Memory Wall

Photo-Electroactive Non-Volatile Memories for Data Storage and Neuromorphic Computing summarizes advances in the development of photo-electroactive memories and neuromorphic computing systems, suggests possible solutions to the challenges of device design, and evaluates the prospects for commercial applications. Sections covers developments in electro-photoactive memory, and photonic neuromorphic and in-memory computing, including discussions on design concepts, operation principles and basic storage mechanism of optoelectronic memory devices, potential materials from organic molecules, semiconductor quantum dots to two-dimensional materials with desirable electrical and optical properties, device challenges, and possible strategies. This comprehensive, accessible and up-to-date book will be of particular interest to graduate students and researchers in solid-state electronics. It is an invaluable systematic introduction to the memory characteristics, operation principles and storage mechanisms of the latest reported electro-photoactive memory devices. Reviews the most promising materials to enable emerging computing memory and data storage devices, including one- and two-dimensional materials, metal oxides, semiconductors, organic materials, and more Discusses fundamental mechanisms and design strategies for two- and three-terminal device structures Addresses device challenges and strategies to enable translation of optical and optoelectronic technologies
This book constitutes the thoroughly refereed post-conference proceedings of the workshops held at the 37th International Symposium on Computer Architecture, ISCA 2010, in Saint-Malo, France, in June 2010. The 28 revised full papers presented were carefully reviewed and selected from

the lectures given at 5 of these workshops. The papers address topics ranging from novel memory architectures to emerging application design and performance analysis and encompassed the following workshops: A4MMC, applications for multi- and many-cores; AMAS-BT, 3rd workshop on architectural and micro-architectural support for binary translation; EAMA, the 3rd Workshop for emerging applications and many-core architectures; WEED, 2nd Workshop on energy efficient design, as well as WIOSCA, the annual workshop on the interaction between operating systems and computer architecture.

This is a practical book for computer engineers who want to understand or implement hardware/software systems. It focuses on problems that require one to combine hardware design with software design – such problems can be solved with hardware/software codesign. When used properly, hardware/software co- sign works better than hardware design or software design alone: it can improve the overall performance of digital systems, and it can shorten their design time. Hardware/software codesign can help a designer to make trade-offs between the ?exibility and the performanceof a digital system. To achieve this, a designer needs to combine two radically different ways of design: the sequential way of dec- position in time, using software, with the parallel way of decomposition in space, using hardware. Intended Audience This book assumes that you have a basic understandingof hardware that you are - miliar with standard digital hardware componentssuch as registers, logic gates, and components such as multiplexers and arithmetic operators. The book also assumes that you know how to write a program in C. These topics are usually covered in an introductory course on computer engineering or in a combination of courses on digital design and software engineering.

While the concept of integration or an interdisciplinary curriculum has been around for decades, the purposeful practice of integration is a relatively new educational endeavor. Though classroom teachers often say they “integrate,” there generally seems to be a lack of understanding of what this thing called integration is (theory) and what it is supposed to look like in the classroom (practice). Arguably, no other discipline has felt the pressure to integrate more than social studies. Marginalized by federal initiatives such as No Child Left Behind and suffering from a general crisis of credibility, social studies has been pushed further and further to the proverbial back burner of educational importance. Yet regardless of perspective or position, social studies remains ripe for integration. The crux of this book is to provide educators insights and strategies into how to integrate social studies with other discipline areas. Calling upon national experts in their respective fields, each chapter chronicles the broad relationship between individual content areas and social studies. Multiple examples of integrative opportunities are included. At the end of each chapter is a series of grade-specific integrative lesson plans ready for implementation. This book was purposefully designed as a how-to, hands-on, ready-reference guide for educators at all stages and all levels of teaching.

This book proposes novel memory hierarchies and software optimization techniques for the optimal utilization of memory hierarchies. It presents a wide range of optimizations, progressively increasing in the complexity of analysis and of memory hierarchies. The final chapter covers optimization techniques for applications consisting of multiple processes found in most modern embedded devices.

This book provides a structured introduction of the key concepts and techniques that enable in-/near-memory computing. For decades, processing-in-memory or near-

memory computing has been attracting growing interest due to its potential to break the memory wall. Near-memory computing moves compute logic near the memory, and thereby reduces data movement. Recent work has also shown that certain memories can morph themselves into compute units by exploiting the physical properties of the memory cells, enabling in-situ computing in the memory array. While in- and near-memory computing can circumvent overheads related to data movement, it comes at the cost of restricted flexibility of data representation and computation, design challenges of compute capable memories, and difficulty in system and software integration. Therefore, wide deployment of in-/near-memory computing cannot be accomplished without techniques that enable efficient mapping of data-intensive applications to such devices, without sacrificing accuracy or increasing hardware costs excessively. This book describes various memory substrates amenable to in- and near-memory computing, architectural approaches for designing efficient and reliable computing devices, and opportunities for in-/near-memory acceleration of different classes of applications.

This book explores how prominent sites across the National Mall remember US history, both individually and in concert with other sites throughout the Mall. Collectively, these sites reveal how the nation remembers itself and convey key elements of its collective nature.

Parallel Computing Architectures and APIs: IoT Big Data Stream Processing commences from the point high-performance uniprocessors were becoming increasingly complex, expensive, and power-hungry. A basic trade-off exists between the use of one or a small number of such complex processors, at one extreme, and a moderate to very large number of simpler processors, at the other. When combined with a high-bandwidth, interprocessor

communication facility leads to significant simplification of the design process. However, two major roadblocks prevent the widespread adoption of such moderately to massively parallel architectures: the interprocessor communication bottleneck, and the difficulty and high cost of algorithm/software development. One of the most important reasons for studying parallel computing architectures is to learn how to extract the best performance from parallel systems. Specifically, you must understand its architectures so that you will be able to exploit those architectures during programming via the standardized APIs. This book would be useful for analysts, designers and developers of high-throughput computing systems essential for big data stream processing emanating from IoT-driven cyber-physical systems (CPS). This pragmatic book: Devolves uniprocessors in terms of a ladder of abstractions to ascertain (say) performance characteristics at a particular level of abstraction Explains limitations of uniprocessor high performance because of Moore's Law Introduces basics of processors, networks and distributed systems Explains characteristics of parallel systems, parallel computing models and parallel algorithms Explains the three primary categorical representatives of parallel computing architectures, namely, shared memory, message passing and stream processing Introduces the three primary categorical representatives of parallel programming APIs, namely, OpenMP, MPI and CUDA Provides an overview of Internet of Things (IoT), wireless sensor networks (WSN), sensor data processing, Big Data and stream processing Provides introduction to 5G communications, Edge and Fog computing Parallel Computing Architectures and APIs: IoT Big Data Stream Processing discusses stream processing that enables the gathering, processing and analysis of high-volume, heterogeneous, continuous Internet of Things (IoT) big data streams, to extract insights and actionable results in real time.

Application domains requiring data stream management include military, homeland security, sensor networks, financial applications, network management, web site performance tracking, real-time credit card fraud detection, etc.

Set on four continents, stories about memory.

This book constitutes the refereed proceedings of the 9th International Conference on Security, Privacy, and Applied Cryptography Engineering, SPACE 2019, held in Gandhinagar, India, in December 2019. The 12 full papers presented were carefully reviewed and selected from 24 submissions. This annual event is devoted to various aspects of security, privacy, applied cryptography, and cryptographic engineering. This is a very challenging field, requiring the expertise from diverse domains, ranging from mathematics to solid-state circuit design.

This book presents the results of the study of the wall paintings from the Northeast Bastion at Ayia Irini, situating them within the wider social context of Kea and the Aegean world. Like the spectacularly well-preserved Akrotiri on Thera, with which these paintings are contemporary, Ayia Irini thrived 3,500 years ago. But unlike Akrotiri, Ayia Irini was not protected by a layer of volcanic ash. When the site was excavated in the 1960s-1970s by the University of Cincinnati under the auspices of the American School of Classical Studies at Athens, the paintings had long since collapsed, fractured into thousands of small pieces. This study attempts to bring the wall paintings back to life. Within the Northeast Bastion was a miniature frieze and, in the adjacent room, large-scale panels of plants. Human action set within

townscapes, landscapes, and the sea presents a vivid account of the social life and environment of the people for whom this harbor town was vital within the trading network of the time. This book explores the social implications of the fascinating and often unique iconography of the paintings whose setting within a fortification wall is quite extraordinary. The volume is profusely illustrated with color drawings, visualizations, and photographs.

Message from the General Co-chairs It is our honor and pleasure as General Co-chairs to welcome you to the p- ceedings of HiPEAC 2010 which was held in Pisa. This was the ?fth HiPEAC conference, following in the strong tradition of the ?rst conference in Barcelona in 2005 and the subsequent conferences in Ghent (2007), Goteborg (2008), and Paphos (2009). HiPEAC2010o?eredarichan ddiversesetoftechnicalandnon-technical- tivities. The technical activities included most importantly another strong te- nical program, and in addition, eight workshops and ?ve tutorials, all central to the HiPEAC network roadmap. The workshops explored multi-cores, simu- tion and performance evaluation, compiler and optimizations, design reliability, recon?gurable computing, interconnection networks,operating system and c- puter architecture codesign. The tutorials dealt with statistical methodology to evaluate program speed-ups, design for reliability, how to teach introductory computer architecture and programming, programming FPGA-based accele- tors and adaptability. We were particularly fortunate to have two keynote addresses, one by Bob Iannucci, formerly from Nokia, on how data

center thinking can be e?ectively ushered into the embedded system domain, and one by Roger Espasa from Intel on the Larrabee Architecture. The non-technical activities re?ected the academic, historical, and cultural charm of Pisa, a major center of Tuscany, and we hope the participants took advantage of our scheduled guided tour of historical Pisa and the conference banquet in a historic villa.

The Committee on the Future of Supercomputing was tasked to assess prospects for supercomputing technology research and development in support of U.S. needs, to examine key elements of context--the history of supercomputing, the erosion of research investment, the changing nature of problems demanding supercomputing, and the needs of government agencies for supercomputing capabilities--and to assess options for progress. This interim report establishes context--including the history and current state of supercomputing, application requirements, technology evolution, the socioeconomic context--to identify some of the issues that may be explored in more depth in the second phase of the study.

This book defines and explores the problem of placing the instances of dynamic data types on the components of the heterogeneous memory organization of an embedded system, with the final goal of reducing energy consumption and improving performance. It is one of the first to cover the problem of placement for dynamic data objects on embedded systems with heterogeneous memory architectures, presenting a complete methodology that can be easily adapted to real cases

and work flows. The authors discuss how to improve system performance and energy consumption simultaneously. Discusses the problem of placement for dynamic data objects on embedded systems with heterogeneous memory architectures; Presents a complete methodology that can be adapted easily to real cases and work flows; Offers hints on how to improve system performance and energy consumption simultaneously.

This book constitutes the thoroughly refereed post-conference proceedings of the 29th International Workshop on Languages and Compilers for Parallel Computing, LCPC 2016, held in Rochester, NY, USA, in September 2016. The 20 revised full papers presented together with 4 short papers were carefully reviewed. The papers are organized in topical sections on large scale parallelism, resilience and persistence, compiler analysis and optimization, dynamic computation and languages, GPUs and private memory, and runt-time and performance analysis.

Verification of real-time requirements in systems-on-chip becomes more complex as more applications are integrated. Predictable and composable systems can manage the increasing complexity using formal verification and simulation. This book explains the concepts of predictability and composability and shows how to apply them to the design and analysis of a memory controller, which is a key component in any real-time system.

"This book includes more than 80 games to help you break down barriers, communicate better, and

generate new ideas, insights, and strategies. The authors have identified tools and techniques from some of the world's most innovative professionals, whose teams collaborate and make great things happen. Gamestorming is the result: a unique collection of games that encourage engagement and creativity while bringing more structure and clarity to the workplace"--Page 4 of cover.

Containing over 300 entries in an A-Z format, the Encyclopedia of Parallel Computing provides easy, intuitive access to relevant information for professionals and researchers seeking access to any aspect within the broad field of parallel computing. Topics for this comprehensive reference were selected, written, and peer-reviewed by an international pool of distinguished researchers in the field. The Encyclopedia is broad in scope, covering machine organization, programming languages, algorithms, and applications. Within each area, concepts, designs, and specific implementations are presented. The highly-structured essays in this work comprise synonyms, a definition and discussion of the topic, bibliographies, and links to related literature. Extensive cross-references to other entries within the Encyclopedia support efficient, user-friendly searchers for immediate access to useful information. Key concepts presented in the Encyclopedia of Parallel Computing include; laws and metrics; specific numerical and non-numerical

algorithms; asynchronous algorithms; libraries of subroutines; benchmark suites; applications; sequential consistency and cache coherency; machine classes such as clusters, shared-memory multiprocessors, special-purpose machines and dataflow machines; specific machines such as Cray supercomputers, IBM's cell processor and Intel's multicore machines; race detection and auto parallelization; parallel programming languages, synchronization primitives, collective operations, message passing libraries, checkpointing, and operating systems. Topics covered: Speedup, Efficiency, Isoefficiency, Redundancy, Amdahls law, Computer Architecture Concepts, Parallel Machine Designs, Benmarks, Parallel Programming concepts & design, Algorithms, Parallel applications. This authoritative reference will be published in two formats: print and online. The online edition features hyperlinks to cross-references and to additional significant research. Related Subjects: supercomputing, high-performance computing, distributed computing

Teaching fundamental design concepts and the challenges of emerging technology, this textbook prepares students for a career designing the computer systems of the future. In-depth coverage of complexity, power, reliability and performance, coupled with treatment of parallelism at all levels, including ILP and TLP, provides the state-of-the-art

training that students need. The whole gamut of parallel architecture design options is explained, from core microarchitecture to chip multiprocessors to large-scale multiprocessor systems. All the chapters are self-contained, yet concise enough that the material can be taught in a single semester, making it perfect for use in senior undergraduate and graduate computer architecture courses. The book is also teeming with practical examples to aid the learning process, showing concrete applications of definitions. With simple models and codes used throughout, all material is made open to a broad range of computer engineering/science students with only a basic knowledge of hardware and software.
Die grundlegenden Prinzipien von Mikrocontrollern und Mikroprozessoren: Detailliert erläutern die Autoren den neuesten Stand der Technik, alle wichtigen Entwicklungstendenzen und den aktuellen Forschungsstand. Ferner analysieren sie in der Praxis häufig verwendete Mikrocontroller und Mikroprozessoren in ihrer Funktionsweise und schildern zukunftsweisende Technologien. Dieses Buch eignet sich besonders für Studierende der Informatik oder Elektrotechnik im fortgeschrittenen Grundstudium oder zu Beginn des Hauptstudiums. Es ist ebenso empfehlenswert für Fachleute, die sich mit Planung, Entwicklung und Einsatz dieser Hardware-Bausteine befassen.
Die Leistungsfahigkeit moderner Rechensysteme

erhoht sich stetig. Allerdings zeigen Untersuchungen, dass die Rechenleistung von Prozessoren starker steigt, als die Ein-/Ausgabeleistung zu den Sekundarspeichern, auf denen die Daten zur Berechnung abgelegt werden. Dies fuhrt dazu, dass Prozessoren zukunftig nicht ihr gesamtes Rechenpotential ausschopfen konnen, da sie auf Daten der Sekundarspeicher warten mussen. Damit die Leistung der Speicher nicht zu einem leistungsbegrenzenden Faktor des gesamten Systems wird, ist die Leistungsanalyse und -optimierung der Sekundarspeicher notwendig. Die Leistungsmessung von Sekundarspeichersystemen wird typischerweise mit Softwarewerkzeugen durchgefuhrt, die eine Last auf dem Sekundarspeicher erzeugen und anhand dieser Last eine Leistungsermittlung durchfuhren. Diese sogenannten I/O-Benchmarks haben zahlreiche Probleme, die im Rahmen dieser Dissertation aufgezeigt und gelost werden. Es wird ein neuer Ansatz entwickelt, der realitatsnahes, nutzerrelevantes, vergleichbares und dennoch einfaches I/O-Benchmarking insbesondere in Hinblick auf die Leistungsermittlung beim Zugriff auf Sekundarspeicher mittels der MPI-IO-Schnittstelle ermoglicht. Ausgehend von den notwendigen Schritten bei der Leistungsanalyse wird eine neue Benchmark-Architektur entwickelt, die insbesondere Losungen fur die gefundenen Probleme der geringen

Reprasentativitat von Benchmarkergebnissen und der fehlenden Nutzerunterstutzung beim Benchmarking bietet und damit uber vorhandene Arbeiten in diesem Themenbereich deutlich hinausgeht. Es wird ein Benchmark-System erstellt, das nutzerrelevante Ergebnisse ermittelt, indem es dem Nutzer ermoglicht, das Lastverhalten MPI-IO-basierter Applikationen als Messgrundlage zu verwenden. Ausserdem wird eine realitatsnahe und einfach nutzbare I/O-Lastbeschreibung prasentiert, deren Moglichkeiten existierende I/O-Lastbeschreibungen in Hinblick auf Genauigkeit bei Verwendung komplexer paralleler I/O-Lasten ubersteigen. Die Funktionsfahigkeit und Genauigkeit des I/O-Benchmarking-Ansatzes wird mit Messungen anhand von Beispielapplikationen gezeigt.

Memory WallNovelleC.H.Beck

The book includes the insights that reflect ‘Advances in Computer and Computational Sciences’ from upcoming researchers and leading academicians across the globe. It contains the high-quality peer-reviewed papers of ‘International Conference on Computer, Communication and Computational Sciences (IC4S 2017), held during 11–12 October, 2017 in Thailand. These papers are arranged in the form of chapters. The content of this book is divided into two volumes that cover variety of topics such as intelligent hardware and software

design, advanced communications, intelligent computing techniques, intelligent image processing, and web and informatics. This book helps the perspective readers' from computer industry and academia to derive the advances of next generation computer and communication technology and shape them into real life applications.

Comprehensive and timely, Cloud Computing: Concepts and Technologies offers a thorough and detailed description of cloud computing concepts, architectures, and technologies, along with guidance on the best ways to understand and implement them. It covers the multi-core architectures, distributed and parallel computing models, virtualization, cloud developments, workload and Service-Level-Agreements (SLA) in cloud, workload management. Further, resource management issues in cloud with regard to resource provisioning, resource allocation, resource mapping and resource adaptation, ethical, non-ethical and security issues in cloud are followed by discussion of open challenges and future directions. This book gives students a comprehensive overview of the latest technologies and guidance on cloud computing, and is ideal for those studying the subject in specific modules or advanced courses. It is designed in twelve chapters followed by laboratory setups and experiments. Each chapter has multiple choice questions with answers, as well as review questions and critical thinking

questions. The chapters are practically-focused, meaning that the information will also be relevant and useful for professionals wanting an overview of the topic.

This conference marked the ?rst time that the Asia-Paci?c Computer Systems Architecture Conference was held outside Australasia (i. e. Australia and New Zealand), and was, we hope, the start of what will be a regular event. The conference started in 1992 as a workshop for computer architects in Australia and subsequently developed into a full-?edged conference covering Austra- sia. Two additional major changes led to the present conference. The ?rst was a change from “computer architecture” to “computer systems architecture”, a change that recognized the importance and close relationship to computer arc- tecture of certain levels of software (e. g. operating systems and compilers) and of other areas (e. g. computer networks). The second change, which re?ected the increasing number of papers being submitted from Asia, was the replacement of “Australasia” with “Asia-Paci?c”. This year’s event was therefore particularly signi?cant, in that it marked the beginning of a truly “Asia-Paci?c” conference. It is intended that in the future the conference venue will alternate between Asia and Australia/New Zealand and, although still small, we hope that in time the conference will develop into a major one that represents Asia to the same - tent as existing major computer-architecture conferences in North America and Europe represent those regions.

In this book, a global team of experts from academia,

research institutes and industry presents their vision on how new nano-chip architectures will enable the performance and energy efficiency needed for AI-driven advancements in autonomous mobility, healthcare, and man-machine cooperation. Recent reviews of the status quo, as presented in CHIPS 2020 (Springer), have prompted the need for an urgent reassessment of opportunities in nanoelectronic information technology. As such, this book explores the foundations of a new era in nanoelectronics that will drive progress in intelligent chip systems for energy-efficient information technology, on-chip deep learning for data analytics, and quantum computing. Given its scope, this book provides a timely compendium that hopes to inspire and shape the future of nanoelectronics in the decades to come.

Exa-scale computing needs to re-examine the existing hardware platform that can support intensive data-oriented computing. Since the main bottleneck is from memory, we aim to develop an energy-efficient in-memory computing platform in this book. First, the models of spin-transfer torque magnetic tunnel junction and racetrack memory are presented. Next, we show that the spintronics could be a candidate for future data-oriented computing for storage, logic, and interconnect. As a result, by utilizing spintronics, in-memory-based computing has been applied for data encryption and machine learning. The implementations of in-memory AES, Simon cipher, as well as interconnect are explained in details. In addition, in-memory-based machine learning and face recognition are also illustrated in this book.

This book constitutes the refereed proceedings of the 10th International Symposium on Parallel Architectures, Algorithms and Programming, PAAP 2019, held in Guangzhou, China, in December 2019. The 39 revised full papers and 8 revised short papers presented were carefully reviewed and selected from 121 submissions. The papers deal with research results and development activities in all aspects of parallel architectures, algorithms and programming techniques.

In this important book, one of Latin America's foremost critical theorists examines the use and abuse of memory in the wake of the social and political trauma of Pinochet's Chile. Focusing on the period 1990–2015, Nelly Richard denounces the politics and aesthetics of forgetting that have underpinned both the protracted transition out of dictatorship and the denial of justice to its survivors and victims. What are the perils and social costs of a culture of forgetting? What forms do memories of injustice take in newly formed democracies? How might a history of violence and an ethics of reparation be reconciled in post-autocratic societies? In addressing these and other questions, Richard exposes the abuses of the past and the present while also attending to the residues of memory that are manifested in street protests, literature, and the media, and in artistic practices from architecture and urban design to installation and film. While cultural artifacts can be powerful devices for resistance and critique, Richard argues that they can also be complicit in reproducing and collaborating with forms of institutional and political oblivion. Both within Chile and beyond, Richard offers a

trenchant critique of how authoritarian regimes and neoliberal states whittle away at memory's critical capacity. At a time of seismic political realignments in Latin America and internationally, Eruptions of Memory makes a powerful case for the ethical, political, and aesthetic value of memory.

Unser Leben, unsere Welt werden durch unsere Erinnerungen zusammengehalten. Was geschieht mit uns, wenn wir sie verlieren, und welche Möglichkeiten tun sich auf, wenn andere unsere Erinnerungen wiederbeleben können? Der 74-jährigen Alma Konachek, die in einem Vorort von Kapstadt lebt, widerfährt genau dies. Sie verliert ihr Gedächtnis. Unbekannte brechen mehrfach in ihr Haus ein, auf der Suche nach Hinweisen zu einem spektakulären Fossilienfund ihres plötzlich verstorbenen Mannes. Denn Alma hat eine Wand voller Fotos, Gedächtnisstützen, Speichermedien, in der sich irgendwo der fehlende Hinweis zu dem gesuchten Fossil befindet. In dieser lichten, wunderschönen Novelle gelangt schließlich ein Junge in den Besitz des Geheimnisses dieser alten Frau und ihres Mannes, einer Episode aus ihrer Vergangenheit mit der Macht, ein Leben zum Guten zu wenden. Der Junge reist dazu in die Karoo-Wüste und setzt sich dieser wilden Landschaft aus. Wie alle Werke Doerrs zeugt auch dieses von der Größe des Lebens – von der geheimnisvollen Schönheit der Fossilien, Wolken, Blätter – vom atemberaubenden Glück, in diesem Universum zu leben. Die Vorstellungskraft und Sprachmacht, das Einfühlungsvermögen und die Erzählkunst Anthony Doerrs sind unvergleichlich.

Technological Advances and Problems of High Performance Communications An ecosystem of solutions along a stack of technology layers Cohesively collecting state-of-the-art contributions from leading researchers in industry, national laboratories, and academia, Attaining High Performance Communications: A Vertical Approach discusses various issues pertaining to high performance communications in a particular layer of a vertical stack. It explores efficient interconnection hardware, the architectural aspects of network adapters and their integration with processor cores, the design of scalable and robust high performance end-to-end communications services and protocols, and system services and tools for new multi-core environments. No single solution applied at one particular layer can help applications solve all performance-related issues with communication services. Instead, this book shows that a coordinated effort is needed among the layers. It covers many different types of technologies and layers across the stack, from the architectural features of the hardware, through the protocols and their implementation in operating system kernels, to the manner in which application services and middleware are using underlying platforms. The book also describes key developments in high-end platforms, high performance interconnection fabrics and communication libraries, and multi- and many-core systems. This volume addresses the challenges involved in emerging types of communications applications, platforms, and services. Examining each layer in the vertical stack, it illustrates how to eliminate bottlenecks and provide optimization

opportunities.
Computing systems are undergoing a transformation from logic-centric towards memory-centric architectures, where overall performance and energy efficiency at the system level are determined by the density, performance, functionality and efficiency of the memory, rather than the logic sub-system. This is driven by the requirements of data-intensive applications in artificial intelligence, autonomous systems, and edge computing. We are at an exciting time in the semiconductor industry where several innovative device and technology concepts are being developed to respond to these demands, and capture shares of the fast growing market for AI-related hardware. This special issue is devoted to highlighting, discussing and presenting the latest advancements in this area, drawing on the best work on emerging memory devices including magnetic, resistive, phase change, and other types of memory. The special issue is interested in work that presents concepts, ideas, and recent progress ranging from materials, to memory devices, physics of switching mechanisms, circuits, and system applications, as well as progress in modeling and design tools. Contributions that bridge across several of these layers are especially encouraged.
Programming multi-core and many-core computing systems Sabri Pllana, Linnaeus University, Sweden Fatos Xhafa, Technical University of Catalonia, Spain Provides state-of-the-art methods for programming multi-core and many-core systems The book comprises a selection of twenty two chapters covering: fundamental techniques and algorithms; programming approaches; methodologies and frameworks; scheduling and management; testing and evaluation

methodologies; and case studies for programming multi-core and many-core systems. Program development for multi-core processors, especially for heterogeneous multi-core processors, is significantly more complex than for single-core processors. However, programmers have been traditionally trained for the development of sequential programs, and only a small percentage of them have experience with parallel programming. In the past, only a relatively small group of programmers interested in High Performance Computing (HPC) was concerned with the parallel programming issues, but the situation has changed dramatically with the appearance of multi-core processors on commonly used computing systems. It is expected that with the pervasiveness of multi-core processors, parallel programming will become mainstream. The pervasiveness of multi-core processors affects a large spectrum of systems, from embedded and general-purpose, to high-end computing systems. This book assists programmers in mastering the efficient programming of multi-core systems, which is of paramount importance for the software-intensive industry towards a more effective product-development cycle. Key features: Lessons, challenges, and roadmaps ahead. Contains real world examples and case studies. Helps programmers in mastering the efficient programming of multi-core and many-core systems. The book serves as a reference for a larger audience of practitioners, young researchers and graduate level students. A basic level of programming knowledge is required to use this book.

Information ist der gefragteste Rohstoff des neuen Jahrtausends. Doch Daten alleine reichen nicht aus, Business Intelligence ist gefragt: Hinter diesem vielversprechenden Begriff stehen Methoden, Prozesse, Werkzeuge und Strategien zur systematischen Analyse von Daten. Unternehmen wollen mit BI Erkenntnisse gewinnen,

mit denen sie bessere strategische und operative Managemententscheidungen treffen können. Dieses Ziel zu erreichen ist angesichts explodierender Datenmengen herausfordernd, aber zugleich lohnender denn je. Um dieses Ziel zu erreichen, müssen zunächst folgende Fragen beantwortet werden: Welchen Wertbeitrag bietet Business Intelligence konkret für Unternehmen? Wie kann Business Intelligence strategisch im Unternehmen positioniert werden (z. B. mit Hilfe eines Business Intelligence Competence Centers oder eines Chief Digital Officers)? Wie können BI-Projekte erfolgreich und agil umgesetzt werden? Wie lassen sich BI-Landschaften zukunftsfähig gestalten ? Wie können neue Herausforderungen (wie z. B. Schatten-BI und Datenschutz) bewältigt werden? Antworten auf all diese Fragen finden Sie im Handbuch Business Intelligence. Mit diesem Werk erhalten Sie das notwendige Wissen, um Business Intelligence als strategischen Wettbewerbsfaktor für Ihr Unternehmen zu nutzen. Bei den Autoren dieses Buches handelt es sich um 17 ausgewiesene Experten aus dem Business Intelligence-Bereich.

An engrossing middle-grade novel set in a high-fantasy video game world that’s part Kathryn Erskine’s Mockingbird, part Patrick Ness’s A Monster Calls. Wellhall is an immersive online fantasy world full of giants, sorcerers, and elves—and it’s junior-high-schooler Nick’s only escape from real life. Nick and his mom used to play the online video game together before her early-onset Alzheimer’s forced her to enter an assisted-living facility. At first, Nick seeks distraction in the game, but he soon becomes convinced that his mom is playing the game as a character named Reunne, and dropping him hints about her diagnosis and how he can help her return home. Even as Nick becomes more and more certain that Reunne is actually his mother, Nick’s father and his new friend encourage Nick to confront the possibility that

the game is just a game, and that he needs to be prepared to say goodbye to his mother as he knows her. . . . “Readers—gamers and nongamers alike—will cheer the resolution of Nick’s transformative journey. Thoughtful, earnest, and gratifying.” —Kirkus Reviews “A lovely, heartwarming story of a young man negotiating personal crises with the help of games, friends, and family, perfect for readers who appreciate a blend of fantasy and realism.” —The Bulletin “A complex, emotional story about grief and acceptance. . . . A strong, thought-provoking novel.” —Publishers Weekly

The energy consumption issue in distributed computing systemsraises various monetary, environmental and system performanceconcerns. Electricity consumption in the US doubled from 2000to 2005. From a financial and environmental standpoint,reducing the consumption of electricity is important, yet thesereforms must not lead to performance degradation of the computingsystems. These contradicting constraints create a suite ofcomplex problems that need to be resolved in order to lead to'greener' distributed computing systems. This book bringstogether a group of outstanding researchers that investigate thedifferent facets of green and energy efficient distributedcomputing. Key features: One of the first books of its kind Features latest research findings on emerging topics bywell-known scientists Valuable research for grad students, postdocs, andresearchers Research will greatly feed into other technologies andapplication domains