

## Parallel Solutions Inc

Parallel robots are closed-loop mechanisms presenting very good performances in terms of accuracy, rigidity and ability to manipulate large loads. Parallel robots have been used in a large number of applications ranging from astronomy to flight simulators and are becoming increasingly popular in the field of machine-tool industry. This book presents a complete synthesis of the latest results on the possible mechanical architectures, analysis and synthesis of this type of mechanism. It is intended to be used by students (with over 100 exercises and numerous Internet addresses), researchers (with over 500 references and anonymous ftp access to the code of some algorithms presented in this book) and engineers (for which practical results and applications are presented).

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 searches 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, Benchmarks, 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

This book constitutes the refereed proceedings of 10 international workshops held in conjunction with the merged 1998 IPPS/SPDP symposia, held in Orlando, Florida, US in March/April 1998. The volume comprises 118 revised full papers presenting cutting-edge research or work in progress. In accordance with the workshops covered, the papers are organized in topical sections on reconfigurable architectures, run-time systems for parallel programming, biologically inspired solutions to parallel processing problems, randomized parallel computing, solving combinatorial optimization problems in parallel, PC based networks of workstations, fault-tolerant parallel and distributed systems, formal methods for parallel programming, embedded HPC systems and applications, and parallel and distributed real-time systems.

Massively Parallel Systems (MPSs) with their scalable computation and storage space promises are becoming increasingly important for high-performance computing. The growing acceptance of MPSs in academia is clearly apparent. However, in industrial companies, their usage remains low. The programming of MPSs is still the big obstacle, and solving this software problem is sometimes referred to as one of the most challenging tasks of the 1990's. The 1994 working conference on "Programming Environments for Massively Parallel Systems" was the latest event of the working group WG 10.3 of the International Federation for Information Processing (IFIP) in this field. It succeeded the 1992 conference in Edinburgh on "Programming Environments for Parallel Computing". The research and development work discussed at the conference addresses the entire spectrum of software problems including virtual machines which are less cumbersome to program; more convenient programming models; advanced programming languages, and especially more sophisticated programming tools; but also algorithms and applications.

ParCFD 2001, the thirteenth international conference on Parallel Computational Fluid Dynamics took place in Egmond aan Zee, the Netherlands, from May 21-23, 2001. The specialized, high-level ParCFD conferences are organized yearly on traveling locations all over the world. A strong back-up is given by the central organization located in the USA <http://www.parcfd.org>. These proceedings of ParCFD 2001 represent 70% of the oral lectures presented at the meeting. All published papers were subjected to a refereeing process, which resulted in a uniformly high quality. The papers cover not only the traditional areas of the ParCFD conferences, e.g. numerical schemes and algorithms, tools and environments, interdisciplinary topics, industrial applications, but, following local interests, also environmental and medical issues. These proceedings present an up-to-date overview of the state of the art in parallel computational fluid dynamics.

Company Profiles: Parallel Solutions Inc Parallel Solution of Partial Differential Equations Springer Science & Business Media

Foreword by Bjarne Stroustrup Software is generally acknowledged to be the single greatest obstacle preventing mainstream adoption of massively-parallel computing. While sequential applications are routinely ported to platforms ranging from PCs to mainframes, most parallel programs only ever run on one type of machine. One reason for this is that most parallel programming systems have failed to insulate their users from the architectures of the machines on which they have run. Those that have been platform-independent have usually also had poor performance. Many researchers now believe that object-oriented languages may offer a solution. By hiding the architecture-specific constructs required for high performance inside platform-independent abstractions, parallel object-oriented programming systems may be able to combine the speed of massively-parallel computing with the comfort of sequential programming. Parallel Programming Using C++ describes fifteen parallel programming systems based on C++, the most popular object-oriented language of today. These systems cover the whole spectrum of parallel programming paradigms, from data parallelism through dataflow and distributed shared memory to message-passing control parallelism. For the parallel programming community, a common parallel application is discussed in each chapter, as part of the description of the system itself. By comparing the implementations of the polygon overlay

problem in each system, the reader can get a better sense of their expressiveness and functionality for a common problem. For the systems community, the chapters contain a discussion of the implementation of the various compilers and runtime systems. In addition to discussing the performance of polygon overlay, several of the contributors also discuss the performance of other, more substantial, applications. For the research community, the contributors discuss the motivations for and philosophy of their systems. As well, many of the chapters include critiques that complete the research arc by pointing out possible future research directions. Finally, for the object-oriented community, there are many examples of how encapsulation, inheritance, and polymorphism can be used to control the complexity of developing, debugging, and tuning parallel software.

???????????

Contributed presentations were given by over 50 researchers representing the state of parallel CFD art and architecture from Asia, Europe, and North America. Major developments at the 1999 meeting were: (1) the effective use of as many as 2048 processors in implicit computations in CFD, (2) the acceptance that parallelism is now the 'easy part' of large-scale CFD compared to the difficulty of getting good per-node performance on the latest fast-clocked commodity processors with cache-based memory systems, (3) favorable prospects for Lattice-Boltzmann computations in CFD (especially for problems that Eulerian and even Lagrangian techniques do not handle well, such as two-phase flows and flows with exceedingly multiple-connected domains with a lot of holes in them, but even for conventional flows already handled well with the continuum-based approaches of PDEs), and (4) the nascent integration of optimization and very large-scale CFD. Further details of Parallel CFD'99, as well as other conferences in this series, are available at <http://www.parcfd.org>

ICA3PP 2000 was an important conference that brought together researchers and practitioners from academia, industry and governments to advance the knowledge of parallel and distributed computing. The proceedings constitute a well-defined set of innovative research papers in two broad areas of parallel and distributed computing: (1) architectures, algorithms and networks; (2) systems and applications.

This book constitutes the refereed proceedings of the Third International Euro-Par Conference, held in Passau, Germany, in August 1997. The 178 revised papers presented were selected from more than 300 submissions on the basis of 1101 reviews. The papers are organized in accordance with the conference workshop structure in tracks on support tools and environments, routing and communication, automatic parallelization, parallel and distributed algorithms, programming languages, programming models and methods, numerical algorithms, parallel architectures, HPC applications, scheduling and load balancing, performance evaluation, instruction-level parallelism, database systems, symbolic computation, real-time systems, and an ESPRIT workshop.

The demand for more computing power has been a constant trend in many fields of science, engineering and business. Now more than ever, the need for more and more processing power is emerging in the resolution of complex problems from life sciences, financial services, drug discovery, weather forecasting, massive data processing for e-science, e-commerce and e-government etc. Grid and P2P paradigms are based on the premise to deliver greater computing power at less cost, thus enabling the solution of such complex problems. Parallel Programming, Models and Applications in Grid and P2P Systems presents recent advances for grid and P2P paradigms, middleware, programming models, communication libraries, as well as their application to the resolution of real-life problems. By approaching grid and P2P paradigms in an integrated and comprehensive way, we believe that this book will serve as a reference for researchers and developers of the grid and P2P computing communities. Important features of the book include an up-to-date survey of grid and P2P programming models, middleware and communication libraries, new approaches for modeling and performance analysis in grid and P2P systems, novel grid and P2P middleware as well as grid and P2P-enabled applications for real-life problems. Academics, scientists, software developers and engineers interested in the grid and P2P paradigms will find the comprehensive coverage of this book useful for their academic, research and development activity.

This book constitutes the refereed proceedings of the 4th European Parallel Virtual Machine and Message Passing Interface Users' Group Meeting, PVM/MPI '97, held in Cracow, Poland in November 1997. Parallel Virtual Machine and Message Passing Interface are the most popular tools for programming in accordance with the message passing paradigm which, at present, is considered to be the best way to develop effective parallel programs. The book presents 63 carefully selected papers covering the whole range of PVM/MPI issues. The papers are organized in sections on evaluation and performance, extensions and improvements, implementation, tools, algorithms, and applications in science and engineering.

This book follows the previously published title, Solving Large-scale Problems in Mechanics, edited by M. Papadrakakis. This first volume to be published in the Wiley Series in Solving Large-scale Problems in Mechanics is devoted to high-performance computing using the new generation of computers with parallel and distributed computing capabilities. Parallel and distributed processing is a rapidly growing area of high technology where engineering applications lagged behind hardware advances. New algorithms and codes are required in order to exploit effectively modern computer architectures, as programs suitable for conventional computers achieve very modest performances on these new machines. There is therefore an urgent need to develop and test powerful solution and data handling techniques capable of exploiting the potential of modern computers and of accomplishing the solution of complex engineering problems in an acceptable computing time. This volume intends capturing the latest developments in the field and to serve as an essential reference book on the subject. It comprises a comprehensive state-of-the-art treatment of theory and practice, illustrated by extensive numerical examples.

Plastic films are high-performance materials which play an essential part in modern life. The plastics films industry uses state-of-the-art manufacturing processes and is continuously seeking out new technologies to improve its performance. The understanding of the nature of plastic films, their production techniques, applications and their characterisation is essential for producing new types of plastic films. This handbook has been written to discuss the production and main uses of plastic films. Plastic films are high-performance materials which play an essential part in modern life. Plastic films are mostly used in packaging applications but as will be seen from this book they are also used in the agricultural, medical and engineering fields. The plastics films industry uses state-of-the-art manufacturing processes and is continuously seeking out new technologies to improve its performance. The understanding of the nature of plastic films, their production techniques, applications and their characterisation is essential for producing new types of plastic films. This handbook has been written to discuss the production and main uses of plastic films.

Study the past, if you would divine the future. -CONFUCIUS A well written, organized, and concise survey is an important tool in any newly emerging field of study. This present text is the first of a new series that has been established to promote the publications of such survey books. A survey serves several needs. Virtually every new research area has its roots in several diverse areas and many of the initial fundamental results are dispersed across a wide range of journals, books, and conferences in many different sub fields. A good survey should bring together these results. But just a collection of articles is not enough. Since terminology and notation take many years to become standardized, it is often difficult to master the early papers. In addition, when a new research field has its foundations outside of computer science, all the papers may be difficult to read. Each field has its own view of elegance and its own method of presenting results. A good survey overcomes such difficulties by presenting results in a notation and terminology that is familiar to most computer scientists. A good survey can give a feel for the whole field. It helps identify trends, both successful and unsuccessful, and it should point new researchers in the right direction.

Parallel CFD 2000, the Twelfth in an International series of meetings featuring computational fluid dynamics research on parallel computers, was held May 22-25, 2000 in Trondheim, Norway. Following the trend of the past conferences, areas such as numerical schemes and algorithms, tools and environments, load balancing, as well as interdisciplinary topics and various kinds of industrial applications were all well represented in the work presented. In addition, for the first time in the Parallel CFD conference series, the organizing committee chose to draw special attention to certain subject areas by organizing a number of special sessions. We feel the emphasis of the papers presented at the conference reflect the direction of the research within parallel CFD at the beginning of the new millennium. It seems to be a clear tendency towards increased industrial exploitation of parallel CFD. Several presentations also demonstrated how new insight is being achieved from complex simulations, and how powerful parallel computers now make it possible to use CFD within a broader interdisciplinary setting. Obviously, successful application of parallel CFD still rests on the underlying fundamental principles. Therefore, numerical algorithms, development tools, and parallelization techniques are still as important as when parallel CFD was in its infancy. Furthermore, the novel concepts of affordable parallel computing as well as metacomputing show that exciting developments are still taking place. As is often pointed out however, the real power of parallel CFD comes from the combination of all the disciplines involved: Physics, mathematics, and computer science. This is probably one of the principal reasons for the continued popularity of the Parallel CFD Conferences series, as well as the inspiration behind much of the excellent work carried out on the subject. We hope that the papers in this book, both on an individual basis and as a whole, will contribute to that inspiration. Further details of Parallel CFD'99, as well as other conferences in this series, are available at <http://www.parcfd.org>

This volume gives an overview of the state-of-the-art with respect to the development of all types of parallel computers and their application to a wide range of problem areas. The international conference on parallel computing ParCo97 (Parallel Computing 97) was held in Bonn, Germany from 19 to 22 September 1997. The first conference in this biannual series was held in 1983 in Berlin. Further conferences were held in Leiden (The Netherlands), London (UK), Grenoble (France) and Gent (Belgium). From the outset the aim with the ParCo (Parallel Computing) conferences was to promote the application of parallel computers to solve real life problems. In the case of ParCo97 a new milestone was reached in that more than half of the papers and posters presented were concerned with application aspects. This fact reflects the coming of age of parallel computing. Some 200 papers were submitted to the Program Committee by authors from all over the world. The final programme consisted of four invited papers, 71 contributed scientific/industrial papers and 45 posters. In addition a panel discussion on Parallel Computing and the Evolution of Cyberspace was held. During and after the conference all final contributions were refereed. Only those papers and posters accepted during this final screening process are included in this volume. The practical emphasis of the conference was accentuated by an industrial exhibition where companies demonstrated the newest developments in parallel processing equipment and software. Speakers from participating companies presented papers in industrial sessions in which new developments in parallel computing were reported.

This volume contains the papers presented at the 5th International Workshop on Advanced Parallel Processing Technologies, APPT 2003. This series of workshops is designed to strengthen the cooperation between the German and Chinese institutions active in the area of these technologies. It has continued to grow, providing an excellent forum for reporting advances in parallel processing technologies. The 5th workshop itself addressed the entire gamut of related topics, ranging from the architectural aspects of parallel computer hardware and system software to the applied technologies for novel applications. For this workshop, we received over 191 full submissions from researchers all over the world. All the papers were peer-reviewed in depth and qualitatively graded on their relevance, originality, significance, presentation, and the overall appropriateness for their acceptance. Any concerns raised were discussed in the program committee. The organizing committee did an excellent job in selecting 78 papers (Among them, 21 were short ones) for presentation. In short, the papers included here represent the forefront of research from China, Germany, and the other countries.

Content Description #Includes bibliographical references and index.

Parallel-Algorithms for Regular Architectures is the first book to concentrate exclusively on algorithms and paradigms for programming parallel computers such as the hypercube, mesh, pyramid, and mesh-of-trees. Algorithms are given to solve fundamental tasks such as sorting and matrix operations, as well as problems in the field of image processing, graph theory, and computational geometry. The first chapter defines the computer models, problems to be solved, and notation that will be used throughout the book. It also describes fundamental abstract data movement operations that serve as the foundation to many of the algorithms presented in the book. The remaining chapters describe efficient implementations of these operations for specific models of computation and present algorithms (with asymptotic analyses) that are often based on these operations. The algorithms presented are the most efficient known, including a number of new algorithms for the hypercube and mesh-of-trees that are better than

those that have previously appeared in the literature. The chapters may be read independently, allowing anyone interested in a specific model to read the introduction and then move directly to the chapter(s) devoted to the particular model of interest. Russ Miller is Assistant Professor in the Department of Computer Science, State University of New York at Buffalo. Quentin F. Stout is Associate Professor in the Department of Electrical Engineering and Computer Science at the University of Michigan. Parallel Algorithms for Regular Architectures is included in the Scientific Computation series, edited by Dennis Gannon.

This volume reviews, in the context of partial differential equations, algorithm development that has been specifically aimed at computers that exhibit some form of parallelism. Emphasis is on the solution of PDEs because these are typically the problems that generate high computational demands. The authors discuss architectural features of these computers inasmuch as they influence algorithm performance, and provide insight into algorithm characteristics that allow effective use of hardware.

This book constitutes the refereed proceedings of the 5th European Meeting of the Parallel Virtual Machine and Message Passing Interface Users' Group, PVM/MPI '98, held in Liverpool, UK, in September 1998. The 49 contributed and invited papers presented were carefully reviewed and revised for inclusion in the volume. All current aspects of PVM and MPI are addressed. The papers are organized in topical sections on evaluation and performance, extensions and improvements, implementation issues, tools, and algorithms.

This book constitutes the refereed proceedings of the 7th International Workshop on Advanced Parallel Processing Technologies, APPT 2007, held in Guangzhou, China, in November 2007. The 78 revised full papers presented were carefully reviewed and selected from 346 submissions. All current aspects in parallel and distributed computing are addressed ranging from hardware and software issues to algorithmic aspects and advanced applications. The papers are organized in topical sections.

Discover how to streamline complex bioinformatics applications with parallel computing This publication enables readers to handle more complex bioinformatics applications and larger and richer data sets. As the editor clearly shows, using powerful parallel computing tools can lead to significant breakthroughs in deciphering genomes, understanding genetic disease, designing customized drug therapies, and understanding evolution. A broad range of bioinformatics applications is covered with demonstrations on how each one can be parallelized to improve performance and gain faster rates of computation. Current parallel computing techniques and technologies are examined, including distributed computing and grid computing. Readers are provided with a mixture of algorithms, experiments, and simulations that provide not only qualitative but also quantitative insights into the dynamic field of bioinformatics. Parallel Computing for Bioinformatics and Computational Biology is a contributed work that serves as a repository of case studies, collectively demonstrating how parallel computing streamlines difficult problems in bioinformatics and produces better results. Each of the chapters is authored by an established expert in the field and carefully edited to ensure a consistent approach and high standard throughout the publication. The work is organized into five parts: \* Algorithms and models \* Sequence analysis and microarrays \* Phylogenetics \* Protein folding \* Platforms and enabling technologies Researchers, educators, and students in the field of bioinformatics will discover how high-performance computing can enable them to handle more complex data sets, gain deeper insights, and make new discoveries.

This book presents the refereed proceedings of the Eighth Annual Workshop on Languages and Compilers for Parallel Computing, held in Columbus, Ohio in August 1995. The 38 full revised papers presented were carefully selected for inclusion in the proceedings and reflect the state of the art of research and advanced applications in parallel languages, restructuring compilers, and runtime systems. The papers are organized in sections on fine-grain parallelism, interprocedural analysis, program analysis, Fortran 90 and HPF, loop parallelization for HPF compilers, tools and libraries, loop-level optimization, automatic data distribution, compiler models, irregular computation, object-oriented and functional parallelism.

Proceedings -- Parallel Computing.

We are proud to introduce the proceedings of the Sixth International Conference on Parallel Problem Solving from Nature, PPSN VI, held in Paris, France, on 18-20 September 2000. PPSN VI was organized in association with the Genetic and Evolutionary Computing Conference (GECCO'2000) and the Congress on Evolutionary Computation (CEC'2000), reflecting the beneficial interaction between the conference activities in Europe and in the USA in the field of natural computation. Starting in 1990 in Dortmund, Germany (Proceedings, LNCS vol. 496, Springer, 1991), this biannual meeting has been held in Brussels, Belgium (Proceedings, Elsevier, 1992), Jerusalem, Israel (Proceedings, LNCS vol. 866, Springer, 1994), Berlin, Germany (Proceedings, LNCS vol. 1141, Springer, 1996), and Amsterdam, The Netherlands (Proceedings, LNCS vol. 1498, Springer, 1998), where it was decided that Paris would be the location of the 2000 conference with Marc Schoenauer as the general chair. The scientific content of the PPSN conference focuses on problem solving paradigms gleaned from natural models.

Characteristic for Natural Computing is the metaphorical use of concepts, principles and mechanisms underlying natural systems, such as evolutionary processes involving mutation, recombination, and selection in natural evolution, annealing or punctuated equilibrium processes of many-particle systems in physics, growth processes in nature and economics, collective intelligence in biology, DNA-based computing in molecular chemistry, and multi-cellular behavioral processes in neural and immune networks.

This IMA Volume in Mathematics and its Applications ALGORITHMS FOR PARALLEL PROCESSING is based on the proceedings of a workshop that was an integral part of the 1996-97 IMA program on "MATHEMATICS IN HIGH-PERFORMANCE COMPUTING." The workshop brought together algorithm developers from theory, combinatorics, and scientific computing. The topics ranged over models, linear algebra, sorting, randomization, and graph algorithms and their analysis. We thank Michael T. Heath of University of Illinois at Urbana (Computer Science), Abhiram Ranade of the Indian Institute of Technology (Computer Science and Engineering), and Robert S. Schreiber of Hewlett Packard Laboratories for their excellent work in organizing the workshop and editing the proceedings. We also take this opportunity to thank the National Science Foundation (NSF) and the Army Research Office (ARO), whose financial support made the workshop possible. A vner Friedman Robert Gulliver v PREFACE The Workshop on Algorithms for Parallel Processing was held at the IMA September 16 - 20, 1996; it was the first workshop of the IMA year dedicated to the mathematics of high performance computing. The workshop organizers were Abhiram Ranade of The Indian Institute of Technology, Bombay, Michael Heath of the University of Illinois, and Robert Schreiber of Hewlett Packard Laboratories. Our idea was to bring together researchers who do innovative, exciting, parallel algorithms research on a wide range of topics, and by sharing insights, problems, tools, and methods to learn something of value from one another.

This book constitutes thoroughly refereed post-conference proceedings of the workshops of the 17th International Conference on Parallel Computing, Euro-Par 2011, held in Bordeaux,

France, in August 2011. The papers of these 12 workshops CCPI, CGWS, HeteroPar, HiBB, HPCVirt, HPPC, HPSS HPCF, PROPER, CCPI, and VHPC focus on promotion and advancement of all aspects of parallel and distributed computing.

This IMA Volume in Mathematics and its Applications PARALLEL SOLUTION OF PARTIAL DIFFERENTIAL EQUATIONS is based on the proceedings of a workshop with the same title. The workshop was an integral part of the 1996-97 IMA program on "MATHEMATICS IN HIGH-PERFORMANCE COMPUTING." I would like to thank Petter Bjørstad of the Institutt for Informatikk, University of Bergen and Mitchell Luskin of the School of Mathematics, University of Minnesota for their excellent work as organizers of the meeting and for editing the proceedings. I also take this opportunity to thank the National Science Foundation (NSF), Department of Energy (DOE), and the Army Research Office (ARO), whose financial support made the workshop possible. Willard Miller, Jr., Professor and Director v PREFACE The numerical solution of partial differential equations has been of major importance to the development of many technologies and has been the target of much of the development of parallel computer hardware and software. Parallel computers offer the promise of greatly increased performance and the routine calculation of previously intractable problems. The papers in this volume were presented at the IMA workshop on the Parallel Solution of PDE held during June 9-13, 1997. The workshop brought together leading numerical analysts, computer scientists, and engineers to assess the state-of-the-art and to consider future directions.

Asia has become the new battleground for the war against HIV/AIDS. The magnitude of the potential public health problems caused by AIDS in this populous continent may become a catastrophic disaster. A 10% rate of prevalence of HIV-1 in India and China alone would mean more than 200 million people are infected with HIV. AIDS in Asia is useful as a comprehensive, up-to-date AIDS reference book for public health and medical professionals. This volume provides concrete information on the diagnosis, treatment, care, prevention and impact of AIDS. Part I contains 'Snapshots of HIV/AIDS in Asia.' Countries and regions included in this section are: Thailand, India, China, Japan, Vietnam, Laos, Cambodia, Myanmar, Malaysia, Indonesia, the Philippines, Pakistan, Nepal, Sri Lanka, Hawaii and the Pacific Islands, Hong Kong and Taiwan. Part II addresses the molecular epidemiology of HIV/AIDS. Part III deals with the diagnosis, surveillance and projected scenarios of the AIDS epidemic. Part IV outlines prevention efforts and treatment options. Part V provides an overview of the ongoing collaborative efforts involved in several different nations in the worldwide war against AIDS. This volume will be invaluable to all the public health professionals and researchers working in this field. "...the book is a useful addition to the HIV/AIDS literature." "AIDS in Asia offers a comprehensive, interesting overview of the epidemic there and of general issues that will influence its progression." -Roger Detels, MD, MS, University of California-Los Angeles The Journal of the American Medical Association, Book Review, 293:15

This book constitutes the thoroughly refereed post-proceedings of the 8th International Workshop on Applied Parallel Computing, PARA 2006. It covers partial differential equations, parallel scientific computing algorithms, linear algebra, simulation environments, algorithms and applications for blue gene/L, scientific computing tools and applications, parallel search algorithms, peer-to-peer computing, mobility and security, algorithms for single-chip multiprocessors.

Data Mining: Opportunities and Challenges presents an overview of the state of the art approaches in this new and multidisciplinary field of data mining. The primary objective of this book is to explore the myriad issues regarding data mining, specifically focusing on those areas that explore new methodologies or examine case studies. This book contains numerous chapters written by an international team of forty-four experts representing leading scientists and talented young scholars from seven different countries.

A survey of products and research projects in the field of highly parallel, optical and neural computers in the USA. It covers operating systems, language projects and market analysis, as well as optical computing devices and optical connections of electronic parts.

Brings together, analyzes, and contextualizes the latest findings and practical applications Polyphosphazenes, an emerging class of polymers, include macromolecules, which have been proven to be biocompatible, biodegradable, and bioactive. Their unprecedented structural diversity and unique properties make them suitable as vaccine adjuvants, microencapsulating agents, biodegradable materials, scaffolds for tissue engineering, biocompatible coatings, and carriers for gene delivery. Polyphosphazenes for Biomedical Applications offers a thorough review of polyphosphazene research findings in the life sciences, chemistry, and chemical engineering. It emphasizes biomedical applications as well as recent advances in polyphosphazene development such as high-throughput discovery and the latest controlled methods of synthesis. The book brings together, analyzes, and contextualizes a wealth of knowledge that previously could only be found scattered throughout the scientific literature. Following two introductory chapters, the book reviews: Vaccine delivery and immunomodulation Biomaterials Drug delivery systems Biodetection Well-defined polyphosphazenes: synthetic aspects and novel molecular architectures All the chapters have been written by leading researchers in the field. Editor Alexander Andrianov, who has led the effort to commercialize polyphosphazenes for biomedical applications, has carefully reviewed and edited all chapters to ensure readability, accuracy, and thoroughness. Polyphosphazenes for Biomedical Applications is not only intended for researchers working in polyphosphazene chemistry, but also for all researchers seeking solutions to problems arising in the areas of biomaterials, drug delivery systems, and controlled release formulations.

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.

[Copyright: ad8e07f1b7675e156cd4b7c9138e610d](http://www.elsevier.com/locate/0167-1875(2013)00000-0)