

# Lecture Notes in Computer Science

Edited by G. Goos, J. Hartmanis and J. van Leeuwen

1808

**Springer**

*Berlin*

*Heidelberg*

*New York*

*Barcelona*

*Hong Kong*

*London*

*Milan*

*Paris*

*Singapore*

*Tokyo*

Santosh Pande Dharma P. Agrawal (Eds.)

# Compiler Optimizations for Scalable Parallel Systems

Languages, Compilation Techniques,  
and Run Time Systems



Springer

## Series Editors

Gerhard Goos, Karlsruhe University, Germany  
Juris Hartmanis, Cornell University, NY, USA  
Jan van Leeuwen, Utrecht University, The Netherlands

## Volume Editors

Santosh Pande  
Georgia Institute of Technology, College of Computing  
801 Atlantic Drive, Atlanta, GA 30332, USA  
E-mail: santosh@cc.gatech.edu

Dharma P. Agrawal  
University of Cincinnati, Department of ECECS  
P.O. Box 210030, Cincinnati, OH 45221-0030, USA  
E-mail: dpa@ececs.uc.edu

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Compiler optimizations for scalable parallel systems : languages,  
compilation techniques, and run time systems / Santosh Pande ; Dharma  
P. Agrawal (ed.). - Berlin ; Heidelberg ; New York ; Barcelona ; Hong  
Kong ; London ; Milan ; Paris ; Singapore ; Tokyo : Springer, 2001  
(Lecture notes in computer science ; 1808)  
ISBN 3-540-41945-4

CR Subject Classification (1998): D.3, D.4, D.1.3, C.2, F.1.2, F.3

ISSN 0302-9743

ISBN 3-540-41945-4 Springer-Verlag Berlin Heidelberg New York

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 New York  
a member of BertelsmannSpringer Science+Business Media GmbH

<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2001  
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Boller Mediendesign  
Printed on acid-free paper SPIN: 10720238 06/3142 5 4 3 2 1 0

# Preface

Santosh Pande<sup>1</sup> and Dharma P. Agrawal<sup>2</sup>

<sup>1</sup> College of Computing  
801 Atlantic Drive,  
Georgia Institute of Technology,  
Atlanta, GA 30332

<sup>2</sup> Department of ECECS, ML 0030,  
PO Box 210030,  
University of Cincinnati,  
Cincinnati, OH 45221-0030

We are very pleased to publish this monograph on Compiler Optimizations for Scalable Distributed Memory Systems. Distributed memory systems offer a challenging model of computing and pose fascinating problems regarding compiler optimizations ranging from language design to run time systems. Thus, the research done in this area serves as foundational to many challenges from memory hierarchy optimizations to communication optimizations encountered in both stand-alone and distributed systems. It is with this motivation that we present a compendium of research done in this area in the form of this monograph.

This monograph is divided into five sections : section one deals with languages, section two deals with analysis, section three with communication optimizations, section four with code generation, and section five with run time systems. In the editorial we present a detailed summary of each of the chapters in these sections.

We would like to express our sincere thanks to many who contributed to this monograph. First we would like to thank all the authors for their excellent contributions which really make this monograph one of a kind; as readers will see, these contributions make the monograph thorough and insightful (for an advanced reader) as well as highly readable and pedagogic (for students and beginners). Next, we would like to thank our graduate student Haixiang He for all his help in organizing this monograph and for solving latex problems. Finally we express our sincere thanks to the LNCS Editorial at Springer-Verlag for putting up with our schedule and for all their help and understanding. Without their invaluable help we would not have been able to put this monograph into its beautiful final shape!!! We sincerely hope the readers find the monograph truly useful in their work and be it further research or practice.

# Introduction

Santosh Pande<sup>1</sup> and Dharma P. Agrawal<sup>2</sup>

<sup>1</sup> College of Computing  
801 Atlantic Drive,  
Georgia Institute of Technology,  
Atlanta, GA 30332

<sup>2</sup> Department of ECECS, ML 0030,  
PO Box 210030,  
University of Cincinnati,  
Cincinnati, OH 45221-0030

## 1. Compiling for Distributed Memory Multiprocessors

### 1.1 Motivation

The distributed memory parallel systems offer elegant architectural solutions for highly parallel data intensive applications primarily because:

- They are highly scalable. These systems currently come in a variety of architectures like 3D torus, mesh and hypercube that allow addition of extra processors should the computing demands increase. Scalability is an important issue especially for high performance servers such as parallel video servers, data mining and imaging applications.
- With increase in parallelism, there is insignificant degradation in memory performance since memories are isolated and decoupled from direct accesses from processors. This is especially good for data intensive applications such as parallel databases and data mining that demand considerable memory bandwidths. In contrast, the memory bandwidths may not match the increase in number of processors in shared memory systems. In fact, the overall system performance may degrade due to increased memory contention. This in turn jeopardizes scalability of application beyond a point.
- Spatial parallelism in large applications such as Fluid Flow, Weather Modeling and Image Processing, in which the problem domains are perfectly decomposable, is easy to map on these systems. The achievable speedups are almost linear and this is primarily due to fast accesses to the data maintained in local memory.
- The interprocessor communication speeds and bandwidths have dramatically improved due to very fast routing. The performance ratings offered by newer distributed memory systems have improved although they are not comparable to shared memory systems in terms of Mflops.
- Medium grained parallelism can be effectively mapped onto the newer systems like the Meiko CS-2, Cray T3D, IBM SP1/SP2 and EM4 due to a

low ratio of communication/computation speeds. Communication bottleneck has decreased compared with earlier systems and this has opened up parallelization of newer applications.

## 1.2 Complexity

However, programming distributed memory systems remains very complex. Most of the current solutions mandate that the users of such machines must manage the processor allocation, data distribution and inter-processor communication in their parallel programs. Programming these systems for achieving the desired high performance is very complex. In spite of frantic demands by programmers, current solutions provided by (semi-automatic) parallelizing compilers are rather constrained. As a matter of fact, for many applications the only practical success has been through hand parallelization of codes with communication managed through MPI. In spite of a tremendous amount of research in this area, applicability of many of the compiler techniques remains rather limited and the achievable performance enhancement remains less than satisfactory. The main reasons for the restrictive solutions offered by parallelizing compilers is the enormous complexity of the problem. Orchestrating computation and communication by suitable analysis and optimizing their performance through judicious use of underlying architectural features demands a true sophistication on the part of the compiler. It is not even clear whether these complex problems are solvable within the realm of compiler analysis and sophisticated restructuring transformations. Perhaps they are much deeper in nature and go right into the heart of design of parallel algorithms for such an underlying model of computation.

The primary purpose of this monograph is to provide an insight into current approaches and point to potentially open problems that could have an impact. The monograph is organized in terms of issues ranging from programming paradigms (languages) to effective run time systems.

## 1.3 Outline of the Monograph

Language design is largely a matter of legacy and language design for distributed memory systems is no exception to the rule. In section I of the monograph we examine three important approaches (one imperative, one object-oriented and one functional) in this domain that have made a significant impact. The first chapter on HPF 2.0 provides an in-depth view of data parallel language which evolved from Fortran 90. They present HPF 1.0 features such as **BLOCK** distribution and **FORALL** loop as well as new features in HPF 2.0 such as **INDIRECT** distribution and **ON** directive. They also point to the complementary nature of MPI and HPF and discuss features such as **EXTRINSIC** interface mechanism. HPF 2.0 has been a major commercial success with many vendors such as Portland Group and Applied Parallel Research providing highly optimizing compiler support which generates

message passing code. Many research issues especially related to supporting irregular computation could prove valuable to domains such as sparse matrix computation etc. The next chapter on Sisal 90 provides a functional view of implicit parallelism specification and mapping. Shared memory implementation of Sisal is discussed, which involves optimizations such as *update in place copy elimination* etc. Sisal 90 and a distributed memory implementation which uses message passing are also discussed. Finally multi-threaded implementations of Sisal are discussed, with a focus on multi-threaded optimizations. The newer optimizations which perform memory management in hardware through dynamically scheduled multi-threaded code should really prove beneficial for the performance of functional languages (including Sisal) which have an elegant programming model. The next chapter on HPC++ provides an object oriented view as well as details on a library and compiler strategy to support HPC++ level 1 release. The authors discuss interesting features related to multi-threading, barrier synchronization and remote procedure invocation. They also discuss library features that are especially useful for scientific programming. Extensions of this work relating to newer portable languages such as Java is currently an active area of research. We also have a chapter on concurrency models of OO paradigms. The authors specifically address a problem called *inheritance anomaly* which arises when synchronization constraints are implemented within methods of a class and an attempt is made to specialize methods through inheritance mechanisms. They propose a solution to this problem by separating the specification of synchronization from the method specification. The synchronization construct is not a part of the method body and is handled separately. It will be interesting to study the compiler optimizations on this model related to strength reduction of barriers, and issues such as data partitioning vs. barrier synchronizations.

In section II of the monograph, we focus on various analysis techniques. Parallelism detection is very important and the first chapter presents a very interesting comparative study of different loop parallelization algorithms by Allen and Kennedy, Wolf and Lam, Darte and Vivien and by Feautrier. They provide comparisons in terms of their performance (ability to parallelize as well as quality of schedules generated for code generation) as well as complexity. The comparison also focusses on the type of dependence information available. Further extensions could involve run-time parallelization given more precise dependence information. Array data-flow is of utmost importance in optimizations : both sequential as well as parallel. The first chapter on array data-flow analysis examines this problem in detail and presents techniques for exact data flow as well as for approximate data flow. The exact solution is shown for static control programs. Authors also show applications to interprocedural cases and some important parallelization techniques such as privatization. Some interesting extensions could involve run-time data flow analysis. The next chapter discusses interprocedural analysis based on guarded (predicated) array regions. This is a framework based on path-sensitive predi-



cated data-flow which provides summary information. The authors also show application of their work to improve array privatization based on symbolic propagation. Extensions of these to newer object oriented languages such as Java (which have clean class hierarchy and inheritance model) could be interesting since these programs really need such summary MOD information for performing any optimization. We finally present a very important analysis/optimization technique for array privatization. Array privatization involves removing memory-related dependences which have a significant impact on communication optimizations, loop scheduling etc. The authors present a demand-driven data-flow formulation of the problem; an algorithm which performs single pass propagation of symbolic array expressions is also presented. This comprehensive framework implemented in a Polaris compiler is making a significant impact in improving many other related optimizations such as load balancing, communication etc.

The next section is focussed on communication optimization. The communication optimization can be achieved through data (and iteration space) distribution, statically or dynamically. These approaches further classify into data and code alignment or simply iteration space transformations such as in tiling. The communication can also be optimized in data-parallel programs through array region analysis. Finally one could tolerate some communication latency through novel techniques such as multi-threading. We have chapters which cover these broad range of topics about communication in depth.

The first chapter in this section focusses on tiling for cache-coherent multicomputers. This work derives optimal tile parameters for minimal communication in loops with all  $n$  index expressions. The authors introduce a notion of data footprints and tile the iteration spaces so that the volume of communication is minimized. They develop an important lattice theoretic framework to precisely determine the sizes of data footprints which are very valuable not only in tiling but in many array distribution transformations. The next two chapters deal with the important problem of communication free loop partitioning.

The second chapter in this section focusses on comparing different methods of achieving communication-free partitioning for DOALL loops. This chapter discusses several variants of the communication-free partitioning problem involving duplication or non-duplication of data, load balancing of iteration space and aspects such as statement level vs. loop level partitioning. Several aspects such as trading parallelism to avoid inter-loop data distribution are also touched upon. Extending these techniques to broader classes of DOALL loops could enhance their applicability.

The next chapter by Pingali et al. proposes a very interesting framework which first determines a set of constraints on data and loop iteration placement. They then determine which constraints should be left unsatisfied to relax an overconstrained system to find a solution involving a large amount of parallelism. Finally, the remaining constraints are solved for data and code

distribution. The systematic linear algebraic framework improves over many ad-hoc loop partitioning approaches.

These approaches trade parallelism for codes that allow *decoupling* the issues of parallelism and communication by relaxing an appropriate constraint of the problem. However, for many important problems such as image processing applications such a relaxation is not possible. That is, one must resort to a different partitioning solution based on relative costs of communication and computation. In the next chapter, for solving such a problem, a new approach has been proposed to partition iteration space by determining directions which maximally cover the communication by minimally trading parallelism. This approach allows mapping of general medium grained DOALL loops. However, the communication resulting from this iteration space partitioning can not be easily aggregated without sophisticated `pack/unpack` mechanisms present at send/receive ends. Such extensions are desirable since aggregating communication has as significant impact as reducing the volume.

The static data distribution and alignment typically solve the problems of communication on a loop nest by loop nest basis but rarely in an intraprocedural scope. Most of the inter-loop nest level and interprocedural boundaries require dynamic data redistribution. Banerjee et al. develop techniques that can be used to automatically determine which data partitions are most beneficial over specific sections of the program by accounting for redistribution overhead. They determine split points and phases of communication and redistribution are performed at split points.

When communication must take place, it should be optimized. Also, any redundancies must be captured and eliminated. Manish Gupta in the next chapter proposes a comprehensive approach for performing global (interprocedural) communication optimizations such as vectorization, PRE, coalescing, hoisting etc. Such an interprocedural approach to communication optimization is highly profitable in substantially improving the performance. Extending this work to irregular communication could be interesting.

Finally, we present a multi-threaded approach which could hide the communication latency. Two representative applications involving bitonic sort and FFT are chosen and using fine grained multi-threading on EM-X it is shown that multi-threading can substantially help in overlapping computation with communication to hide latencies up to 35 %. These methods could be especially useful for irregular computation.

The final phase of compiling for distributed memory systems involves solving many code generation problems. Code generation problems involve, determining communication generation and doing address calculation to map global references to local ones. The next section deals with these issues. The first chapter presents structures and techniques for communication generation. They focus on issues such as flexible computation partitioning (going beyond owner computes rule), communication adaptation based upon manipulating integer sets through abstract inequalities and control flow simpli-

location based on these. One good property of this work is that it can work with many different front ends (not just data parallel languages) and the code generator has more opportunities to perform low level optimizations due to simplified control flow.

The second chapter discusses basis vector based address calculation mechanisms for efficient traversals of partitioned data. While one important issue of code generation is communication generation, a very important issue is to map global address space to local address space efficiently. The problem is complicated due to data distributions and access strides. Ramanujam et al. present closed form expressions for basis vectors for several cases. Using the closed form expressions for the basis vectors, they derive a non-unimodular linear transformation.

The final section is on supporting task parallelism and dynamic data structures. We also present a run-time system to manage irregular computation. The first chapter by Darbha et al. presents a task scheduling approach that is optimal for many practical cases. The authors evaluate its performance for many practical applications such as the Bellman-Ford algorithm, Cholesky decomposition, the Systolic algorithm etc. They show that schedules generated by their algorithm are optimal for some cases and near optimal for most others. With HPF 2.0 supporting task parallelism, this could open up many new application domains.

The next two chapters describe language supports for dynamic data structures such as pointers in distributed address space. Gupta describes several extensions to C with declarations such as *TREE*, *ARRAY*, *MESH* to declare dynamic data structures. He then describes name generation and distribution strategies for name generation and distribution strategies. Finally he describes support for both regular as well as irregular dynamic structures. The second chapter by Rogers et al. presents an approach followed in their *Olden* project which uses a distributed heap. The remote access is handled by software caching or computation migration. The selection of these mechanisms is done automatically through a compile time heuristic. They provide a data layout annotation to the programmer called *local path lengths* which allows programmers to give hints regarding expected data layout thereby fixing these mechanisms. Both of these chapters provide highly useful insights into supporting dynamic data structures which are very important for scalable domains of computation supported by these machines. Thus, these works should have a significant impact on future scalable applications supported by these systems.

Finally, we present a run-time system called *CHAOS* which provides efficient support for irregular computations. Due to indirection in many sparse matrix computations, the communication patterns are unknown at compile time in these applications. Indirection patterns have to be preprocessed, and the sets of elements to be sent and received by each processor precomputed,

in order to optimize communication. In this work, the authors provide details of efficient run time support for an *inspector-executor* model.

## 1.4 Future Directions

The two important bottlenecks for the use of distributed memory systems are the limited application domains and the fact that the performance is less than satisfactory. The main bottleneck seems to be handling communication. Thus, efficient solutions must be developed. Application domains beyond regular communication can be handled by supporting a general run-time communication model. This run-time communication model must be latency hiding and should give sufficient flexibility to the compiler to defer the hard decisions to run time yet allow static optimizations involving communication motion etc. One of the big problems compilers face is that estimating cost of communication is almost impossible. They can however gauge criticality (or relative importance) of communication. Developing such a model will allow compilers to more effectively deal with issues of *relative* importance between computation and communication and communication and communication.

Probably *the* best reason to use distributed memory systems is to benefit from *scalability* even though application domains and performance might be somewhat weaker. Thus, new research must be done in *scalable code generation*. In other words, as size of the problem and number of processors increase, should there be a change in data/code partition or should it remain the same? What code generation issues are related to this? How could one potentially handle the "hot spots" that inevitably (although at much lower levels than shared memory systems) arise? Can one benefit from the above communication model and dynamic data ownerships discussed earlier?

# Table of Contents

## Preface

Santosh Pande and Dharma P. Agrawal	V
-------------------------------------	---

## Introduction

Santosh Pande and Dharma P. Agrawal	XXI
-------------------------------------	-----

1	Compiling for Distributed Memory Multiprocessors	XXI
1.1	Motivation	XXI
1.2	Complexity	XXII
1.3	Outline of the Monograph	XXII
1.4	Future Directions	XXVII

## Section I : Languages

### Chapter 1. High Performance Fortran 2.0

Ken Kennedy and Charles Koelbel >>&

## Chapter 2. The Sisal Project: Real World Functional Programming

Jean-Luc Gaudiot, Tom DeBoni, John Feo, Wim B hm,

[illegible]

1	Introduction . . . . .	45
2	The Sisal Language: A Short Tutorial . . . . .	46
3	An Early Implementation: The Optimizing Sisal Compiler . . . . .	49
	3.1 Update in Place and Copy Elimination . . . . .	49
	3.2 Build in Place . . . . .	50
	3.3 Reference Counting Optimization . . . . .	51
	3.4 Vectorization . . . . .	51
	3.5 Loop Fusion, Double Buffering Pointer Swap, and Inversion . . . . .	51
4	Sisal90 . . . . .	53
	4.1 The Foreign Language Interface . . . . .	54
5	A Prototype Distributed-Memory SISAL Compiler . . . . .	58
	5.1 Base Compiler . . . . .	59
	5.2 Rectangular Arrays . . . . .	59
	5.3 Block Messages . . . . .	60
	5.4 Multiple Alignment . . . . .	60
	5.5 Results . . . . .	61
	5.6 Further Work . . . . .	62
6	Architecture Support for Multithreaded Execution . . . . .	62
	6.1 Blocking and Non-blocking Models . . . . .	63
	6.2 Code Generation . . . . .	64
	6.3 Summary of Performance Results . . . . .	68
7	Conclusions and Future Research . . . . .	69

## Chapter 3. HPC++ and the HPC++Lib Toolkit

Dennis Gannon, Peter Beckman, Elizabeth Johnson, Todd Green,

[illegible]

1	Introduction . . . . .	73
2	The HPC++ Programming and Execution Model . . . . .	74
	2.1 Level 1 HPC++ . . . . .	75
	2.2 The Parallel Standard Template Library . . . . .	76
	2.3 Parallel Iterators . . . . .	77
	2.4 Parallel Algorithms . . . . .	77
	2.5 Distributed Containers . . . . .	78
3	A Simple Example: The Spanning Tree of a Graph . . . . .	78
4	Multi-threaded Programming . . . . .	82
	4.1 Synchronization . . . . .	84
	4.2 Examples of Multi-threaded Computations . . . . .	92
5	Implementing the HPC++ Parallel Loop Directives . . . . .	96

6	Multi-context Programming and Global Pointers .....	99
	6.1 Remote Function and Member Calls .....	101
	6.2 Using Corba IDL to Generate Proxies .....	103
7	The SPMD Execution Model .....	105
	7.1 Barrier Synchronization and Collective Operations .....	105
8	Conclusion .....	106

## Chapter 4. A Concurrency Abstraction Model for Avoiding Inheritance Anomaly in Object-Oriented Programs

Sandeep Kumar and Dharma P. Agrawal		109
1	Introduction	109
2	Approaches to Parallelism Specification	113
2.1	Issues in Designing a COOPL	113
2.2	Issues in Designing Libraries	114
3	What Is the Inheritance Anomaly?	115
3.1	State Partitioning Anomaly ( <i>SPA</i> )	116
3.2	History Sensitiveness of Acceptable States Anomaly ( <i>HSASA</i> )	118
3.3	State Modification Anomaly ( <i>SMA</i> )	118
3.4	Anomaly A	119
3.5	Anomaly B	120
4	What Is the Reusability of Sequential Classes?	120
5	A Framework for Specifying Parallelism	121
6	Previous Approaches	122
7	The Concurrency Abstraction Model	123
8	The CORE Language	126
8.1	Specifying a Concurrent Region	126
8.2	Defining an AC	126
8.3	Defining a Parallel Block	127
8.4	Synchronization Schemes	129
9	Illustrations	129
9.1	Reusability of Sequential Classes	130
9.2	Avoiding the Inheritance Anomaly	131
10	The Implementation Approach	133
11	Conclusions and Future Directions	134

## Section II : Analysis

## Chapter 5. Loop Parallelization Algorithms

Alain Darte, Yves Robert, and Frédéric Vivien	141
1 Introduction	141
2 Input and Output of Parallelization Algorithms	142
2.1 Input: Dependence Graph	143
2.2 Output: Nested Loops	144







3.4	Profitability of Privatization .....	257
3.5	Last Value Assignment .....	258
4	Demand-Driven Symbolic Analysis .....	261
4.1	Gated Single Assignment .....	263
4.2	Demand-Driven Backward Substitution .....	264
4.3	Backward Substitution in the Presence of Gating Functions ...	266
4.4	Examples of Backward Substitution .....	267
4.5	Bounds of Symbolic Expression .....	269
4.6	Comparison of Symbolic Expressions .....	269
4.7	Recurrence and the $\theta$ Function .....	272
4.8	Bounds of Monotonic Variables .....	273
4.9	Index Array .....	274
4.10	Conditional Data Flow Analysis .....	275
4.11	Implementation and Experiments .....	276
5	Related Work .....	277

## Section III : Communication Optimizations

### Chapter 9. Optimal Tiling for Minimizing Communication in Distributed Shared-Memory Multiprocessors

Anant Agarwal, David Kranz, Rajeev Barua, and Venkat Natarajan ▷▷▷ 285

1	Introduction .....	285
1.1	Contributions and Related Work .....	286
1.2	Overview of the Paper .....	288
2	Problem Domain and Assumptions .....	289
2.1	Program Assumptions .....	289
2.2	System Model .....	291
3	Loop Partitions and Data Partitions .....	292
4	A Framework for Loop and Data Partitioning .....	295
4.1	Loop Tiles in the Iteration Space .....	296
4.2	Footprints in the Data Space .....	298
4.3	Size of a Footprint for a Single Reference .....	300
4.4	Size of the Cumulative Footprint .....	304
4.5	Minimizing the Size of the Cumulative Footprint .....	311
5	General Case of $\mathbf{G}$ .....	314
5.1	$\mathbf{G}$ Is Invertible, but Not Unimodular .....	314
5.2	Columns of $\mathbf{G}$ Are Dependent and the Rows Are Independent .	316
5.3	The Rows of $\mathbf{G}$ Are Dependent .....	316
6	Other System Environments .....	318
6.1	Coherence-Related Cache Misses .....	318
6.2	Effect of Cache Line Size .....	320
6.3	Data Partitioning in Distributed-Memory Multicomputers .....	320



5	Heuristics .....	398
	5.1 Lessons from Some Common Computational Kernels .....	399
	5.2 Implications for Alignment Heuristic .....	402
6	Conclusion .....	402
A	Reducing the Solution Matrix .....	404
	A.1 Unrelated Constraints .....	404
	A.2 General Procedure .....	405
B	A Comment on A[math display="block">\square ne Encoding .....	408

## Chapter 12. A Compilation Method for Communication-Efficient Partitioning of DOALL Loops

Santosh Pande and Tareq Bali >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>	413
1 Introduction.....	413
2 DOALL Partitioning .....	414
2.1 Motivating Example.....	415
2.2 Our Approach .....	419
3 Terms and Definitions .....	421
3.1 Example .....	422
4 Problem .....	423
4.1 Compatibility Subsets .....	423
4.2 Cyclic Directions.....	424
5 Communication Minimization .....	427
5.1 Algorithm : Maximal Compatibility Subsets .....	427
5.2 Algorithm : Maximal Fibonacci Sequence .....	428
5.3 Data Partitioning .....	428
6 Partition Merging .....	429
6.1 Granularity Adjustment .....	431
6.2 Load Balancing .....	431
6.3 Mapping .....	432
7 Example : Texture Smoothing Code .....	432
8 Performance on Cray T3D .....	435
8.1 Conclusions .....	440

## Chapter 13. Compiler Optimization of Dynamic Data Distributions for Distributed-Memory Multicomputers

Daniel J. Palermo, Eugene W. Hodges IV, and Prithviraj Banerjee >>> 445

1	Introduction . . . . .	445
2	Related Work . . . . .	447
3	Dynamic Distribution Selection . . . . .	449
3.1	Motivation for Dynamic Distributions . . . . .	449
3.2	Overview of the Dynamic Distribution Approach . . . . .	450
3.3	Phase Decomposition . . . . .	451
3.4	Phase and Phase Transition Selection . . . . .	457



## Chapter 15. Tolerating Communication Latency through Dynamic Thread Invocation in a Multithreaded Architecture

Andrew Sohn, Yuetsu Kodama, Jui-Yuan Ku, Mitsuhsa Sato, and

Yoshinori Yamaguchi		525
1	Introduction . . . . .	525
2	Multithreading Principles and Its Realization . . . . .	527
	2.1 The Principle . . . . .	527
	2.2 The EM-X Multithreaded Distributed-Memory Multiprocessor .	530
	2.3 Architectural Support for Fine-Grain Multithreading . . . . .	533
3	Designing Multithreaded Algorithms . . . . .	535
	3.1 Multithreaded Bitonic Sorting . . . . .	535
	3.2 Multithreaded Fast Fourier Transform . . . . .	538
4	Overlapping Analysis . . . . .	540
5	Analysis of Switches . . . . .	544
6	Conclusions . . . . .	547

## Section IV : Code Generation

## Chapter 16. Advanced Code Generation for High Performance Fortran

[illegible]

1	Introduction . . . . .	553
2	Background: The Code Generation Problem for HPF . . . . .	556
	2.1 Communication Analysis and Code Generation for HPF . . . . .	556
	2.2 Previous Approaches to Communication Analysis and Code Generation . . . . .	558
3	An Integer Set Framework for Data-Parallel Compilation . . . . .	561
	3.1 Primitive Components of the Framework . . . . .	561
	3.2 Implementation of the Framework . . . . .	562
4	Computation Partitioning . . . . .	565
	4.1 Computation Partitioning Models . . . . .	565
	4.2 Code Generation to Realize Computation Partitions . . . . .	567
5	Communication Code Generation . . . . .	573
	5.1 Communication Generation with Message Vectorization and Coalescing . . . . .	577
	5.2 Recognizing In-Place Communication . . . . .	581
	5.3 Implementing Loop-Splitting for Reducing Communication Overhead . . . . .	582
6	Control Flow Simplification . . . . .	584
	6.1 Motivation . . . . .	584
	6.2 Overview of Algorithm . . . . .	588
	6.3 Evaluation and Discussion . . . . .	589
7	Conclusions . . . . .	590

## Chapter 17. Integer Lattice Based Methods for Local Address Generation for Block-Cyclic Distributions

J. Ramanujam	597
1 Introduction	597
2 Background and Related Work	599
2.1 Related Work on One-Level Mapping	600
2.2 Related Work on Two-Level Mapping	602
3 A Lattice Based Approach for Address Generation	603
3.1 Assumptions	603
3.2 Lattices	604
4 Determination of Basis Vectors	605
4.1 Basis Determination Algorithm	607
4.2 Extremal Basis Vectors	609
4.3 Improvements to the Algorithm for $s < k$	612
4.4 Complexity	613
5 Address Sequence Generation by Lattice Enumeration	614
6 Optimization of Loop Enumeration: GO-LEFT and GO-RIGHT	616
6.1 Implementation	620
7 Experimental Results for One-Level Mapping	620
8 Address Sequence Generation for Two-Level Mapping	626
8.1 Problem Statement	626
9 Algorithms for Two-Level Mapping	628
9.1 <i>Itable</i> : An Algorithm That Constructs a Table of O-sets	629
9.2 Optimization of the <i>Itable</i> Method	631
9.3 Search-Based Algorithms	634
10 Experimental Results for Two-Level Mapping	635
11 Other Problems in Code Generation	638
11.1 Communication Generation	639
11.2 Union and Difference of Regular Sections	640
11.3 Code Generation for Complex Subscripts	640
11.4 Data Structures for Runtime Efficiency	640
11.5 Array Redistribution	641
12 Summary and Conclusions	641

## Section V : Task Parallelism, Dynamic Data Structures and Run Time Systems

## Chapter 18. A Duplication Based Compile Time Scheduling Method for Task Parallelism

Sekhar Darbha and Dharma P. Agrawal >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>> 649

1	Introduction . . . . .	649
2	STDS Algorithm . . . . .	652
	2.1 Complexity Analysis . . . . .	663

3	Illustration of the STDS Algorithm .....	664
4	Performance of the STDS Algorithm .....	670
4.1	CRC Is Satisfied .....	670
4.2	Application of Algorithm for Random Data .....	672
4.3	Application of Algorithm to Practical DAGs .....	674
4.4	Scheduling of Diamond DAGs .....	675
4.5	Comparison with Other Algorithms .....	680
5	Conclusions .....	680

**Chapter 19. SPMD Execution in the Presence of Dynamic Data Structures**

Rajiv Gupta	.....	683
1	Introduction .....	683
2	Language Support for Regular Data Structures .....	684
2.1	Processor Structures .....	685
2.2	Dynamic Data Structures .....	685
2.3	Name Generation and Distribution Strategies .....	688
2.4	Examples .....	689
3	Compiler Support for Regular Data Structures .....	693
3.1	Representing Pointers and Data Structures .....	693
3.2	Translation of Pointer Operations .....	694
4	Supporting Irregular Data Structures .....	703
5	Compile-Time Optimizations .....	705
6	Related Work .....	706

**Chapter 20. Supporting Dynamic Data Structures with Olden**

Martin C. Carlisle and Anne Rogers	.....	709
1	Introduction .....	709
2	Programming Model .....	711
2.1	Programming Language .....	711
2.2	Data Layout .....	711
2.3	Marking Available Parallelism .....	714
3	Execution Model .....	715
3.1	Handling Remote References .....	715
3.2	Introducing Parallelism .....	718
3.3	A Simple Example .....	719
4	Selecting Between Mechanisms .....	722
4.1	Using Local Path Lengths .....	723
4.2	Update Matrices .....	724
4.3	The Heuristic .....	726



