


Parallel Machines: Parallel Machine Languages

The Emergence of Hybrid Dataflow Computer Architectures

It is universally accepted today that parallel processing is here to stay but that software for parallel machines is still difficult to develop. However, there is little recognition of the fact that changes in processor architecture can significantly ease the development of software. In the seventies the availability of processors that could address a large name space directly, eliminated the problem of name management at one level and paved the way for the routine development of large programs. Similarly, today, processor architectures that can facilitate cheap synchronization and provide a global address space can simplify compiler development for parallel machines. If the cost of synchronization remains high, the programming of parallel machines will remain significantly less abstract than programming sequential machines. In this monograph Bob Iannucci presents the design and analysis of an architecture that can be a better building block for parallel machines than any von Neumann processor. There is another very interesting motivation behind this work. It is rooted in the long and venerable history of dataflow graphs as a formalism for expressing parallel computation. The field has bloomed since 1974, when Dennis and Misunas proposed a truly novel architecture using dataflow graphs as the parallel machine language. The novelty and elegance of dataflow architectures has, however, also kept us from asking the real question: "What can dataflow architectures buy us that von Neumann architectures can't?" In the following I explain in a round about way how Bob and I arrived at this question.

It is universally accepted today that parallel processing is here to stay but that software for parallel machines is still difficult to develop. However, there is little recognition of the fact that changes in processor architecture can significantly ease the development of software. In the seventies the availability of processors that could address a large name space directly, eliminated the problem of name management at one level and paved the way for the routine development of large programs. Similarly, today, processor architectures that can facilitate cheap synchronization and provide a global address space can simplify compiler development for parallel machines. If the cost of synchronization remains high, the programming of parallel machines will remain significantly less abstract than programming sequential machines. In this monograph Bob Iannucci presents the design and analysis of an architecture that can be a better building block for parallel machines than any von Neumann processor. There is another very interesting motivation behind this work. It is rooted in the long and venerable history of dataflow graphs as a formalism for expressing parallel computation. The field has bloomed since 1974, when Dennis and Misunas proposed a truly novel architecture using dataflow graphs as the parallel machine language. The novelty and elegance of dataflow architectures has, however, also kept us from asking the real question: "What can dataflow architectures buy us that von Neumann architectures can't?" In the following I explain in a round about way how Bob and I arrived at this question.



160,49 €

149,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9780792391012

Medium: Buch

ISBN: 978-0-7923-9101-2

Verlag: Springer

Erscheinungstermin: 31.05.1990

Sprache(n): Englisch

Auflage: Erscheinungsjahr 1990

Serie: The Springer International Series in Engineering and Computer Science

Produktform: Gebunden

Gewicht: 1110 g

Seiten: 198

Format (B x H): 160 x 241 mm

