

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign UrbanaUSA

DOI <https://doi.org/10.1007/978-0-387-09766-4>

Copyright Information Springer
Science+Business Media, LLC 2011

Publisher Name Springer, Boston,
MA

Print ISBN 978-0-387-09844-9

Number Of Entries 796

Contents in Order

b Inatio Molecular Dynamics
Access Anomaly
Actors
Affinity Scheduling
Ajtai-Komlos-Szemerédi Sorting
Network
AKS Network
AKS Sorting Network
Algebraic Multigrid
Algorithm Engineering
Algorithmic Skeletons
All Prefix Sums
Allen and Kennedy Algorithm
Allgather
All-to-All
All-to-All Broadcast
Altivec
AMD Opteron Processor Barcelona
Amdahl's Argument
Amdahl's Law
AMG
Analytics, Massive-Scale
Anomaly Detection
Anton, A Special-Purpose Molecular
Simulation Machine
Application-Specific Integrated
Circuits
Applications and Parallelism
Architecture Independence
Area-Universal Networks
Array Languages
Array Languages, Compiler
Techniques for Asynchronous
Iterations
Asynchronous Iterative Algorithms
Asynchronous Iterative Computations
ATLAS (Automatically Tuned Linear
Algebra Software)
Atomic Operations
Automated Empirical Optimization
Automated Empirical Tuning

Automated Performance Tuning
Automated Tuning
Automatically Tuned Linear Algebra
Software (ATLAS)
Autotuning
Backpressure
Bandwidth-Latency Models (BSP,
Logp)
Banerjee's Dependence Test
Barnes-Hut
Barriers
Basic Linear Algebra Subprograms
(BLAS)
Behavioral Equivalences
Behavioral Relations
Benchmarks
Beowulf Clusters
Beowulf-Class Clusters
Bernstein's Conditions
Bioinformatics
Bisimilarity
Bisimulation
Bisimulation Equivalence
Bitonic Sort
Bitonic Sorting Network
Bitonic Sorting, Adaptive
BLAS (Basic Linear Algebra
Subprograms)
Blocking
Blue CHiP
Blue CHiP Project
Blue Gene/L
Blue Gene/P
Blue Gene/Q
Branch Predictors
Brent's Law
Brent's Theorem
Broadcast
BSP
BSP (Bulk Synchronous Parallelism)
Bulk Synchronous Parallelism (BSP)
Bus: Shared Channel
Buses and Crossbars
Butterfly
C*

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign Urbana USA

Cache Affinity Scheduling
Cache Coherence
Cache-Only Memory Architecture (COMA)
Caches, NUMA
Calculus of Mobile Processes
Carbon Cycle Research
Car-Parrinello Method
CDC 6600
Cedar Multiprocessor
CELL
Cell Broadband Engine Processor
Cell Processor
Cell/B.E.
Cellular Automata
Chaco
Chapel (Cray Inc. HPCS Language)
Charm++
Checkpoint/Restart
Checkpointing
Checkpoint-Recovery
CHiP Architecture
CHIP Computer
Cholesky Factorization
Cilk
Cilk Plus
Cilk++
Cilk-1
Cilk-5
Cilkscreen
Cluster File Systems
Cluster of Workstations
Clusters
CM Fortran
Cm* The First Non-Uniform Memory Access Architecture
CML
CM-Lisp
CnC
Coarray Fortran
Code Generation
Collect
Collective Communication
Collective Communication, Network Support For COMA (Cache-Only Memory Architecture) Combinatorial Search
Commodity Clusters
Communicating Sequential Processes (CSP)
Community Atmosphere Model (CAM)
Community Climate Model (CCM)
Community Climate System Model
Community Climate System Model (CCSM)
Community Earth System Model (CESM)
Community Ice Code (CICE)
Community Land Model (CLM)
Compiler Optimizations for Array Languages
Compilers
Complete Exchange
Complex Brent Processing
Computational Biology
Computational Chemistry
Computational Models
Computational Sciences
Computer Graphics
Computing Surface
Concatenation
Concurrency Control
Concurrent Collections Programming Model
Concurrent Logic Languages
Concurrent ML
Concurrent Prolog
Configurable, Highly Parallel
Computer
Congestion Control
Congestion Management
Connected Components Algorithm
Connection Machine
Connection Machine Fortran
Connection Machine Lisp
Consistent Hashing
Control Data 6600
Coordination
Copy
Core2-Duo /Core2-Quad Processors
COW
Crash Simulation
Cray MTA
Cray Red Storm
Cray SeaStar Interconnect
CRAY T3E

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign UrbanaUSA

Cray Vector Computers
Cray XMT
Cray XT Series
Cray XTS
Cray XT3 and Cray XT Series of Supercomputers
Cray XT4
Cray XT4 and Seastar 3-0 Torus Interconnect
Cray AT5
Cray XT6
Critical Race
Critical Sections
Crossbar
CS-2
CSP (Communicating Sequential Processes)
Cyclops
Cyclops-64
Cydra 5
DAG Scheduling
Data Analytics
Data Centers
Data Distribution
Data Flow Computer Architecture
Data Flow Graph
Data Mining
Data Race Detection
Data Starvation Crisis
Dataflow Supercomputer
Data-Parallel Execution Extensions
Deadlock Detection
Deadlocks
Debugging
DEC Alpha
Decentralization
Decomposition
Deep Analytics
Denelcor HEP
Dense Linear System Solvers
Dependence Abstractions
Dependence Accuracy
Dependence Analysis
Dependence Approximation
Dependence Cone
Dependence Direction Vector
Dependence Level
Dependence Polyhedron
Dependences
Detection of DOALL Loops
Determinacy
Determinacy Race
Determinism
Deterministic Parallel Java
Direct Schemes
Distributed Computer
Distributed Hash Table (DHT)
Distributed Logic Languages
Distributed Memory Computers
Distributed Process Management
Distributed Switched Networks
Distributed-Memory Multiprocessor
Ditonic Sorting
DLPAR
Doall Loops
Domain Decomposition
DPJ
DR
Dynamic Logical Partitioning for POWER Systems
Dynamic LPAR
Dynamic Reconfiguration
Earth Simulator
Eden
Eigenvalue and Singular-Value Problems
EPIC Processors
Erlangen General purpose Array (EGPA)
ES
Ethernet
Event Stream Processing
Eventual Values
Exact Dependence
Exaflop Computing
Exaop Computing
Exascale Computing
Execution Ordering
Experimental Parallel Algorithmics
Extensional Equivalences
Fast Fourier Transform (FFT)
Fast Multipole Method (FMM)
Fast Poisson Solvers
Fat Tree
Fault Tolerance
Fences
FFT (Fast Fourier Transform)
Fast Algorithm for the Discrete Fourier

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign UrbanaUSA

Transform (DFT)
FFTW
File Systems
Fill-Reducing Orderings
First-Principles Molecular Dynamics
Fixed-Size Speedup
FLAME
Floating Point Systems FPS-120B and Derivatives
Flow Control
Flynn's Taxonomy
Forall Loops
FORGE
Formal Methods-Based Tools for Race, Deadlock, and Other Errors
Fortran 90 and Its Successors
Fortran, Connection Machine
Fortress (Sun HPCS Language)
Forwarding
Fujitsu Vector Computers
Fujitsu Vector Processors
Fujitsu VPP Systems
Functional Decomposition Functional Languages
Futures
GA
Gather
Gather-to-All
Gaussian Elimination
GCD Test
Gene Networks Reconstruction
Gene Networks Reverse-Engineering
Generalized Meshes and Tori
Genome Assembly
Genome Sequencing
3GIO
Glasgow Parallel Haskell (GpH)
Global Arrays
Global Arrays Parallel Programming Toolkit
Gossiping
GpH (Glasgow Parallel Haskell)
GRAPE
Graph Algorithms
Graph Analysis Software
Graph Partitioning
Graph Partitioning Software
Graphics Processing Unit
Green Flash: Climate Machine
(LBNL)
Grid Partitioning
Gridlock
Group Communication
Gustafson's Law
Gustafson-Barsis Law
Half Vector Length
Hang
Harmful Shared-Memory Access
Haskell
Hazard (in Hardware)
HDFS
HEP, Denelcor
Heterogeneous Element Processor
Hierarchical Data Format
High Performance Fortran (HPF)
High-Level I/O Library
High-Performance I/O
Homology to Sequence Alignment, From
Horizon
HPC Challenge Benchmark
HPF (High Performance Fortran)
HPS Microarchitecture
HT
HT3.10
Hybrid Programming With SIMPLE
Hypercube
Hypercubes and Meshes
Hypergraph Partitioning
Hyperplane Partitioning
HyperTransport
IBM Blue Gene Supercomputer
IBM Power
IBM Power Architecture
IBM PowerPC
IBM RS/6000 SP
IBM SP
IBM SPI
IBM SP2
IBM SP3
IBM System/360 Model 91
IEEE 802.3
Illegal Memory Access
Illiac IV
ILUPACK
Impass
Implementations of Shared Memory in Software

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign UrbanaUSA

Index
InfiniBand
Instant Replay
Instruction-Level Parallelism
Instruction Systolic Arrays
Intel Celeron
Intel Core Microarchitecture, x86
Processor Family
Intel Parallel Inspector Intel® Parallel
Studio
Intel® Thread Profiler
Intel® Threading Building Blocks
(TBB)
Interactive Parallelization
Interconnection Network
Interconnection Networks
Internet Data Centers
Inter-Process Communication
I/O
iPSC
Isoefficiency
JANUS FPGA-Based Machine
Java
JavaParty
Job Scheduling
k-ary n-cube
k-ary n-fly
k-ary n-tree
Knowledge Discovery
KSR
LANai
Languages
LAPACK
Large-Scale Analytics
Latency Hiding
Law of Diminishing Returns
Laws
Layout, Array
LBNL Climate Computer
libflame
Libraries, Numerical
Linda
Linear Algebra Software
Linear Algebra, Numerical
Linear Equations Solvers
Linear Least Squares and Orthogonal
Factorization
Linear Regression
LINPACK Benchmark
Linux Clusters
*Lisp
Lisp, Connection Machine
Little's Law
Little's Lemma
Little's Principle
Little's Result
Little's Theorem
Livermore Loops
Load Balancing
Load Balancing, Distributed Memory
Locality of Reference and Parallel
Processing Lock-Free Algorithms
Locks
Logarithmic-Depth Sorting Network
Logic Languages
LogP Bandwidth-Latency Model
Loop Blocking
Loop Nest Parallelization
Loop Tiling
Loops, Parallel
LU Factorization
Manycore
MapReduce
MasPar
Massively Parallel Processor (MPP)
Massive-Scale Analytics
Matrix Computations
Maude
Media Extensions
Meiko
Memory Consistency Models
Memory Models
Memory Ordering
Memory Wall
MEMSY
Mesh
Mesh Partitioning
Message Passing
Message Passing Interface (MPI)
Message-Passing Performance Models
METIS and ParMETIS
Metrics
Microprocessors
MILC
MIMD (Multiple Instruction, Multiple
Data) Machines
MIMD Lattice Computation
MIN

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign UrbanaUSA

ML
MMX
Model Coupling Toolkit (MCT)
Models for Algorithm Design and Analysis Models of Computation, Theoretical
Modulo Scheduling and Loop
Pipelining
Molecular Evolution
Monitors
Monitors, Axiomatic Verification of
Moore's Law
MPI (Message Passing Interface)
MP1-2V0
MPHO
MPP
Mul-T
Multicomputers
Multicore Networks
Multiflow Computer
Multifrontal Method
Multi-Level Transactions
Multilisp
Multimedia Extensions
Multiple-Instruction Issue
Multiprocessor Networks
Multiprocessor Synchronization
Multiprocessors
Multiprocessors, Symmetric
MultiScheme
Multistage Interconnection Networks
Multi-Streamed Processors
Multi-Threaded Processors
Mumps
MUMPS
Mutual Exclusion
Myri-10G
Myricom
Myrinet
NAMD (NAnoscale Molecular Dynamics)
NAnoscale Molecular Dynamics (NAMD)
NAS Parallel Benchmarks
nCUBE
N-Body Computational Methods
NEC SX Series Vector Computers
NESL
Nested Loops Scheduling
Nested Spheres of Control
NetCDF UO Library, Parallel
Network Adapter
Network Architecture
Network interfaces
Network Obliviousness
Network of Workstations
Network offload
Networks Direct
Networks, Fault-Tolerant
Networks, Multistage
NI (Network Interface)
NIC (Network Interface Controller or Network Interface Card)
Node Allocation
Non-Blocking Algorithms
Nondeterminator
Nonuniform Memory Access (NUMA)
Machines
NOW
NUMA Caches
Numerical Algorithms
Numerical Libraries
Numerical Linear Algebra
NVIDIA GPU
NWChem
Omega Calculator
Omega Library
Omega Project
Omega Test
One-to-All Broadcast
Open Distributed Systems
OpenMP
OpenMP Profiling with Ompp
OpenSHMEM-Toward a Unified
RMA Model
Operating System Strategies
Optimistic Loop Parallelization
Orthogonal Factorization
OS Jitter
OS, Light-Weight
Out-of-Order Execution Processors
Overdetermined Systems
Overlay Network
Owicki-Gries Method of Axiomatic Verification
Paraphrase
Parallel Communication Models
Parallel Computing

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign Urbana USA

Parallel U/O Library (PIO)
Parallel Ocean Program (POP)
Parallel Operating System
Parallel Prefix Algorithms
Parallel Prefix Sums
Parallel Random Access Machines (PRAM)
Parallel Skeletons
Parallel Tools Platform
Parallelism Detection in Nested Loops,
Optimal
Parallelization
Parallelization, Automatic
Parallelization, Basic Block
Parallelization, Loop Nest
ParaMETIS
PARDISO
PARSEC Benchmarks
Partial Computation
Particle Dynamics
Particle Methods
Partitioned Global Address Space (PGAS) Languages
PASM Parallel Processing System
Path Expressions
PaToH (Partitioning Tool for Hypergraphs)
Partitioning Tool for Hypergraphs (PaToH)
PC Clusters
PCI Express
PCIe
PCI-E
PCI-Express
Peer-to-peer
Pentium
PERCS System Architecture
Perfect Benchmarks
Performance Analysis Tools
Performance Measurement
Performance Metrics
Periscope
Personalized All-to-All Exchange
Petaflop Barrier
Petascale Computer
Petri Nets
PETS (Portable, Extensible Toolkit for Scientific Computation)
PGAS (Partitioned Global Address Space) Languages
Space) Languages
Phylogenetic inference
Phylogenetics
Pi-Calculus
Pipelining
Place-Transition Nets
PLAPACK
PLASMA
PMPI Tools
Pnetcdf
Point-to-Point Switch
Polaris
Polyhedra Scanning Polyhedron Model
Polytope Model
Position Tree
POSIX Threads (Pthreads)
Power Wall
PRAM (Parallel Random-Access Machines)
Preconditioners for Sparse Iterative Methods
Prefix
Prefix Reduction
Problem Architectures
Process Algebras
Process Calculi
Process Description Languages
Process Synchronization
Processes, Tasks, and Threads
Processor Allocation
Processor Arrays
Processors-in-Memory
Profiling
Profiling with OmpP, OpenMP
Program Graphs
Programmable Interconnect Computer
Programming Languages
Programming Models
Prolog
Prolog Machines
Promises
Protein Docking
Pthreads (POSIX Threads)
PVM (Parallel Virtual Machine)
QCD apeNEXT Machines
QCD (Quantum Chromodynamics) Computations
QCD Machines
QCDSM and QCDOC Computers

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign UrbanaUSA

QsNet
Quadrics
Quantum Chemistry
Quantum Chromodynamics (QCD) Computations
Quicksort
Race
Race Conditions
Race Detection Techniques
Race Detectors for Cilk and Cilk++
Programs
Race Hazard
Radix Sort
Rapid Elliptic Solvers
Reconfigurable Computer
Reconfigurable Computers
Reconstruction of Evolutionary Trees
Reduce and Scan
Relaxed Memory Consistency Models
Reliable Networks
Rendezvous
Reordering
Resource Affinity Scheduling
Resource Management for Parallel Computers
Rewriting Logic
Ring
Roadrunner Project, Los Alamos
Router Architecture Router-Based Networks
Routing (Including Deadlock Avoidance)
R-Stream Compiler
Run Time Parallelization
Runtime System Scalability
Scalable Coherent Interface (SCI)
ScaLAPACK
Scalasca
Scaled Speedup
Scan for Distributed Memory, Message-Passing Systems
Scan, Reduce and Scatter
Scheduling Scheduling Algorithms
SCI (Scalable Coherent Interface)
Semantic Independence
Semaphores
Sequential Consistency
Server Farm
Shared Interconnect
Shared Virtual Memory Shared-Medium Network
Shared-Memory Multiprocessors
SHMEM
SIGMA-1
SIMD (Single Instruction, Multiple Data) Machines
SIMD Extensions
SIMD ISA
Single System Image
Singular Value Decomposition (SVD)
Sisal
Small World Network Analysis and Partitioning (SNAP) Framework
SNAP (Small World Network Analysis and Partitioning) Framework
SoC (System on Chip)
Social Networks
Software Autotuning
Software Distributed Shared Memory
Sorting
Space-Filling Curves
SPAI (SParse Approximate Inverse)
Spanning Tree, Minimum Weight
Sparse Approximate Inverse Matrix
Sparse Direct Methods
Sparse Gaussian Elimination
Sparse Iterative Methods, Preconditioners for
SPEC Benchmarks
SPEC HPC2002
SPEC HPC96
SPEC MP12007
SPEC OMP2001
Special-Purpose Machines
Speculation
Speculation, Thread-Level
Speculative Multithreading (SM)
Speculative Parallelization
Speculative Parallelization of Loops
Speculative Run-Time Parallelization
Speculative Threading
Speculative Thread-Level Parallelization
Speedup
SPIKE
Spiral
SPMD Computational Model
SSE

Encyclopedia of Parallel Computing

David Padua - University of Illinois at Urbana-Champaign UrbanaUSA

Stalemate
State Space Search
Stream Processing
Stream Programming Languages
Strong Scaling
Suffix Trees
Superlinear Speedup
SuperLU
Supernode Partitioning
Superscalar Processors
SWARM: A Parallel Programming
Framework for Multicore Processors
Switch Architecture
Switched Medium Network
Switching Techniques
Symmetric Multiprocessors
Synchronization
System Integration

System on Chip (SoC)
Systems Biology, Network Inference
in
Systolic Architecture
Systolic Arrays
Task Graph Scheduling
Task Mapping, Topology Aware
Tasks
TAU
TAU Performance System®
TBB (Intel Threading Building
Blocks)
Tensilica
Tera MTA
Terrestrial Ecosystem Carbon
Modeling
Terrestrial Ecosystem Modeling
The High-Performance Substrate
Theory of Mazurkiewicz-Traces
Thick Ethernet
Thin Ethernet
Thread Level Speculation (TLS)
Parallelization
Thread-Level Data Speculation
(TLDS)
Thread-Level Speculation
Threads Tiling
Titanium
TLS
TOP500

Topology Aware Task Mapping
Torus
Total Exchange
Trace Scheduling
Trace Theory
Tracing
Transactional Memories
Transactions, Nested
Transpose
TStreams
Tuning and Analysis Utilities
Ultracomputer, NYU
Uncertainty Quantification
Underdetermined Systems
Unified Parallel C
Unimodular Transformations
Universal VLSI Circuits
Universality in VLSI Computation
UPC
Use-Def Chains
Vampir
Vampir 7
Vampir NG
VampirServer
VampirTrace
Vector Extensions, Instruction-Set
Architecture (ISA)
Vectorization
Verification of Parallel Shared-
Memory Programs,
Owicki-Gries Method of Axiomatic
View from Berkeley
Virtual Shared Memory
VLIW Processors
VLSI Algorithmics
VLSI Computation
VNG
Warp and iWarp
Wavefront Arrays
Weak Scaling Whole Program
Analysis
Work-Depth Model
Workflow Scheduling
Z-Level Programming Language
ZPL