Linux 4 14 Kernel Benchmarks With The Intel Core I9 7980xe

This book shows the conference proceedings of CloudComp 2009 held in Munich, Germany, in October 2009.
High Performance Parallelism Pearls Volume 2 offers another set of examples that demonstrate how to leverage parallelism. Similar to Volume 1, the techniques included here explain how to use processors and coprocessors with the same programming – illustrating the most effective ways to combine Xeon Phi coprocessors with Xeon and other multicore processors. The book includes examples of successful programming efforts, drawn from across industries and domains such as biomed, genetics, finance, manufacturing, imaging, and more. Each chapter in this edited work includes detailed explanations of the programming techniques used, while showing high performance results on both Intel Xeon Phi coprocessors and multicore processors. Learn from dozens of new examples and case studies illustrating "success stories" demonstrating not just the features of Xeon-powered systems, but also how to leverage parallelism across these heterogeneous systems. Promotes write-once, run-anywhere coding, showing how to code for high performance on multicore processors and Xeon Phi Examples from multiple vertical domains illustrating real-world use of Xeon Phi coprocessors Source code available for download to facilitate further exploration
This book constitutes the refereed proceedings of the 9th International Conference on High Performance Computing, HiPC 2002, held in Bangalore, India in December 2002. The 57 revised full contributed papers and 9 invited papers presented together with various keynote abstracts were carefully reviewed and selected from 145 submissions. The papers are organized in topical sections on algorithms, architecture, systems software, networks, mobile computing and databases, applications, scientific computation, embedded systems, and biocomputing.
This volume contains the lectures given in honor to Georg Färber as tribute to his contributions in the area of real-time and embedded systems. The chapters of many leading scientists cover a wide range of aspects, like robot or automotive vision systems or medical aspects.
This book constitutes the refereed proceedings of the 22nd International Symposium on Model Checking Software, SPIN 2015, held in Stellenbosch, South Africa, in August 2015. The 18 papers presented – 14 regular papers and 4 tool or new idea papers – were carefully reviewed and selected from 27 submissions. They cover the field between theoretical advances and practical considerations and are organized in topical sections such as abstraction, refinement, translation; Büchi automata and hashing; embedded systems; heuristics and benchmarks; SAT/SMT- based approaches; software validation and verification.
This book constitutes the proceedings of the 10th IFIP International Conference on Network and Parallel Computing, NPC 2013, held in Guiyang, China, in September 2013. The 34 papers presented in this volume were carefully reviewed and selected from 109 submissions. They are organized in topical sections named: parallel programming and algorithms; cloud resource management; parallel architectures; multi-core computing and GPU; and miscellaneous.
This book contains the revised selected papers of 4 workshops held in conjunction with the International Conference on High Performance Computing, Networking, Storage and Analysis (SC) in November 2017 in Denver, CO, USA, and in November 2018 in Dallas, TX, USA: the 6th and 7th International Workshop on Extreme-Scale Programming Tools, ESPT 2017 and ESPT 2018, and the 4th and 5th International Workshop on Visual Performance Analysis, VPA 2017 and VPA 2018. The 11 full papers of ESPT 2017 and ESPT 2018 and the 6 full papers of VPA 2017 and VPA 2018 were carefully reviewed and selected for inclusion in this book. The papers discuss the requirements for exascale-enabled tools as well as new approaches of applying visualization and visual analytic techniques to large-scale applications. Topics of interest include: programming tools; methodologies for performance engineering; tool technologies for extreme-scale challenges (e.g., scalability, resilience, power); tool support for accelerated architectures and large-scale multi-cores; tool infrastructures and environments; evolving/future application requirements for programming tools and technologies; application developer experiences with programming and performance tools; scalable displays of performance data; case studies demonstrating the use of performance visualization in practice; data models to enable scalable visualization; graph representation of unstructured performance data; presentation of high-dimensional data; visual correlations between multiple data sources; human-computer interfaces for exploring performance data; and multi-scale representations of performance data for visual exploration.
This book contains extended versions of key papers from the 2nd International Conference on High-Performance Embedded Architectures and Compilers (HiPEAC 2007). It also covers such topics as microarchitecture, code generation, and performance modeling.
High Performance Computing (HPC) remains a driver that offers huge potentials and benefits for science and society. However, a profound understanding of the computational matters and specialized software is needed to arrive at effective and efficient simulations. Dedicated software tools are important parts of the HPC software landscape, and support application developers. Even though a tool is by definition not a part of an application, but rather a supplemental piece of software, it can make a fundamental difference during the development of an application. Such tools aid application developers in the context of debugging, performance analysis, and code optimization, and therefore make a major contribution to the development of robust and efficient parallel software. This book introduces a selection of the tools presented and discussed at the 9th International Parallel Tools Workshop held in Dresden, Germany, September 2-3, 2015, which offered an established forum for discussing the latest advances in parallel tools.
The aim of this book is to explain to high-performance computing (HPC) developers how to utilize the Intel® Xeon PhiTM series products efficiently. To that end, it introduces some computing grammar, programming technology and optimization methods for using many-integrated-core (MIC) platforms and also offers tips and tricks for actual use, based on the authors' first-hand optimization experience. The material is organized in three sections. The first section, "Basics of MIC", introduces the fundamentals of MIC architecture and programming, including the specific Intel MIC programming environment. Next, the section on "Performance Optimization" explains general MIC optimization techniques, which are then illustrated step-by-step using the classical parallel programming example of matrix multiplication. Finally, "Project development" presents a set of practical and experience-driven methods for using parallel computing in application projects, including how to determine if a serial or parallel CPU program is suitable for MIC and how to transplant a program onto MIC. This book appeals to two main audiences: First, software developers for HPC applications – it will enable them to fully exploit the MIC architecture and thus achieve the extreme performance usually required in biological genetics, medical imaging, aerospace, meteorology and other areas of HPC. Second,

students and researchers engaged in parallel and high-performance computing – it will guide them on how to push the limits of system performance for HPC applications.
This book constitutes the symposia and workshops of the 10th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP. Each of the sympois and workshops focuses on a particular theme and complements the spectrum of the main conference.
This book constitutes the thoroughly refereed post-proceedings of the 12th International Workshop on Job Scheduling Strategies for Parallel Processing, JSSPP 2006, held in Saint-Malo, France in June 2006 in conjunction with the Joint International Conference on Measurement and Modeling of Computer Systems SIGMETRICS/Performance 2006. The 12 revised full research papers cover all current issues of job scheduling strategies for parallel processing.
Increasing size and complexity of software and hardware systems makes it harder to ensure their reliability. At the same time, the issues of autonomous computing become more critical as we more and more rely on software systems in our daily life. Such complexity is getting even more critical with the ubiquitous computing of embedded devices and other pervasive systems. These trends ask for techniques and tools for developing reliable and autonomous software which can support software engineers in their efforts. This book summarizes the state of the art of research in the diverse fields concerned, including novel designs, case studies and experimental as well as theoretical results.
This IBM® Redbooks® publication is a refresh of IBM Technical Computing Clouds, SG24-8144, Enhance Inbound and Outbound Marketing with a Trusted Single View of the Customer, SG24-8173, and IBM Platform Computing Integration Solutions, SG24-8081, with a focus on High Performance and Technical Computing on IBM Power SystemsTM. This book describes synergies across the IBM product portfolio by using case scenarios and showing solutions such as IBM SpectrumTM Scale (formerly GPFSTM). This book also reflects and documents the IBM Platform Computing Cloud Services as part of IBM Platform Symphony® for analytics workloads and IBM Platform LSF® (with new features, such as a Hadoop connector, a MapReduce accelerator, and dynamic cluster) for job scheduling. Both products are used to help customers schedule and analyze large amounts of data for business productivity and competitive advantages. This book is targeted at technical professionals (consultants, technical support staff, IT Architects, and IT Specialists) that are responsible for delivering cost-effective cloud services and big data solutions on IBM Power Systems to uncover insights among client data so that they can take actions to optimize business results, product development, and scientific discoveries.
This book constitutes the refereed proceedings of the 20th International Conference on Parallel and Distributed Computing, Euro-Par 2014, held in Porto, Portugal, in August 2014. The 68 revised full papers presented were carefully reviewed and selected from 267 submissions. The papers are organized in 15 topical sections: support tools environments; performance prediction and evaluation; scheduling and load balancing; high-performance architectures and compilers; parallel and distributed data management; grid, cluster and cloud computing; green high performance computing; distributed systems and algorithms; parallel and distributed programming; parallel numerical algorithms; multicore and manycore programming; theory and algorithms for parallel computation; high performance networks and communication; high performance and scientific applications; and GPU and accelerator computing.
This book constitutes the refereed proceedings of the 11th International Conference on Parallel Computing, Euro-Par 2005, held in Lisbon, Portugal, in August/September 2005. The 120 revised papers presented together with 4 invited papers were carefully reviewed and selected from 388 submissions. The papers are organized in topical sections on support tools and environments, performance prediction and evaluation, scheduling and load balancing, compilers for high performance, parallel and distributed databases, data mining and knowledge discovery, grid and cluster computing: models, middleware and architectures, parallel computer architecture and instruction distributed systems and algorithms, parallel programming: models, methods, and languages, parallel numerical algorithms, distributed and high-performance multimedia, theory and algorithms for parallel computation, routing and communication in interconnection networks, mobile and ubiquitous computing, peer-to-peer and web computing, and applications of high-performance and grid computing.
Parallel Virtual Machine (PVM) and Message Passing Interface (MPI) are the most frequently used tools for programming according to the message passing paradigm, which is considered one of the best ways to develop parallel applications. This volume comprises 42 revised contributions presented at the Seventh European PVM/MPI Users' Group Meeting, which was held in Balatonfr ed, Hungary, 10 13 September 2000. The conference was organized by the Laboratory of Parallel and Distributed Systems of the Computer and Automation Research Institute of the Hungarian Academy of Sciences. This conference was previously held in Barcelona, Spain (1999), Liverpool, UK (1998) and Cracow, Poland (1997). The first three conferences were devoted to PVM and were held at the Technische Universit t M nchen, Germany (1996), Ecole Normale Superieure Lyon, France (1995), and University of Rome, Italy (1994). This conference has become a forum for users and developers of PVM, MPI, and other message passing environments. Interaction between those groups has proved to be very useful for developing new ideas in parallel computing and for applying existing ideas to new practical fields. The main topics of the meeting were evaluation and performance of PVM and MPI, extensions and improvements to PVM and MPI, algorithms using the message passing paradigm, and applications in science and engineering based on message passing. The conference included four tutorials and five invited talks on advances in MPI, cluster computing, network computing, grid computing, and SGI parallel computers and programming systems.
BPF and related observability tools give software professionals unprecedented visibility into software, helping them analyze operating system and application performance, troubleshoot code, and strengthen security. BPF Performance Tools: Linux System and Application Observability is the industry's most comprehensive guide to using these tools for observability. Brendan Gregg, author of the industry's definitive guide to system performance, introduces powerful new methods and tools for doing analysis that leads to more robust, reliable, and safer code. This authoritative guide: Explores a wide spectrum of software and hardware targets Thoroughly covers open source BPF tools from the Linux Foundation iovisor project's bcc and bpftrace repositories Summarizes performance engineering and kernel internals you need to understand Provides and discusses 150+ bpftrace tools, including 80 written specifically for this book: tools you can run as-is, without programming — or customize and develop further, using

diverse interfaces and the bpftrace front-end You'll learn how to use BPF (eBPF) tracing tools to analyze CPUs, memory, disks, file systems, networking, languages, applications, containers, hypervisors, security, and the Linux kernel. You'll move from basic to advanced tools and techniques, producing new metrics, stack traces, custom latency histograms, and more. It's like having a superpower: with Gregg's guidance and tools, you can analyze virtually everything that impacts system performance, so you can improve virtually any Linux operating system or application.

This book constitutes the refereed proceedings of the 8th Asian Computing Science Conference, ASIAN 2003, held in Mumbai, India in December 2003.The 16 revised full papers presented together with 2 invited papers were carefully reviewed and selected from 53 submissions. Among the topics addressed are type theory for operating systems protection, self configurable servers, network services, predicate detection, hierarchical specification, proof theory, electronic auctions, secure mobile computing, programming casculi, access control, middleware, program logic, real-time systems, and probabilistic distributed systems.

This book constitutes the refereed post-conference proceedings of 13 workshops held at the 33rd International ISC High Performance 2018 Conference, in Frankfurt, Germany, in June 2018: HPC I/O in the Data Center, HPC-IODC 2018; Workshop on Performance and Scalability of Storage Systems, WOPSSS 2018; 13th Workshop on Virtualization in High-Performance Cloud Computing, VHPC 2018; Third International Workshop on In Situ Visualization, WOIV 2018; 4th International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale, ExaComm 2018; International Workshop on OpenPOWER for HPC, IWOPH 2018; IXPUG Workshop: Many-Core Computing on Intel Processors; Workshop on Sustainable Ultrascale Computing Systems; Approximate and Transprecision Computing on Emerging Technologies, ATCET 2018; First Workshop on the Convergence of Large-Scale Simulation and Artificial Intelligence; Third Workshop for Open Source Supercomputing, OpenSuCo 2018; First Workshop on Interactive High-Performance Computing; Workshop on Performance Portable Programming Models for Accelerators, P^3MA 2018. The 53 full papers included in this volume were carefully reviewed and selected from 80 submissions. They cover all aspects of research, development, and application of large-scale, high performance experimental and commercial systems. Topics include HPC computer architecture and hardware; programming models, system software, and applications; solutions for heterogeneity, reliability, power efficiency of systems; virtualization and containerized environments; big data and cloud computing; and artificial intelligence.

This book constitutes the refereed proceedings of the 30th International Conference, ISC High Performance 2015, [formerly known as the International Supercomputing Conference] held in Frankfurt, Germany, in July 2015. The 27 revised full papers presented together with 10 short papers were carefully reviewed and selected from 67 submissions. The papers cover the following topics: cost-efficient data centers, scalable applications, advances in algorithms, scientific libraries, programming models, architectures, performance models and analysis, automatic performance optimization, parallel I/O and energy efficiency.

This book focuses on the core question of the necessary architectural support provided by hardware to efficiently run virtual machines, and of the corresponding design of the hypervisors that run them. Virtualization is still possible when the instruction set architecture lacks such support, but the hypervisor remains more complex and must rely on additional techniques. Despite the focus on architectural support in current architectures, some historical perspective is necessary to appropriately frame the problem. The first half of the book provides the historical perspective of the theoretical framework developed four decades ago by Popek and Goldberg. It also describes earlier systems that enabled virtualization despite the lack of architectural support in hardware. As is often the case, theory defines a necessary—but not sufficient—set of features, and modern architectures are the result of the combination of the theoretical framework with insights derived from practical systems. The second half of the book describes state-of-the-art support for virtualization in both x86-64 and ARM processors. This book includes an in-depth description of the CPU, memory, and I/O virtualization of these two processor architectures, as well as case studies on the Linux/KVM, VMware, and Xen hypervisors. It concludes with a performance comparison of virtualization on current-generation x86- and ARM-based systems across multiple hypervisors.

This volume represents the 18th International Conference on Information Technology - New Generations (ITNG), 2021. ITNG is an annual event focusing on state of the art technologies pertaining to digital information and communications. The applications of advanced information technology to such domains as astronomy, biology, education, geosciences, security, and health care are the among topics of relevance to ITNG. Visionary ideas, theoretical and experimental results, as well as prototypes, designs, and tools that help the information readily flow to the user are of special interest. Machine Learning, Robotics, High Performance Computing, and Innovative Methods of Computing are examples of related topics. The conference features keynote speakers, a best student award, poster award, service award, a technical open panel, and workshops/exhibits from industry, government and academia. This publication is unique as it captures modern trends in IT with a balance of theoretical and experimental work. Most other work focus either on theoretical or experimental, but not both. Accordingly, we do not know of any competitive literature.

This book constitutes the refereed proceedings of the International Conferences on Security Technology, SecTech 2012, on Control and Automation, CA 2012, and CES-CUBE 2012, the International Conference on Circuits, Control, Communication, Electricity, Electronics, Energy, System, Signal and Simulation; all held in conjunction with GST 2012 on Jeju Island, Korea, in November/December 2012. The papers presented were carefully reviewed and selected from numerous submissions and focus on the various aspects of security technology, and control and automation, and circuits, control, communication, electricity, electronics, energy, system, signal and simulation.

The two-volume set LNCS 12043 and 12044 constitutes revised selected papers from the 13th International Conference on Parallel Processing and Applied Mathematics, PPAM 2019, held in Bialystok, Poland, in September 2019. The 91 regular papers presented in these volumes were selected from 161 submissions. For regular tracks of the conference, 41 papers were selected from 89 submissions. The papers were organized in topical sections named as follows: Part I: numerical algorithms and parallel scientific computing; emerging HPC architectures; performance analysis and scheduling in HPC systems; environments and frameworks for parallel/distributed/cloud computing; applications of parallel computing; parallel non-numerical algorithms; soft computing with applications; special session on GPU computing; special session on parallel matrix factorizations. Part II: workshop on language-based parallel programming models (WLPP 2019); workshop on models algorithms and methodologies for hybrid parallelism in new HPC systems; workshop on power and energy aspects of computations (PEAC 2019); special session on tools for energy efficient computing; workshop on scheduling for parallel computing (SPC 2019); workshop on applied high performance numerical algorithms for PDEs; minisymposium on HPC applications in physical sciences; minisymposium on high performance computing interval methods; workshop on complex collective systems. Chapters "Parallel adaptive cross approximation for the multi-trace formulation of scattering problems" and "A High-Order Discontinuous Galerkin Solver with Dynamic Adaptive Mesh Refinement to Simulate Cloud Formation Processes" of LNCS 12043 are available open access under a Creative Commons Attribution 4.0 International License via link.springer.com.

This book contains thoroughly refereed and revised papers from the 8th International Andrei Ershov Memorial Conference on Perspectives of System Informatics, PSI 2011, held in

Akademgorodok, Novosibirsk, Russia, in June/July 2011. The 18 revised full papers and 10 revised short papers presented were carefully reviewed and selected from 60 submissions. The volume also contains 5 invited papers covering a range of hot topics in computer science and informatics. The papers are organized in topical sections on foundations of program and system development and analysis, partial evaluation, mixed computation, abstract interpretation, compiler construction, computer models and algorithms for bioinformatics, programming methodology and software engineering, information technologies, knowledge-based systems, and knowledge engineering.

This book constitutes the proceedings of the 27th International Conference on Parallel and Distributed Computing, Euro-Par 2021, held in Lisbon, Portugal, in August 2021. The conference was held virtually due to the COVID-19 pandemic. The 38 full papers presented in this volume were carefully reviewed and selected from 136 submissions. They deal with parallel and distributed computing in general, focusing on compilers, tools and environments; performance and power modeling, prediction and evaluation; scheduling and load balancing; data management, analytics and machine learning; cluster, cloud and edge computing; theory and algorithms for parallel and distributed processing; parallel and distributed programming, interfaces, and languages; parallel numerical methods and applications; and high performance architecture and accelerators.

This book constitutes the refereed post-proceedings of the Second International Conference on High Performance Networking, Computing, and Communication systems, ICHCC 2011, held in Singapore in May 2011. The conference was held together with the Second International Conference on Theoretical and Mathematical Foundations of Computer Science, ICTMF 2011, which proceedings are published in CCIS 164. The 84 revised selected papers presented were carefully reviewed and selected for inclusion in the book. The topics covered range from computational science, engineering and technology to digital signal processing, and computational biology to game theory, and other related topices.

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.

This book constitutes revised selected papers from 7 workshops that were held in conjunction with the ISC High Performance 2016 conference in Frankfurt, Germany, in June 2016. The 45 papers presented in this volume were carefully reviewed and selected for inclusion in this book. They stem from the following workshops: Workshop on Exascale Multi/Many Core Computing Systems, E-MuCoCoS; Second International Workshop on Communication Architectures at Extreme Scale, ExaComm; HPC I/O in the Data Center Workshop, HPC-IODC; International Workshop on OpenPOWER for HPC, IWOPH; Workshop on the Application Performance on Intel Xeon Phi – Being Prepared for KNL and Beyond, IXPUG; Workshop on Performance and Scalability of Storage Systems, WOPSSS; and International Workshop on Performance Portable Programming Models for Accelerators, P3MA.

This book constitutes the refereed proceedings of the First International Conference on High-Performance Computing and Communications, HPCC 2005, held in Sorrento, Italy in September 2005. The 76 revised full papers and 44 revised short papers presented were carefully reviewed and selected from 273 submissions. The papers are organized in topical sections on network protocols, routing, and algorithms; languages and compilers for HPC; parallel and distributed system architectures; embedded systems; parallel and distributed algorithms, wireless and mobile computing, Web services and Internet computing; peer-to-peer computing, grid and cluster computing, reliability, fault-tolerance, and security; performance evaluation and measurement; tools and environments for software development; distributed systems and applications; high performance scientific and engineering computing; database applications and data mining; HPSRF; pervasive computing and communications; and LMS.

Euro-Par is an annual series of international conferences dedicated to the p- motion and advancementof allaspects of parallelcomputing. The major themes can be divided into four broad categories: theory, high-performance, cluster and grid,distributedandmobilecomputing.Thesecategoriescomprise14topicsthat focus on particular issues. The objective of Euro-Paris to provide a forum within which to promote the development of parallel computing both as an industrial technique and an a- demic discipline, extending the frontier of both the state of the art and the state of practice. The main audience for and participants in Euro-Par are researchers inacademicdepartments,governmentlaboratories,andindustrialorganizations. Euro-Par2010 was the 16th conference in the Euro-Parseries, and was or- nizedbythe Institute forHigh-PerformanceComputingandNetworking(ICAR) of the Italian National Research Council (CNR), in Ischia, Italy. Previous Euro- ParconferencestookplaceinStockholm,Lyon,Passau,Southampton,Toulouse, Munich, Manchester, Padderborn, Klagenfurt, Pisa, Lisbon, Dresden, Rennes, Las Palmas, and Delft. Next year the conference will take place in Bordeaux, France. More information on the Euro-Par conference series and organization is available on the wesite http://www.europar.org. As mentioned before, the conference was organized in 14 topics. The paper review process for each topic was managed and supervised by a committee of at least four persons: a Global Chair, a Local Chair, and two members. Some speci'c topics with a high number of submissions were managed by a larger committeewithmoremembers.The'naldecisionsontheacceptanceorrejection ofthesubmitted papersweremadein ameeting ofthe ConferenceCo-chairsand Local Chairs of the topics.

This book constitutes the thoroughly refereed post-proceedings of the 7th International Workshop on Job Scheduling Strategies for Parallel Processing, JSSPP 2001, held in Cambridge, MA, USA, in June

2001. The 11 revised full papers presented were carefully selected and improved during two rounds of reviewing and revision, and present state-of-the-art results in the area.
"Large-scale enterprise, cloud, and virtualized computing systems have introduced serious performance challenges. Now, internationally renowned performance expert Brendan Gregg has brought together proven methodologies, tools, and metrics for analyzing and tuning even the most complex environments. Systems Performance: Enterprise and the Cloud focuses on Linux® and Unix® performance, while illuminating performance issues that are relevant to all operating systems. You'll gain deep insight into how systems work and perform, and learn methodologies for analyzing and improving system and application performance. Gregg presents examples from bare-metal systems and virtualized cloud tenants running Linux-based Ubuntu®, Fedora®, CentOS, and the illumos-based Joyent® SmartOSTM and OmniTI OmniOS®. He systematically covers modern systems performance, including the "traditional" analysis of CPUs, memory, disks, and networks, and new areas including cloud computing and dynamic tracing. This book also helps you identify and fix the "unknown unknowns" of complex performance: bottlenecks that emerge from elements and interactions you were not aware of. The text concludes with a detailed case study, showing how a real cloud customer issue was analyzed from start to finish."--Back cover.
"This book discusses non-distributed operating systems that benefit researchers, academicians, and practitioners"--Provided by publisher.
The 17th annual International Symposium on High Performance Systems and Applications (HPCS 2003) and the first OSCAR Symposium were held in Sherbrooke, Quebec Canada, May 11-14, 2003. The proceedings cover various areas of High Performance Computing, from specific scientific applications to computer architecture. OSCAR is an Open Source clustering software suite for building, maintaining, and using high performance clusters.