


Data Processing on FPGAs

Roughly a decade ago, power consumption and heat dissipation concerns forced the semiconductor industry to radically change its course, shifting from sequential to parallel computing. Unfortunately, improving performance of applications has now become much more difficult than in the good old days of frequency scaling. This is also affecting databases and data processing applications in general, and has led to the popularity of so-called data appliances—specialized data processing engines, where software and hardware are sold together in a closed box. Field-programmable gate arrays (FPGAs) increasingly play an important role in such systems. FPGAs are attractive because the performance gains of specialized hardware can be significant, while power consumption is much less than that of commodity processors. On the other hand, FPGAs are way more flexible than hard-wired circuits (ASICs) and can be integrated into complex systems in many different ways, e.g., directly in the network for a high-frequency trading application. This book gives an introduction to FPGA technology targeted at a database audience. In the first few chapters, we explain in detail the inner workings of FPGAs. Then we discuss techniques and design patterns that help mapping algorithms to FPGA hardware so that the inherent parallelism of these devices can be leveraged in an optimal way. Finally, the book will illustrate a number of concrete examples that exploit different advantages of FPGAs for data processing. Table of Contents: Preface / Introduction / A Primer in Hardware Design / FPGAs / FPGA Programming Models / Data Stream Processing / Accelerated DB Operators / Secure Data Processing / Conclusions / Bibliography / Authors' Biographies / Index

Roughly a decade ago, power consumption and heat dissipation concerns forced the semiconductor industry to radically change its course, shifting from sequential to parallel computing. Unfortunately, improving performance of applications has now become much more difficult than in the good old days of frequency scaling. This is also affecting databases and data processing applications in general, and has led to the popularity of so-called data appliances—specialized data processing engines, where software and hardware are sold together in a closed box. Field-programmable gate arrays (FPGAs) increasingly play an important role in such systems. FPGAs are attractive because the performance gains of specialized hardware can be significant, while power consumption is much less than that of commodity processors. On the other hand, FPGAs are way more flexible than hard-wired circuits (ASICs) and can be integrated into complex systems in many different ways, e.g., directly in the network for a high-frequency trading application. This book gives an introduction to FPGA technology targeted at a database audience. In the first few chapters, we explain in detail the inner workings of FPGAs. Then we discuss techniques and design patterns that help mapping algorithms to FPGA hardware so that the inherent parallelism of these devices can be leveraged in an optimal way. Finally, the book will illustrate a number of concrete examples that exploit different advantages of FPGAs for data processing. Table of Contents: Preface / Introduction / A Primer in Hardware Design / FPGAs / FPGA Programming Models / Data Stream Processing / Accelerated DB Operators / Secure Data Processing / Conclusions / Bibliography / Authors' Biographies / Index


35,30 €

32,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage
Artikelnummer: 9783031007217

Medium: Buch

ISBN: 978-3-031-00721-7

Verlag: Springer

Erscheinungstermin: 27.06.2013

Sprache(n): Englisch

Auflage: Erscheinungsjahr 2013

Serie: Synthesis Lectures on Data Management

Produktform: Kartoniert

Gewicht: 241 g

Seiten: 104

Format (B x H): 191 x 235 mm
