_R/D \geq \lfloor B_R/D \rfloor \geq (B_R - Ad_e)/D - \partial(S),$$

Q.E.D.

By the above lemma, each uniform dependence may be translated to a linear constraint on the  $a$  and  $b$ 's coefficients. For more complicated dependences, one has to resort to the Farkas algorithm [Fea92a], but the result is still a set of linear constraints. One then selects a particular solution according to some objective function. Of particular interest for fine grain scheduling are the minimum latency schedules, in which one minimizes first the initiation interval  $a$ , and then the  $b_R$ .

## 2.2 Resource constraints

In operation research, a *resource* is an entity which may or may not be used by tasks or operations. To each resource is associated a constraint: namely, that the execution intervals of two operations which use the same resource cannot overlap. One may have resource *classes*. In that case, at any given time, the number of active operations which use a given resource cannot exceed the number of resources in the class. I will suppose here that all operations which are instances of the same instruction use the same resource class. For simplicity, I will assume that each operation uses only one resource. This restriction can be easily lifted in case of need. In fact, in this work resource classes will simply be sets of statements. If  $\rho$  is a resource class,  $S \in \rho$  means that statement  $S$  uses a resource from class  $\rho$ .

In the case of unique resources, the non overlap constraint may be translated to simple inequalities on schedules. Suppose that  $S$  and  $T$  use the same resource. If  $\langle S, i \rangle$  is scheduled before  $\langle T, j \rangle$ , then we must have:

$$\theta(T, j) \geq \theta(S, i) + \partial(S),$$

while in the opposite situation the constraint is:

$$\theta(S, i) \geq \theta(T, j) + \partial(T).$$

Since the two situations are exclusive, we may write the resource constraint as:

$$\forall i, j : \theta(T, j) - \theta(S, i) \geq \partial(S) \vee \theta(S, i) - \theta(T, j) \geq \partial(T). \quad (6)$$

Beside that, two operations which are instance of the same instruction necessarily use the same resource and cannot overlap:

$$\forall i, j : i < j \Rightarrow |\theta(S, i) - \theta(S, j)| \geq \partial(S). \quad (7)$$

This condition gives a very simple bound on  $a$ . Suppose a large number  $N$  of iterations of the loop body are executed in time  $t$ . The total usage of resource  $\rho$  will be:

$$t_\rho \approx N \sum_{S \in \rho} \partial(S).$$

Suppose there are  $P_\rho$  copies of  $\rho$ :

$$t \approx Na \leq N \sum_{S \in \rho} \partial(S) / P_\rho,$$

from which one deduces the lower bound for  $a$ :

$$a \geq \max_{\rho} \sum_{S \in \rho} \partial(S) / P_\rho. \quad (8)$$

If the initiation interval satisfies the above constraint, (7) is automatically satisfied.

In actual processors, resource utilization may be a much more complicated affair than the simplified scheme above. Pipelined resources, for instance, do not appear to be busy for the whole duration of one operation. This is easily taken care of by replacing  $\partial(S)$  in (6) by another timing characteristics, the *stalling time* of operation  $S$ , noted  $\sigma(S)$ . The resource constraint is now:

$$\forall i, j : \theta(T, j) - \theta(S, i) \geq \sigma(S) \vee \theta(S, i) - \theta(T, j) \geq \sigma(T). \quad (9)$$

An ordinary functional unit will have  $\partial(S) = \sigma(S)$ , while a pipelined unit will have  $\sigma(S) \ll \partial(S)$ .

There may be links between resources, as for instance when one cannot use a functional unit unless there is a free data path to it. That kind of constraint must be handled heuristically.

The problem is more complicated if some resource class has more than one element. A resource is in use at time  $t$  if some statement  $S$  which uses it has been initiated less than  $\sigma(S)$  time units before  $t$ . If we identify a resource class with the set of statements which use it, we may write the constraint for resource  $\rho$  as:

$$\text{Card} \{(S, i) \mid S \in \rho \wedge t - \sigma(S) < \theta(S, i) \leq t\} \leq P_\rho. \quad (10)$$

### 3 Two Scheduling Algorithms

Basically, the scheduling method of [Fea92a] works by replacing (4), which represents a potentially infinite system of affine inequalities, by a finite set of constraints on the coefficients  $a$  and  $b_R$ . The first problem is to find a similar reduction for (9). Due to the non-convexity of (9), the result is non linear. Hence, one cannot directly use linear programming to solve the problem. However, the problem lends itself to a simple and efficient solution by searching the space of possible values for  $a$ .

#### 3.1 The singular resource case

For schedules of the form (2), the floor function in the expression of (6) may be ignored:

**Theorem 3.** *Let  $\tau(S, i) = \frac{Ai+Bs}{D}$  and  $\theta(S, i) = \lfloor \tau(S, i) \rfloor$ . Then the two conditions:*

$$\forall i, j : \theta(T, j) - \theta(S, i) \geq \sigma(S) \vee \theta(S, i) - \theta(T, j) \geq \sigma(T). \quad (11)$$

and

$$\forall i, j : \tau(T, j) - \tau(S, i) \geq \sigma(S) \vee \tau(S, i) - \tau(T, j) \geq \sigma(T). \quad (12)$$

are equivalent.

*Proof.* Suppose first that (12) is true. Let us be given two arbitrary integers  $i$  and  $j$ . We may suppose, without loss of generality, that  $\tau(S, i) - \tau(T, j) > 0$ . We have, successively:

$$\begin{aligned} \lfloor \tau(T, j) \rfloor &\leq \tau(T, j), \\ \lfloor \tau(T, j) \rfloor + \sigma(T) &\leq \tau(T, j) + \sigma(T) \leq \tau(S, i), \end{aligned}$$

and, since the left hand side is an integer,

$$\lfloor \tau(T, j) \rfloor + \sigma(T) \leq \lfloor \tau(S, i) \rfloor,$$

Q.E.D.

Conversely, suppose that (12) is false for some values of  $i$  and  $j$ . Set  $x = i - j$ . We have both:  $Ax + Bs - B_T \leq \sigma(T)$  and  $B_T - Bs - Ax \leq \sigma(S)$ . Set  $B = B_T - Bs$  for short. We may suppose that  $Ax - B > 0$ . The other case is handled in a symmetrical fashion. We have, for all  $j$ :

$$\tau(S, j + x) - \tau(T, j) = \frac{Ax - B}{D} \leq \sigma(T).$$

Since  $A$  and  $D$  are relatively prime, a  $j$  may be selected in such a way that  $\tau(S, j + x)$  is an integer. We then have:

$$\begin{aligned} \lfloor \tau(S, j + x) \rfloor &= \tau(S, j + x) \leq \tau(T, j) + \sigma(T), \\ \lfloor \tau(S, j + x) \rfloor &\leq \lfloor \tau(T, j) \rfloor + \sigma(T), \end{aligned}$$

(11) is also false, Q.E.D.



With the help of this result, the resource constraint above may be written in the form:

$$\forall x \in \mathbb{Z} : \frac{Ax - B}{D} \geq \sigma(T) \vee \frac{B - Ax}{D} \geq \sigma(S).$$

Now  $\frac{Ax-B}{D}$  is an affine function of  $x$  whose zero is  $x_0 = B/A$ . For all values of  $x > x_0$ , the second inequality is certainly not verified. Hence, the first one must be true, and a necessary and sufficient condition is that:

$$A(\lfloor B/A \rfloor + 1) - B \geq D\sigma(T).$$

The other case is handled similarly and gives:  $B - A \lfloor B/A \rfloor \geq D\sigma(S)$ . These conditions may even be simplified by observing that, if they are true, then there exists a unique integer  $q$  such that:

$$A(q + 1) - B \geq D\sigma(T) \wedge B - Aq \geq D\sigma(S). \quad (13)$$

As a consequence, the resource constraints in the singular case are given by the following rule:

For all statements  $S$  and  $T$  which use the same resource:

- Create a new integer variable  $q_{ST}$ ,
- Write the two constraints:

$$\begin{aligned} A(q_{ST} + 1) - B_T + B_S &\geq D\sigma(T), \\ B_T - B_S - Aq_{ST} &\geq D\sigma(S). \end{aligned} \quad (14)$$

These constraints are to be added to the dependence constraints and solved for  $A$  and the  $B_S$ ,  $A$  being the objective function to be minimized. Now the constraints generated by (14) are clearly non linear. However, they become linear if the value of  $A$  is known. Remember that we have one upper bound for  $a = A/D$  which is simply the sum of the duration of all statements in the loop body – the *sequential* upper bound – and two lower bounds. One of them, the *resource usage* bound, is given by (8), and the other, the *free* bound, is obtained simply by solving the scheduling problem with no resource constraints. The maximum of these two bounds gives the *parallel* lower bound. The problem is that, since  $a$  is a rational number, exploring its possible range of values is not a finite process. As has been observed many times, the schedule (2) has period  $D$ .  $D$  iterations of the loop body are scheduled in  $A$  clock cycles, giving a mean activation interval of  $A/D$ . When generating code from such a schedule, the loop body has to be replicated  $D$  times, which means that  $D$  cannot be too large. In the singular resource case, the resource usage bound is an integer. The free bound may be rational, but the actual value of its denominator is no indication, because simplification may occur depending on the values of the statement durations. A better guess may be obtained by observing that when computing the free schedule, one has to solve a linear programming problem by a process analogous to Gaussian elimination. By the familiar Cramer rule, the denominator of the solution is the determinant of a submatrix of the problem tableau, the basis matrix. The value of this determinant can be easily extracted

from the linear programming code, and is a good candidate for the unrolling factor.

We have found in practice that the following heuristic gives satisfactory results :

1. Compute the free bound, the resource usage bound and the parallel lower bound,  $l$ , which is their maximum.
2.  $D$  is set equal to the determinant of the basis matrix or to 1, depending whether the parallel lower bound is the larger bound or not.
3. Set  $A = \lceil Dl \rceil$ .
4. Solve the complete scheduling problem for  $A$  and  $D$ .
5. If the problem has no solution, increase  $A$  by 1 and start again at step 4.

Let us consider first a very simple example:

```

program A

  do i = 1,n
1      r1 = a(i)-b(i)
2      c(i) = c(i-2) + r1
  end do

```

Suppose that all operations are executed in unit time. Let  $\theta(1, i) = ai + b_1$  and  $\theta(2, i) = ai + b_2$  be the prototype schedules. There are two dependences:

- The first one is from  $\langle 1, i \rangle$  to  $\langle 2, i \rangle$  and gives the constraint:  $b_2 - b_1 \geq 1$ .
- The second one is from  $\langle 2, i - 2 \rangle$  to  $\langle 2, i \rangle$  and gives:  $2a \geq 1$ .

It is easy to see that the minimum latency solution is:

$$\theta(1, i) = i/2, \quad \theta(2, i) = i/2 + 1.$$

Suppose now that both statements of the example are executed on the same resource. This gives the following additional constraints:

$$a(q + 1) + b_1 - b_2 \geq 1, \quad b_2 - b_1 - aq \geq 1.$$

Since there are two statements in the loop and only one resource, we must have  $a \geq 2$ . An attempt to solve the remaining constraints with  $A = 2$ ,  $D = 1$  succeeds and gives:

$$\theta(1, i) = 2i, \quad \theta(2, i) = 2i + 1.$$

Since  $A$  has an upper and a lower bound, it may seem that a binary search for the right value might be a good idea. However, experiment shows that the solution is always near the lower bound. In that case, a simple linear search is sufficient. Let us consider the following example:

```

program B

  do i = 1,n
1      r0 = a(i-2)/2.0
2      r1 = r0+a(i-3)
3      r2 = r0+a(i-4)
4      a(i) = r1*r2
  end do

```

Suppose that the available resources are an adder, a multiplier and a divider, and that addition takes one cycle, multiplication and division taking two cycles. Analysis of resource usage shows that the minimum initiation interval is two cycles. Dependence analysis shows that statement 1 has to be executed first, that 2 and 3 can be executed in parallel, and that 4 is to be executed last. However, since the cycle is closed by a dependence from 4 at iteration  $i$  to 1 at iteration  $i + 2$ , this gives a minimum rate of  $5/2$ , and this is the parallel lower bound. Hence, we set  $D = 2$ . The first value of  $A$  to be tested is 5, and the integer programming algorithm finds that there is no solution.  $A$  is thus increased to 6, and there is a solution. It is easy to see *a posteriori* that this solution is optimal. In fact, since there is only one adder, statements 1 and 2 must be executed sequentially. Hence each iteration of the loop cannot take less than 6 cycles. The resulting initiation interval is  $6/2 = 3$ , indicating that no unrolling is necessary.

Suppose now that the multiplication time is reduced to 1 cycle. The free bound decreases to 2, but the determinant of the basis matrix is still 2. Hence, we set  $D = 2$  and  $A = 4$ . The first solution is found at the second iteration when  $A = 5$ , giving an initiation interval of  $5/2$  with an unrolling factor of two. The schedule is:

$$\begin{aligned}\theta(1, i) &= 5/2i, & \theta(2, i) &= 5/2i + 2, \\ \theta(i, 3) &= 5/2i + 3, & \theta(4, i) &= 5/2i + 4.\end{aligned}$$

To solve this problem, three calls to the integer programming algorithm PIP were needed, which took 0.43 seconds on a low end workstation.

### 3.2 The many resource case

In the many resource case, the resource constraint is given by (10). In the singular case, the scheduler has to guess the value of  $D$  and to search for the value of  $A$ . The many resource case is evidently more complicated. Hence, I will suppose that the algorithm structure is the same, namely that the problem is to test whether,  $A$  and  $D$  being given, there is a possible assignment for the  $B_S$  which meets all the constraints of the problem.

Here again, the first step is to get rid of the floor function. Suppose  $t$  is given, and that we are trying to count how many instances of  $S$  are active at time  $t$ . The iteration counter of the active instances is a positive integer such that:

$$t - \sigma(S) < \left\lfloor \frac{Ai + B_S}{D} \right\rfloor \leq t. \quad (15)$$

All terms in these inequalities are integers. Hence, they can be rewritten as:

$$t - \sigma(S) + 1 \leq \left\lfloor \frac{Ai + B_S}{D} \right\rfloor < t + 1.$$

Now  $t - \sigma(S) + 1 \leq \left\lfloor \frac{Ai + B_S}{D} \right\rfloor$  and  $t - \sigma(S) + 1 \leq \frac{Ai + B_S}{D}$  are equivalent. In one direction, this is because  $\lfloor x \rfloor \leq x$ , and in the other, it results from the monotony of the floor function.

For the other inequality,  $\frac{Ai+B_S}{D} < t+1$  clearly imply  $\lfloor \frac{Ai+B_S}{D} \rfloor < t+1$ . In the reverse direction,  $\lfloor \frac{Ai+B_S}{D} \rfloor \leq t$  implies  $\frac{Ai+B_S}{D} < t+1$  by definition.

As a consequence, the iterations of  $S$  which are active at time  $t$  are solutions of:

$$Dt - D\sigma(S) + D \leq Ai + B_S < Dt + D.$$

The problem is to count the solutions of these inequalities with  $i$  as the unknown as a function of  $t$ .

Introducing an “excess” variable  $x$ , the constraints may be transformed into an equation:

$$Ai + B_S = Dt + D - 1 - x, \quad (16)$$

provided that  $x$  satisfies  $0 \leq x < D\sigma(S)$ . If  $N_S(t, x)$  is the count of solutions of (16) for given  $t$  and  $x$ , then the number of active iterations at time  $t$  is:

$$N_S(t) = \sum_{x=0}^{D\sigma(S)-1} N_S(t, x).$$

The first observation is that equation (16) has at most one solution, which is given by:

$$i = \frac{Dt + D - 1 - x - B_S}{A}.$$

To be a legitimate iteration number, this solution has to be a positive integer.  $i$  is obviously positive for large enough  $t$ . The effect of ignoring the positivity condition, is to overestimate the resource usage for the *prologue* of the loop nest. It is customary in the field to ignore this factor by considering only very long loops, and this is the best one can do at compile time, since, for most loops, the iteration count is a variable. It may be possible to do better under user guidance: for instance, to inhibit software pipelining when the user knows that the iteration count will be small.

The integrity condition is simply:

$$Dt + D - 1 - x - B_S \equiv 0 \pmod{A}. \quad (17)$$

This has to be evaluated for all values of  $t$ . It is clear, however, that the condition depends only on  $t \bmod A$ . It thus has to be tested for  $t \in [0, A-1]$ . Another point is that the correspondance from  $t$  to  $Dt \bmod A$  is bijective, since  $A$  and  $D$  are relatively prime. As a consequence, one may introduce a new variable  $t' = Dt \bmod A$ ,  $0 \leq t' \leq A-1$ . The number of solutions of (16) may be written:

$$N_S(t, x) = \delta((t' + D - 1 - x - B_S) \bmod A),$$

where  $\delta$  is a variant of the Kronecker symbol:

$$\delta(0) = 1, \quad \delta(i) = 0, i \neq 0.$$