

Preface

This book has been distilled over more than a decade. It started as part of a more ambitious project that Utpal Banerjee and Alex Nicolau discussed over a rather ample dinner around the turn of the century. The book, as originally planned, would have consisted of several volumes culminating in a unification of all parallelizing compiler technology in an all-encompassing framework/formalization. Over time, the project gained a sense of proportion, and was reduced to the more realistic goal of a thorough description of Instruction Level Parallelism and a more complete, unifying formalization of the techniques within. After an initial period of re-education about the subject, gathering and agreement on the most representative material to include followed by (very) long, detailed discussions to ascertain that both authors had a common understanding of the material, we finally started writing, some 4–5 years after the original discussion of the project. This writing phase took an additional two years, and resulted in four chapters being written in varying degrees of detail (Basic Block techniques, Definitions, Architectures in close-to, but not quite ready for publication form, and Trace Scheduling re-formalized by Utpal with some interesting new results in the bargain). The rest of the chapters ranged from detailed outlines with some sections in rough form, to little more than titles and subsection titles.

We then decided to add an extra author, who with fresh energy and optimism, would help write the remaining chapters and have the stamina to carefully read, fix, expand and homogenize the whole text. In this way, Arun Kejariwal joined the team, and after six months or so of reading, he quickly (one year) produced important additions to the work, including (but not limited to) an encyclopedic collection of references, Further Study sections for each chapter and the preliminary work on the Loop Scheduling chapters. At that point, while the

manuscript had grown to about 180 pages, it was not in what any of the authors would have considered close to publication-ready. We invited Alex Aiken to join the project owing to his deep understanding and contributions to research in Instruction Level Parallelism. Much to our surprise, Alex graciously agreed to join and not content to just help complete the Loop Scheduling chapters, which we saw as the minimal required effort to make the manuscript publishable, he volunteered to perform a very thorough pass on the whole manuscript and rewrite/revise/unify/amplify/simplify as needed. The resulting book is a much better end product for Alex's efforts. Indeed, the resulting book, after a total of approximately 15 years, is a team effort that we can all be proud of, and, realistically, is also much better than any of us could have produced individually. We can only hope that the reader will agree and find the book useful.

A. AIKEN, U. BANERJEE, A. KEJARIWAL , A. NICOLAU

Instruction Level Parallelism

Aiken, A.; Banerjee, U.; Kejariwal, A.; Nicolau, A.

2016, XXI, 255 p. 78 illus., 30 illus. in color., Hardcover

ISBN: 978-1-4899-7795-3