

Contents

| | |
|--|----------|
| <i>Foreword to the First Edition</i> | xv |
| <i>Preface to the Second Edition</i> | xvii |
| <i>Preface to the First Edition</i> | xxiii |
| Part I Theory of Parallelism | 1 |
| 1. Parallel Computer Models | 3 |
| 1.1 The State of Computing | 3 |
| 1.1.1 Computer Development Milestones | 3 |
| 1.1.2 Elements of Modern Computers | 6 |
| 1.1.3 Evolution of Computer Architecture | 8 |
| 1.1.4 System Attributes to Performance | 12 |
| 1.2 Multiprocessors and Multicomputers | 17 |
| 1.2.1 Shared-Memory Multiprocessors | 17 |
| 1.2.2 Distributed-Memory Multicomputers | 22 |
| 1.2.3 A Taxonomy of MIMD Computers | 24 |
| 1.3 Multivector and SIMD Computers | 25 |
| 1.3.1 Vector Supercomputers | 25 |
| 1.3.2 SIMD Supercomputers | 27 |
| 1.4 PRAM and VLSI Models | 29 |
| 1.4.1 Parallel Random-Access Machines | 30 |
| 1.4.2 VLSI Complexity Model | 33 |
| 1.5 Architectural Development Tracks | 36 |
| 1.5.1 Multiple-Processor Tracks | 36 |
| 1.5.2 Multivector and SIMD Tracks | 38 |
| 1.5.3 Multithreaded and Dataflow Tracks | 39 |
| <i>Summary</i> | 40 |
| <i>Exercises</i> | 41 |

| | |
|--|-----------|
| 2. Program and Network Properties | 44 |
| 2.1 Conditions of Parallelism | 44 |
| 2.1.1 Data and Resource Dependences | 44 |
| 2.1.2 Hardware and Software Parallelism | 49 |
| 2.1.3 The Role of Compilers | 52 |
| 2.2 Program Partitioning and Scheduling | 52 |
| 2.2.1 Grain Sizes and Latency | 52 |
| 2.2.2 Grain Packing and Scheduling | 55 |
| 2.2.3 Static Multiprocessor Scheduling | 58 |
| 2.3 Program Flow Mechanisms | 61 |
| 2.3.1 Control Flow Versus Data Flow | 61 |
| 2.3.2 Demand-Driven Mechanisms | 65 |
| 2.3.3 Comparison of Flow Mechanisms | 65 |
| 2.4 System Interconnect Architectures | 66 |
| 2.4.1 Network Properties and Routing | 67 |
| 2.4.2 Static Connection Networks | 70 |
| 2.4.3 Dynamic Connection Networks | 77 |
| <i>Summary</i> | 83 |
| <i>Exercises</i> | 84 |
| 3. Principles of Scalable Performance | 89 |
| 3.1 Performance Metrics and Measures | 89 |
| 3.1.1 Parallelism Profile in Programs | 89 |
| 3.1.2 Mean Performance | 92 |
| 3.1.3 Efficiency, Utilization, and Quality | 95 |
| 3.1.4 Benchmarks and Performance Measures | 97 |
| 3.2 Parallel Processing Applications | 99 |
| 3.2.1 Massive Parallelism for Grand Challenges | 99 |
| 3.2.2 Application Models of Parallel Computers | 102 |
| 3.2.3 Scalability of Parallel Algorithms | 104 |
| 3.3 Speedup Performance Laws | 108 |
| 3.3.1 Amdahl's Law for a Fixed Workload | 108 |
| 3.3.2 Gustafson's Law for Scaled Problems | 111 |
| 3.3.3 Memory-Bounded Speedup Model | 112 |
| 3.4 Scalability Analysis and Approaches | 116 |
| 3.4.1 Scalability Metrics and Goals | 116 |
| 3.4.2 Evolution of Scalable Computers | 120 |
| 3.4.3 Research Issues and Solutions | 123 |
| <i>Summary</i> | 125 |
| <i>Exercises</i> | 125 |

| | |
|---|------------|
| Part II Hardware Technologies | 131 |
| 4. Processors and Memory Hierarchy | 133 |
| 4.1 Advanced Processor Technology 133 | |
| 4.1.1 Design Space of Processors 133 | |
| 4.1.2 Instruction-Set Architectures 137 | |
| 4.1.3 CISC Scalar Processors 139 | |
| 4.1.4 RISC Scalar Processors 143 | |
| 4.2 Superscalar and Vector Processors 150 | |
| 4.2.1 Superscalar Processors 150 | |
| 4.2.2 The VLIW Architecture 154 | |
| 4.2.3 Vector and Symbolic Processors 156 | |
| 4.3 Memory Hierarchy Technology 160 | |
| 4.3.1 Hierarchical Memory Technology 160 | |
| 4.3.2 Inclusion, Coherence, and Locality 161 | |
| 4.3.3 Memory Capacity Planning 165 | |
| 4.4 Virtual Memory Technology 167 | |
| 4.4.1 Virtual Memory Models 167 | |
| 4.4.2 TLB, Paging, and Segmentation 169 | |
| 4.4.3 Memory Replacement Policies 174 | |
| Summary 177 | |
| Exercises 178 | |
| 5. Bus, Cache, and Shared Memory | 182 |
| 5.1 Bus Systems 182 | |
| 5.1.1 Backplane Bus Specification 182 | |
| 5.1.2 Addressing and Timing Protocols 184 | |
| 5.1.3 Arbitration, Transaction, and Interrupt 186 | |
| 5.1.4 IEEE Futurebus ⁺ and other Standards 189 | |
| 5.2 Cache Memory Organizations 192 | |
| 5.2.1 Cache Addressing Models 192 | |
| 5.2.2 Direct Mapping and Associative Caches 195 | |
| 5.2.3 Set-Associative and Sector Caches 198 | |
| 5.2.4 Cache Performance Issues 202 | |
| 5.3 Shared-Memory Organizations 205 | |
| 5.3.1 Interleaved Memory Organization 205 | |
| 5.3.2 Bandwidth and Fault Tolerance 208 | |
| 5.3.3 Memory Allocation Schemes 210 | |
| 5.4 Sequential and Weak Consistency Models 213 | |
| 5.4.1 Atomicity and Event Ordering 213 | |
| 5.4.2 Sequential Consistency Model 217 | |

| | | | |
|-----------------|--|-----|------------|
| 5.4.3 | Weak Consistency Models | 218 | |
| | <i>Summary</i> | 221 | |
| | <i>Exercises</i> | 222 | |
| 6. | Pipelining and Superscalar Techniques | | 227 |
| 6.1 | Linear Pipeline Processors | 227 | |
| 6.1.1 | Asynchronous and Synchronous Models | 227 | |
| 6.1.2 | Clocking and Timing Control | 229 | |
| 6.1.3 | Speedup, Efficiency, and Throughput | 229 | |
| 6.2 | Nonlinear Pipeline Processors | 232 | |
| 6.2.1 | Reservation and Latency Analysis | 232 | |
| 6.2.2 | Collision-Free Scheduling | 235 | |
| 6.2.3 | Pipeline Schedule Optimization | 237 | |
| 6.3 | Instruction Pipeline Design | 240 | |
| 6.3.1 | Instruction Execution Phases | 240 | |
| 6.3.2 | Mechanisms for Instruction Pipelining | 243 | |
| 6.3.3 | Dynamic Instruction Scheduling | 247 | |
| 6.3.4 | Branch Handling Techniques | 250 | |
| 6.4 | Arithmetic Pipeline Design | 255 | |
| 6.4.1 | Computer Arithmetic Principles | 255 | |
| 6.4.2 | Static Arithmetic Pipelines | 257 | |
| 6.4.3 | Multifunctional Arithmetic Pipelines | 263 | |
| 6.5 | Superscalar Pipeline Design | 266 | |
| | <i>Summary</i> | 273 | |
| | <i>Exercises</i> | 274 | |
| Part III | Parallel and Scalable Architectures | | 279 |
| 7. | Multiprocessors and Multicomputers | | 281 |
| 7.1 | Multiprocessor System Interconnects | 281 | |
| 7.1.1 | Hierarchical Bus Systems | 282 | |
| 7.1.2 | Crossbar Switch and Multiport Memory | 286 | |
| 7.1.3 | Multistage and Combining Networks | 290 | |
| 7.2 | Cache Coherence and Synchronization Mechanisms | 296 | |
| 7.2.1 | The Cache Coherence Problem | 296 | |
| 7.2.2 | Snoopy Bus Protocols | 299 | |
| 7.2.3 | Directory-Based Protocols | 303 | |
| 7.2.4 | Hardware Synchronization Mechanisms | 308 | |
| 7.3 | Three Generations of Multicomputers | 312 | |
| 7.3.1 | Design Choices in the Past | 312 | |

| | | |
|-----------|--|------------|
| 7.3.2 | Present and Future Development | 314 |
| 7.3.3 | The Intel Paragon System | 316 |
| 7.4 | Message-Passing Mechanisms | 318 |
| 7.4.1 | Message-Routing Schemes | 319 |
| 7.4.2 | Deadlock Virtual Channels | 322 |
| 7.4.3 | Flow Control Strategies | 324 |
| 7.4.4 | Multicast Routing Algorithms | 329 |
| | <i>Summary</i> | 334 |
| | <i>Exercises</i> | 335 |
| 8. | Multivector and SIMD Computers | 341 |
| 8.1 | Vector Processing Principles | 341 |
| 8.1.1 | Vector Instruction Types | 341 |
| 8.1.2 | Vector-Access Memory Schemes | 345 |
| 8.1.3 | Early Supercomputers | 347 |
| 8.2 | Multivector Multiprocessors | 352 |
| 8.2.1 | Performance-Directed Design Rules | 352 |
| 8.2.2 | Cray Y-MP, C-90, and MPP | 356 |
| 8.2.3 | Fujitsu VP2000 and VPP500 | 362 |
| 8.2.4 | Mainframes and Minisupercomputers | 365 |
| 8.3 | Compound Vector Processing | 372 |
| 8.3.1 | Compound Vector Operations | 372 |
| 8.3.2 | Vector Loops and Chaining | 374 |
| 8.3.3 | Multipipeline Networking | 378 |
| 8.4 | SIMD Computer Organizations | 382 |
| 8.4.1 | Implementation Models | 383 |
| 8.4.2 | The CM-2 Architecture | 385 |
| 8.4.3 | The MasPar MP-1 Architecture | 388 |
| 8.5 | The Connection Machine CM-5 | 392 |
| 8.5.1 | A Synchronized MIMD Machine | 392 |
| 8.5.2 | The CM-5 Network Architecture | 395 |
| 8.5.3 | Control Processors and Processing Nodes | 397 |
| 8.5.4 | Interprocessor Communications | 399 |
| | <i>Summary</i> | 403 |
| | <i>Exercises</i> | 404 |
| 9. | Scalable, Multithreaded, and Dataflow Architectures | 408 |
| 9.1 | Latency-Hiding Techniques | 408 |
| 9.1.1 | Shared Virtual Memory | 408 |
| 9.1.2 | Prefetching Techniques | 412 |
| 9.1.3 | Distributed Coherent Caches | 413 |

| | | |
|-------|--|-----|
| 9.1.4 | Scalable Coherence Interface | 415 |
| 9.1.5 | Relaxed Memory Consistency | 418 |
| 9.2 | Principles of Multithreading | 421 |
| 9.2.1 | Multithreading Issues and Solutions | 421 |
| 9.2.2 | Multiple-Context Processors | 426 |
| 9.2.3 | Multidimensional Architectures | 431 |
| 9.3 | Fine-Grain Multicomputers | 434 |
| 9.3.1 | Fine-Grain Parallelism | 434 |
| 9.3.2 | The MIT J-Machine | 435 |
| 9.3.3 | The Caltech Mosaic C | 442 |
| 9.4 | Scalable and Multithreaded Architectures | 444 |
| 9.4.1 | The Stanford Dash Multiprocessor | 444 |
| 9.4.2 | The Kendall Square Research KSR-1 | 448 |
| 9.4.3 | The Tera Multiprocessor System | 452 |
| 9.5 | Dataflow and Hybrid Architectures | 458 |
| 9.5.1 | The Evolution of Dataflow Computers | 458 |
| 9.5.2 | The ETL/EM-4 in Japan | 461 |
| 9.5.3 | The MIT/Motorola *T Prototype | 463 |
| | <i>Summary</i> | 465 |
| | <i>Exercises</i> | 466 |

Part IV Software for Parallel Programming

471

10. Parallel Models, Languages, and Compilers

473

| | | |
|--------|---|-----|
| 10.1 | Parallel Programming Models | 473 |
| 10.1.1 | Shared-Variable Model | 473 |
| 10.1.2 | Message-Passing Model | 477 |
| 10.1.3 | Data-Parallel Model | 479 |
| 10.1.4 | Object-Oriented Model | 481 |
| 10.1.5 | Functional and Logic Models | 483 |
| 10.2 | Parallel Languages and Compilers | 484 |
| 10.2.1 | Language Features for Parallelism | 485 |
| 10.2.2 | Parallel Language Constructs | 487 |
| 10.2.3 | Optimizing Compilers for Parallelism | 488 |
| 10.3 | Dependence Analysis of Data Arrays | 491 |
| 10.3.1 | Iteration Space and Dependence Analysis | 491 |
| 10.3.2 | Subscript Separability and Partitioning | 494 |
| 10.3.3 | Categorized Dependence Tests | 496 |
| 10.4 | Code Optimization and Scheduling | 501 |
| 10.4.1 | Scalar Optimization with Basic Blocks | 501 |

| | | |
|---------------|--|------------|
| 10.4.2 | Local and Global Optimizations | 505 |
| 10.4.3 | Vectorization and Parallelization Methods | 508 |
| 10.4.4 | Code Generation and Scheduling | 514 |
| 10.4.5 | Trace Scheduling Compilation | 518 |
| 10.5 | Loop Parallelization and Pipelining | 520 |
| 10.5.1 | Loop Transformation Theory | 520 |
| 10.5.2 | Parallelization and Wavefronting | 523 |
| 10.5.3 | Tiling and Localization | 526 |
| 10.5.4 | Software Pipelining | 531 |
| | <i>Summary</i> | 533 |
| | <i>Exercises</i> | 534 |
| 11. | Parallel Program Development and Environments | 537 |
| 11.1 | Parallel Programming Environments | 537 |
| 11.1.1 | Software Tools and Environments | 537 |
| 11.1.2 | Y-MP, Paragon and CM-5 Environments | 541 |
| 11.1.3 | Visualization and Performance Tuning | 543 |
| 11.2 | Synchronization and Multiprocessing Modes | 545 |
| 11.2.1 | Principles of Synchronization | 545 |
| 11.2.2 | Multiprocessor Execution Modes | 547 |
| 11.2.3 | Multitasking on Cray Multiprocessors | 548 |
| 11.3 | Shared-Variable Program Structures | 552 |
| 11.3.1 | Locks for Protected Access | 553 |
| 11.3.2 | Semaphores and Applications | 556 |
| 11.3.3 | Monitors and Applications | 559 |
| 11.4 | Message-Passing Program Development | 562 |
| 11.4.1 | Distributing the Computation | 562 |
| 11.4.2 | Synchronous Message Passing | 564 |
| 11.4.3 | Asynchronous Message Passing | 565 |
| 11.5 | Mapping Programs onto Multicomputers | 566 |
| 11.5.1 | Domain Decomposition Techniques | 566 |
| 11.5.2 | Control Decomposition Techniques | 570 |
| 11.5.3 | Heterogeneous Processing | 573 |
| | <i>Summary</i> | 577 |
| | <i>Exercises</i> | 578 |
| Part V | Instruction and System Level Parallelism | 583 |
| 12. | Instruction Level Parallelism | 585 |
| 12.1 | Introduction | 585 |
| 12.2 | Basic Design Issues | 587 |

| | | |
|------------|---|------------|
| 12.3 | Problem Definition | 589 |
| 12.4 | Model of a Typical Processor | 594 |
| 12.5 | Compiler-detected Instruction Level Parallelism | 598 |
| 12.6 | Operand Forwarding | 602 |
| 12.7 | Reorder Buffer | 605 |
| 12.8 | Register Renaming | 607 |
| 12.9 | Tomasulo's Algorithm | 610 |
| 12.10 | Branch Prediction | 615 |
| 12.11 | Limitations in Exploiting Instruction Level Parallelism | 618 |
| 12.12 | Thread Level Parallelism | 623 |
| | <i>Summary</i> | 624 |
| | <i>Exercises</i> | 626 |
| 13. | Trends in Parallel Systems | 629 |
| 13.1 | Brief Overview of Technology | 629 |
| 13.1.1 | Semiconductor Technology | 630 |
| 13.1.2 | Display Technology | 632 |
| 13.1.3 | Storage Technology | 633 |
| 13.1.4 | Interconnect and Network Technologies | 635 |
| 13.2 | Forms of Parallelism | 639 |
| 13.2.1 | Structural Parallelism versus Instruction Level Parallelism | 640 |
| 13.2.2 | A Simple Parallel Computation | 642 |
| 13.2.3 | Parallel Algorithms | 646 |
| 13.2.4 | Stream Processing | 651 |
| 13.3 | Case Studies | 654 |
| 13.3.1 | Cray Line of Computer Systems | 654 |
| 13.3.2 | PowerPC Architecture, IBM Power7 & Blue Gene | 657 |
| 13.3.3 | Tilera's TILE64 System | 658 |
| 13.3.4 | Sun UltraSparc T2 Processor | 660 |
| 13.3.5 | AMD Opteron | 662 |
| 13.3.6 | Intel Pentium Processors | 663 |
| 13.4 | Parallel Programming Models and Languages | 665 |
| 13.4.1 | Parallel Programming Language Chapel | 665 |
| 13.4.2 | Function Libraries for Parallel Programming | 669 |
| | <i>Summary</i> | 673 |
| | <i>Exercises</i> | 674 |
| | Answers to Selected Exercises | 679 |
| | Bibliography | 687 |
| | Index | 707 |

Foreword to the First Edition

Kai Hwang has introduced the issues in designing and using high performance parallel computers at a time when a plethora of scalable computers utilizing commodity microprocessors offer higher peak performance than traditional vector supercomputers. These new machines, their operating environments including the operating system and languages, and the programs to effectively utilize them are introducing more rapid changes for researchers, builders, and users than at any time in the history of computer structures.

For the first time since the introduction of Cray 1 vector processor in 1975, it may again be necessary to change and evolve the programming paradigm—provided that massively parallel computers can be shown to be useful outside of research on massive parallelism. Vector processors required modest data parallelism and these operations have been reflected either explicitly in Fortran programs or implicitly with the need to evolve Fortran (e.g., Fortran 90) to build in vector operations.

So far, the main line of supercomputing as measured by the usage (hours, jobs, number of programs, program portability) has been the shared memory, vector multiprocessor as pioneered by Cray Research. Fujitsu, IBM, Hitachi, and NEC all produce computers of this type. In 1993, the Cray C90 supercomputer delivers a peak of 16 billion floating-point operations per second (a Gigaflops) with 16 processors and costs about \$30 million, providing roughly 500 floating-point operations per second per dollar.

In contrast, massively parallel computers introduced in the early 1990s are nearly all based on utilizing the same powerful, RISC-based, CMOS microprocessors that are used in workstations. These scalar processors provide a peak of ≈ 100 million floatingpoint operations per second and cost \$20 thousand, providing an order of magnitude more peak per dollar (5000 flops per dollar). Unfortunately, to obtain peak power requires large-scale problems that can require $O(n^3)$ operations over supers, and this significantly increases the running time when peak power is the goal.

The multicomputer approach interconnects computers built from microprocessors through high-bandwidth switches that introduce latency. Programs are written in either an evolved parallel data model utilizing Fortran or as independent programs that communicate by passing messages. The book describes a variety of multicomputers including Thinking Machines' CM5, the first computer announced that could reach a teraflops using 8K independent computer nodes, each of which can deliver 128 Mflops utilizing four 32-Mflops floating-point units.

The architecture research trend is toward scalable, shared-memory multiprocessors in order to handle general workloads ranging from technical to commercial tasks and workloads, negate the need to explicitly pass messages for communication, and provide memory addressed accessing. KSR's scalable multiprocessor and Stanford's Dash prototype have proven that such machines are possible.

The author starts by positing a framework based on evolution that outlines the main approaches to designing computer structures. He covers both the scaling of computers and workloads, various multiprocessors, vector processing, multicomputers, and emerging scalable or multithreaded multiprocessors. The final three chapters describe parallel programming techniques and discuss the host operating environment necessary to utilize these new computers.

The book provides case studies of both industrial and research computers, including the Illinois Cedar, Intel Paragon, TMC CM-2, MasPar M1, TMC CM-5, Cray Y-MP, C-90, and Cray MPP, Fujitsu VP2000 and VPP500, NEC SX, Stanford Dash, KSR-1, MIT J-Machine, MIT *T, ETL EM-4, Caltech Mosaic C, and Tera Computer.

The book presents a balanced treatment of the theory, technology, architecture, and software of advanced computer systems. The emphasis on parallelism, scalability, and programmability makes this book rather unique and educational.

I highly recommend Dr. Hwang's timely book. I believe it will benefit many readers and be a fine reference.

C. Gordon Bell

Preface to the Second Edition

Technologies underlying computer architecture—*VLSI*, communication, storage, graphics and others—have undergone tremendous advances over the last two decades. High performance computing, which was earlier reserved for very large research projects, has today become far more affordable, and the range of applications of computer systems has expanded enormously. The digital revolution in communication technology, the *world wide web*, and increased user awareness of the huge power of computers have been major contributors to the enormous growth in computer applications.

Against the backdrop of this scenario, the revision of a well-accepted textbook on computer architecture is necessary. The first edition of *Advanced Computer Architecture* by Kai Hwang has been hugely successful with teachers and students—the basic reason being that it explains simply and effectively the key principles and techniques of high performance computing systems.

New to this Edition

The first edition of the book provides a faithful and very useful record of developments in computer architecture in its creative and formative years, beyond the original uniprocessor systems based on the von Neumann model. This was the exciting period when many innovative research ideas were tried out which have had a long-lasting impact on the field. This record of developments remains valuable even today—almost two decades later—because often the same innovative ideas and concepts reappear at a latter stage in a different form, on a more advanced technology platform.

Keeping this factor in mind, the required revisions have been made in chapters 1 to 11 of the first edition. Some revisions are necessitated by changes in technology. For example, the successful *Scalable Coherent Interface* (SCI) and InfiniBand standards grew out of the IEEE Futurebus+ standard, while the latter itself failed to take off. Thus it became necessary to describe this development in the revised edition of the book, since it has had an impact on the evolution of system interconnect technology.

The topic of *Instruction Level Parallelism* has been discussed in a self-contained manner in the newly introduced **Chapter 12**. The basic concepts and techniques of instruction-level parallelism have been described, followed by a discussion of the relevant system design and performance issues which often place a practical upper limit on its successful exploitation.

The trends and advances in underlying technologies have been discussed in another newly added **Chapter 13**, which also touches upon issues such as the design trade-offs involved in multi-core processors. The basic concepts of *parallel algorithms* and *stream processing* have been described. This chapter also includes several specific case-studies of recently introduced processors, systems, multi-core *systems-on-a-chip*, function libraries, and also the parallel programming language Chapel being developed at Cray.

Issues related to *instruction level parallelism*, *processor clock speed* and *power consumption* have defined the recent direction of development of processor design. These issues have been discussed as appropriate in both the newly introduced chapters.

Chapter Summary has been added to all the chapters. In the interests of usability and topical relevance, **Chapter 12** of the first edition and bibliographic notes provided at the end of each chapter in the first edition have been moved to the website associated with the book.

To summarize, the new features are as follows:

- Two new chapters on *Instruction Level Parallelism* and *Trends in Parallel Systems*
- Topical inclusions—
 - Pipelining hazards, data hazards, control hazards
 - PCI bus and PCI Express
 - Cluster computing, interconnection networks and clusters
 - MPI, openMP, PVM, Pthreads
 - Multi-core processors
 - Impact of technology
 - Stream processing
 - Programming language Chapel
 - Introduction to VLSI computing structures
- Updated coverage of recent processors and systems—Intel Pentium IV, UltraSparc, Blue Gene (from IBM), Cray XT series, XT5 and XMT
- Chapters 1 to 11 revised to better reflect the contributions of earlier systems to the development of computer architecture
- *Chapter Summary* added in each chapter
- Strong pedagogical features:
 - 360 Illustrations
 - 114 Solved Examples
 - 251 Exercise Problems
 - 6 (new) Case Studies

Chapter Organization

The book is divided into five parts and 13 chapters.

Part I is on the *Theory of Parallelism* and contains chapters 1 to 3.

Chapter 1 on Parallel Computer Models describes the state of computing, multiprocessors, multicomputers, multivector and SIMD computers, and PRAM and VLSI models. It also gives an overview of architectural development tracks.

Chapter 2 discusses Program and Network Properties. Conditions of parallelism, program partitioning and scheduling, program flow mechanisms and system interconnect architectures are explained in depth.

Chapter 3 deals with the Principles of Scalable Performance. This chapter covers the topics on performance metrics and measures, parallel processing applications, speedup performance laws and scalability analysis and approaches.

Part II deals with *Hardware Technologies* and contains chapters 4 to 6.

Chapter 4 provides coverage on Processors and Memory Hierarchy. Advanced processor technology, memory hierarchy technology and virtual memory technology are discussed in detail.

Chapter 5 presents the concepts of Bus, Cache and Shared Memory. It treats the topics of backplane bus systems, cache memory organizations, shared memory organizations and sequential and weak consistency models in depth.

Chapter 6 is on Pipelining and Superscalar Techniques. Linear and nonlinear pipeline processors; and instruction, arithmetic and superscalar design are explained in detail.

Part III is on *Parallel and Scalable Architectures* and contains chapters 7 to 9.

Multiprocessors and Multicomputers are taken up in **Chapter 7**. Multiprocessor system interconnects, cache coherence and synchronization mechanisms, three generations of multicomputers and message-passing mechanisms are covered in detail.

Chapter 8 is on Multivector and SIMD Computers. This deals with vector processing principles, multivector multiprocessors, compound vector processing and SIMD computer organizations. An overview of the Connection Machine CM-5 is also given here.

Chapter 9 is on Scalable, Multithreaded and Dataflow Architecture. This chapter takes up the concepts of latency-hiding techniques, principles of multithreading, fine-grain multicomputers, scalable, multithreaded, dataflow and hybrid architectures.

Part IV discusses *Software for Parallel Programming*. This part contains chapters 10 and 11.

Chapter 10 covers Parallel Models, Languages and Computers. Parallel programming models, languages and compilers are discussed first. Dependence analysis of data arrays and code optimization and scheduling are taken up thereafter. Finally, the chapter ends with detailed section on loop parallelization and pipelining.

Parallel Program Development and Environments are explained in **Chapter 11**. Here, parallel program environments, synchronization and multiprocessing modes, shared-variable program structures, message-passing program development and mapping programs onto multicomputers are dealt with in depth.

Part V deals with the *Instruction and System Level Parallelism*. This part contains the two newly introduced chapters 12 and 13.

Chapter 12 is on Instruction Level Parallelism. Design issues are brought out using the basic problem definition in the context of a typical processor model. Specific techniques discussed include reorder buffer, register renaming, operand forwarding, branch prediction, and Tomasulo's algorithm. Compiler-detected instruction level parallelism and thread level parallelism are also explained, along with the practical limitations encountered in exploiting instruction level parallelism.

Chapter 13 elucidates the current trends in parallel systems. A brief overview of relevant semiconductor, display, storage, interconnect and network technologies is taken up first, followed by forms of parallelism, some relevant concepts of parallel algorithms, and several case studies of current processors and systems. In the latter part of the chapter, parallel programming models and the parallel programming language Chapel are discussed.

All chapters are supplemented with exhaustive pedagogical features like summary, solved examples and exercise problems. Moreover, case studies are given at relevant places.

Issues related to *instruction level parallelism*, *processor clock speed* and *power consumption* have defined the recent direction of development of processor design. These issues have been discussed as appropriate in both the newly introduced chapters.

Chapter Summary has been added to all the chapters. In the interests of usability and topical relevance, **Chapter 12** of the first edition and bibliographic notes provided at the end of each chapter in the first edition have been moved to the website associated with the book.

To summarize, the new features are as follows:

- Two new chapters on *Instruction Level Parallelism* and *Trends in Parallel Systems*
- Topical inclusions—
 - Pipelining hazards, data hazards, control hazards
 - PCI bus and PCI Express
 - Cluster computing, interconnection networks and clusters
 - MPI, openMP, PVM, Pthreads
 - Multi-core processors
 - Impact of technology
 - Stream processing
 - Programming language Chapel
 - Introduction to VLSI computing structures
- Updated coverage of recent processors and systems—Intel Pentium IV, UltraSparc, Blue Gene (from IBM), Cray XT series, XT5 and XMT
- Chapters 1 to 11 revised to better reflect the contributions of earlier systems to the development of computer architecture
- *Chapter Summary* added in each chapter
- Strong pedagogical features:
 - 360 Illustrations
 - 114 Solved Examples
 - 251 Exercise Problems
 - 6 (new) Case Studies

Chapter Organization

The book is divided into five parts and 13 chapters.

Part I is on the *Theory of Parallelism* and contains chapters 1 to 3.

Chapter 1 on *Parallel Computer Models* describes the state of computing, multiprocessors, multicomputers, multivector and SIMD computers, and PRAM and VLSI models. It also gives an overview of architectural development tracks.

Chapter 2 discusses *Program and Network Properties*. Conditions of parallelism, program partitioning and scheduling, program flow mechanisms and system interconnect architectures are explained in depth.

Chapter 3 deals with the *Principles of Scalable Performance*. This chapter covers the topics on performance metrics and measures, parallel processing applications, speedup performance laws and scalability analysis and approaches.

V S Shanker
Birla Institute of Technology, Mesra
Ranchi, Jharkhand

A P Shanthi
College of Engineering, Anna University
Chennai, Tamil Nadu

K V Madhu Murthy
College of Engineering, S V University
Tirupati, Tamil Nadu

N Subrahmanyam
National Institute of Technology (NIT-W)
Warangal, Andhra Pradesh

Sincere thanks are due to the editorial and production team at Tata McGraw-Hill who have worked on the book—and specifically to Surbhi Suman, Surbhi Shukla, Anjali Razdan and Sohini Mukherjee.

Naresh Jotwani

Feedback

We hope that teachers and students will find the revised edition of this book useful, as they set out to explore the fascinating world of high performance computer architecture. Please send your comments, views and suggestions to tmh.csefeedback@gmail.com, mentioning the title and author name in the subject line. Please, also do report any piracy of the book spotted by you.

