Darte / Robert / Vivien

Scheduling and Automatic Parallelization

Readership This book is devoted to the study of compiler transformations that are needed to expose the parallelism hiddenin a program. This book is notan introductory book to parallel processing, nor is it an introductory book to parallelizing compilers. Weassume thatreaders are familiar withthebooks High Performance Compilers for Parallel Computingby Wolfe [121] and Super compilers for Parallel and Vector Computers by Zima and Chapman [125], and that they want to know more about scheduling transformations. In this book we describe both task graph scheduling and loop nest scheduling. Taskgraphschedulingaims atexecuting tasks linked by prece dence constraints; it is a run-time activity. Loop nest scheduling aims at ex ecutingstatementinstances linked bydata dependences;it is a compile-time activity. We are mostly interested in loop nestscheduling,butwe also deal with task graph scheduling for two main reasons: (i) Beautiful algorithms and heuristics have been reported in the literature recently; and (ii) Several graphscheduling, like list scheduling, are the basis techniques used in task ofthe loop transformations implemented in loop nest scheduling. As for loop nest scheduling our goal is to capture in a single place the fantastic developments of the last decade or so. Dozens of loop trans formations have been introduced (loop interchange, skewing, fusion, dis tribution, etc.) before a unifying theory emerged. The theory builds upon the pioneering papers of Karp, Miller, and Winograd [65] and of Lam port [75], and it relies on sophisticated mathematical tools (unimodular transformations, parametric integer linear programming, Hermite decom position, Smithdecomposition, etc.).

Readership This book is devoted to the study of compiler transformations that are needed to expose the parallelism hiddenin a program. This book is notan introductory book to parallel processing, nor is it an introductory book to parallelizing compilers. Weassume thatreaders are familiar withthebooks High Performance Compilers for Parallel Computingby Wolfe [121] and Super compilers for Parallel and Vector Computers by Zima and Chapman [125], and that they want to know more about scheduling transformations. In this book we describe both task graph scheduling and loop nest scheduling. Taskgraphschedulingaims atexecuting tasks linked by prece dence constraints; it is a run-time activity. Loop nest scheduling aims at ex ecutingstatementinstances linked bydata dependences;it is a compile-time activity. We are mostly interested in loop nestscheduling,butwe also deal with task graph scheduling for two main reasons: (i) Beautiful algorithms and heuristics have been reported in the literature recently; and (ii) Several graphscheduling, like list scheduling, are the basis techniques used in task ofthe loop transformations implemented in loop nest scheduling. As for loop nest scheduling our goal is to capture in a single place the fantastic developments of the last decade or so. Dozens of loop trans formations have been introduced (loop interchange, skewing, fusion, dis tribution, etc.) before a unifying theory emerged. The theory builds upon the pioneering papers of Karp, Miller, and Winograd [65] and of Lam port [75], and it relies on sophisticated mathematical tools (unimodular transformations, parametric integer linear programming, Hermite decom position, Smithdecomposition, etc.).

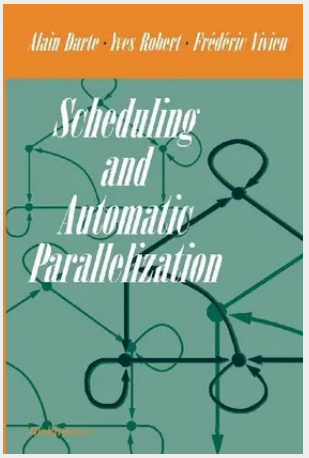

106,99 €
99,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9780817641498
Medium: Buch
ISBN: 978-0-8176-4149-8
Verlag: Birkhäuser
Erscheinungstermin: 30.03.2000
Sprache(n): Englisch
Auflage: 2000
Produktform: Gebunden
Gewicht: 1680 g
Seiten: 264
Format (B x H): 183 x 260 mm

Kundenservice Verlag Dr. Otto Schmidt KG
Neumannstraße 10, 40235 Düsseldorf | kundenservice@otto-schmidt.de | 0221-93738-997 (Inland)