

Parallel Computing: Concepts and keywords

Jesper Larsson Träff

January 8, 2018

This checklist lists concepts and keywords from the bachelor lecture (VU) on Parallel Computing (184.710) roughly in order of appearance. The list can be used to check your understanding of the material; all keywords can come up during the examination.

Models, Performance

“Free lunch”	Prefix-sums problem
Single-core performance	Sorting (mergesort, Quicksort, countingsort, bucketsort)
SPEC benchmarks	Architecture model
Deep pipelining	Execution/cost model
Superscalar execution	RAM (Random Access Machine)
Data parallelism	“von Neumann bottleneck”
Out-of-order execution	Bus
Branch prediction	Vector computer
Cache	Shared-memory model
Instruction set	PRAM (Parallel Random Access Machine)
Parallel architecture	EREW (Exclusive Read Exclusive Write)
Architecture model	PRAM
Programming model	CREW (Concurrent Read Exclusive Write)
Performance portability	PRAM
Parallel computing (definition)	CRCW (Concurrent Read Concurrent Write)
Distributed computing (definition)	PRAM
Concurrent computing (definition)	CRCW COMMON/ARBITRARY/PRIORITY
Coordination	conflict resolution
Synchronization	PRAM fast maximum
Scheduling	UMA (Uniform Memory Access)
Mapping	NUMA (Non-Uniform Memory Access)
Load balancing	Distributed memory model
Communication	Communication network
Explicit parallelism	Atomic operation
Automatic parallelization	Cellular automaton
Matrix-vector multiplication	Flynn’s taxonomy
Matrix-matrix multiplication	SISD (Single Instruction, Single Data)
LU factorization	MISD (Multiple Instruction, Single Data)
Stencil computation	SIMD (Single Instruction, Multiple Data)
5-point stencil	MIMD (Multiple Instruction, Multiple Data)
Data distribution	Programming model
Merging or ordered sequences	SPMD (Single Program, Multiple Data)
	PGAS (Partitioned Global Address Space)

OpenMP	Backwards dependency
MPI (Message Passing Interface)	Nested parallelism
Cilk	Barrier synchronization
Speedup, $S_p(n)$	Map
$T_{\text{seq}}(n)$	Scatter
$T_{\text{par}}(p, n)$	Gather
Work	Reduce
Measured time	Map-Reduce
Worst case complexity	“Seven Dwarves” Berkeley patterns
Problem complexity	Stencil pattern
Best possible/best known algorithm	Data distribution
Absolute speedup	Block data distribution
Relative speedup	Cyclic data distribution
$T_{\infty}(n)$	Block-cyclic data distribution
Parallelism	Domain decomposition
Perfect speedup	Task parallelism
“Embarrassingly parallel”	Functional parallelism
“Pleasantly parallel”	Pipeline, linear
Load imbalance	Pipeline, complex
Simulation (argument)	Pipeline performance
SimGrid	Pipeline balancing
“Modest potential”	Master-worker pattern
Cost	Distributed work pool
Cost optimality	Work-dealing
Work optimality	Work-stealing
Overhead	Termination detection
Idle time	Acyclic task graph (Directed Acyclic Graph, DAG)
Granularity	Immediate dependency
“Fine grained” parallel computation	Dependent tasks
“Coarse grained” parallel computation	Independent tasks
Sequential fraction	Schedule, optimal, Makespan
Parallel Fraction	Scheduling problem
Amdahl’s law	Fork-join parallelism
Sequential bottleneck	Input node(s), output node(s)
Scaled speedup	Topological order
Scalability	Heaviest path
Strong scaling	Work law
Weak scaling	Depth law
Parallel efficiency	Work-time presentation
Iso-efficiency function	Merging ordered sequences in parallel
	Stable merge
	Lexicographic ordering
	Rank, ranking
	Ranking, binary search
	Load imbalance, bad segments
	Co-ranks, co-ranking
	Oblivious merging
	Bitonic sequence
Parallelization patterns	
Parallelization pattern/paradigm	
Loop parallelism	
“Parallel loop”	
Independent iterations	
Forwards dependency	

Bitonic split
 Bitonic merge
 Comparator network
 Sorting networks, depth, size
 “Cole’s parallel merge sort”
 High-Performance Computing (HPC)
 Thinking Machines CM2, CM5
 Intel iPSC
 MasPar
 5th Generation project, Star Wars
 “Grand Challenge Problems”
 “Lost decade”
 SiCortex, Kautz network
 Moore’s “law” (popular version)
 Moore’s “law” (what Moore observed)
 Top500 list, www.top500.org
 High-Performance LINPACK (HPLINPACK)
 LU-factorization
 FLOPS (Tera-, Peta-, Exa-, Zetta-, ...)
 K Computer, Road Runner, Earth Simulator
 Vienna Scientific Cluster (VSC)
 Empirical “law”: observation, forecast, prediction, prophecy, dictate
 Hardware-oriented efficiency
 Peak performance
 Roofline model
 HPCC Benchmarks (DGEMM, STREAM, PTRANS, FFT, random access)
 Green 500 list
 Graph 500 list
 Superlinear speedup
 Observation, experiment, hypothesis
 Experimental setup, experiment design
 Reproducibility
 Statistical testing
 Memory hierarchy
 Superlinear speedup by memory hierarchy effects
 Superlinear speedup by non-obvious algorithmic differences
 Searching, randomization
 Guided tree search
 Reduction problem
 Prefix sums problem (inclusive, exclusive)
 Collective operation pattern
 Reduction-to-one, reduction-to-all, reduction-with-scatter
 Scan (inclusive, exclusive)
 Array compaction

Load balancing
 Quicksort partitioning
 Binomial tree, binomial coefficient
 Prefix sums algorithms (recursive, iterative)
 Hillis-Steele algorithm
 $2n - 2 \leq s + t$ trade-off for prefix sums
 Blocking technique
 List ranking problem
 Parallel breadth first search (BFS)
 Depth first search (DFS)

pthread, OpenMP

Memory consistency
 Program order
 Interleaving (of instructions)
 Cache
 TLB (Translation Lookaside Buffer)
 Write buffer
 Cache coherency problem
 Memory consistency modeling problem
 Cache line
 Cache block
 Temporal locality
 Spatial locality
 Cache hit
 Cache miss
 Write allocate
 Write-through cache
 Write-back cache
 Directly mapped cache
 Fully associative cache
 Set-associative cache
 Cold (compulsory) cache miss
 Capacity miss
 Conflict miss
 Eviction policy
 LRU (Least Recently Used)
 LFU (Least Frequently Used)
 Sequential consistency
 Cache coherency definition
 Cache coherence protocols
 Update based, Invalidation based
 Snooping (bus based)
 Directory based
 False sharing
 Relaxed consistency models
 Memory fence/barrier

Prefetching
 Multithreading
 SIMD extensions (MMX, SSE, AVX)
 Thread model
 POSIX (Portable Operating Systems Interface for uniX)
pthreads
 Fork-join parallelism
 Peer threads
 Spawn
 Join
 Race condition
 Semaphores, locks, monitors, atomic instructions
 Critical section
 Condition variable
 Wait, Signal, Broadcast
 Deadlock freedom
 Starvation freedom
 Bounded waiting
 Fairness
 Thread safety
 State, static variables
 Stratified locks
 Backoff
 Readers-writers lock
 Spurious wakeup
 Spin lock (pragmatics)
 Load imbalance, prime numbers
 Shared counter
 Fetch-and-increment
 Progress guarantee
 Wait-free algorithms
 Lock-free algorithms
 Atomic instructions
 Test-and-set
 Compare-and-swap
 Work-stealing: randomized, deterministic
 C11, C++11 threads
 OpenMP
 Work-sharing (constructs)
 Loop, sections, single, master construct
 Implicit barrier
 OpenMP clauses
 Canonical form loop
 Parallel loop schedule (static, dynamic, guided)
 Independence conditions
 Bernstein's conditions

True (flow) dependency
 Anti-dependency
 Output dependency
 Loop carried dependency
 Reduction
 Ordered iteration

Cilk

Spawn node, spawn edge
 Sync edge
 Fully strict computation
 Finish edge
 Strict computation
 Single control variable canonical loop form
 Cut-off unit
 Mergesort, Quicksort
 Divide-and-conquer merge
 Depth, work analysis

MPI

Communication driven model
 Data distribution driven model
 Network topology
 Communication capabilities
 One-ported
 k -ported
 Uni-directional
 Bi-directional
 Telephone model
 Send-receive bi-directional model
 Routing
 Switching
 Latency-bandwidth model
 Network contention
 Direct network
 Indirect network
 Diameter
 Degree
 Bisection width
 Linear array, ring topology
 Tree topology
 Fat-tree topology
 Torus, mesh topology
 Hypercube topology
 Kautz network
 Fully connected network

Broadcast lower bound, fully connected network	Communication context
Diameter bound	Message non-determinism, wildcards
Deterministic, oblivious routing	Unsafe communication
Adaptive routing	Safe programming
Minimal routing algorithm	Send-receive communication
Transmission cost model	Enforcing completion
Pipelining	Testing completion
Hybrid architectures, heterogeneous networks	Communication progress, progress engine
Message-passing model	MPI datatypes
Shared nothing	Derived (user-defined) datatypes
CSP (Communicating Sequential Processes)	Contiguous, vector, index, struct types
MPI communication models	Origin, target process
Point-to-point communication	Communication window
One-sided communication	Communication epoch
Collective communication	Exposure epoch
MPI-IO	Access epoch
Dynamic process management	Active synchronization
MPI Standard	Fence synchronization
MPI Forum	Generalized, pairwise synchronization
Error handling, fault tolerance	Passive synchronization
MPI error codes	Distributed memory binary search
MPI bindings	Distributed memory merging
Communication domain, communicator	Cartesian communicator
Process group	Collective communication operation
MPI object	Collective operation complexity
Encapsulation, library building	MPI collectives
Communicator splitting	Symmetric vs. non-symmetric (rooted) collectives
Point-to-point semantics	Regular vs. irregular collectives
Blocking communication	Distributed matrix-vector multiplication
Non-blocking (immediate) communication	Distributed memory Quicksort
Non-local completion	Distributed memory counting (bucket) sort
Non-overtaking, ordered communication	
Message envelope	