



Instruction-Level Parallel Processor Architectures

Instruction Scheduling

Local and Global Scheduling

Christoph Kessler, IDA,
Linköpings universitet, 2014.



RISC and Instruction-Level Parallel Target Architectures

Christoph Kessler, IDA,
Linköpings universitet, 2014.

CISC vs. RISC



CISC

- Complex Instruction Set Computer
- Memory operands for arithmetic and logical operations possible
- $M(r1+r2) \leftarrow M(r1+r2) * M(r3+disp)$
- Many instructions
- Complex instructions
- Few registers, not symmetric
- Variable instruction size
- Instruction decoding (often done in microcode) takes much silicon overhead
- Example: 80x86, 680x0

RISC

- Reduced Instruction Set Computer
- Arithmetic/logical operations only on registers
- add r1, r2, r1
load r1, r4
load r3+disp, r5
mul r4, r5
store r5, r1
- Few, simple instructions
- Many registers, all general-purpose typ. 32 ... 256
- Fixed instruction size and format
- Instruction decoding hardwired
- Example: POWER, HP-PA RISC, MIPS, ARM, SPARC

C. Kessler, IDA, Linköpings universitet.

3

Instruction-Level Parallel (ILP) architectures



Single-Issue: (can start at most one instruction per clock cycle)

- Simple, pipelined RISC processors with one or multiple functional units
- e.g. ARM, DLX

Multiple-Issue: (can start several instructions per clock cycle)

- Superscalar processors
 - e.g. Sun SPARC, MIPS R10K, Alpha 21264, IBM Power2, Pentium
- VLIW processors
 - e.g. Multiflow Trace, Cydrome Cydra-5, Intel i860, HP Lx, Transmeta Crusoe; most DSPs, e.g. Philips Trimedia TM32, TI 'C6x
- EPIC processors
 - e.g. Intel Itanium family (IA-64)

C. Kessler, IDA, Linköpings universitet.

4

Pipelined RISC Architectures



- A single instruction is issued per clock cycle
- Possibly several parallel functional units / resources
- Execution of different phases of subsequent instructions overlaps in time. This makes them prone to:
 - data hazards** (may have to delay op until operands ready),
 - control hazards** (may need to flush pipeline after wrongly predicted branch),
 - structural hazards** (required resource(s) must not be occupied)
- Static scheduling (insert NOPs to avoid hazards) vs. Run-time treatment by automatic hazard detection + pipeline stalling

	issue	cycle	PM	Decoder	ALU ₁	DM/ALU ₂	Regs
IF	I ₁	1	IF ₁				
ID	I ₂	2	IF ₂	ID ₁			
EX	I ₃	3	IF ₃	ID ₂	EX ₁		
MEM/EX2	I ₄	4	IF ₄	ID ₃	EX ₂	MEM ₁	
WB	I ₅	5	IF ₅	ID ₄	EX ₃	MEM ₂	WB ₁
	I ₆	6	IF ₆	ID ₅	EX ₄	MEM ₃	WB ₂

C. Kessler, IDA, Linköpings universitet.

Reservation Table, Scheduling Hazards



add:

	read	read	stage	stage	stage	stage	stage	stage	write
	stc1	stc2	0	1	0	1	2	3	result
Time	opnd	opnd							bus
0	X								
1		X							
2			X						
3				X					

mul:

	read	read	stage	stage	stage	stage	stage	stage	write
	stc1	stc2	0	1	0	1	2	3	result
Time	opnd	opnd							bus
0	X								
1		X							
2			X						
3				X					
4					X				
5						X			

Reservation table
specifies required resource
occupations

[Davidson 1975]

t: mul ...

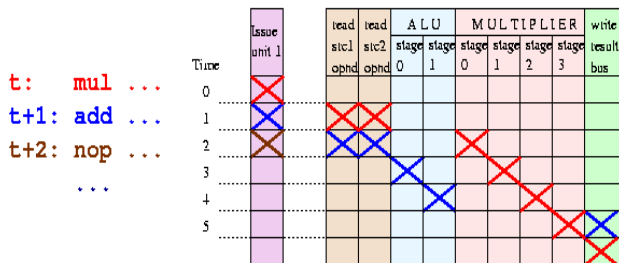
t+1: ...

t+2: add ...

... **structural hazard**
at t=5


Instruction Scheduling (1)

- Map instructions to time slots on issue units (and resources), such that no hazards occur
→ **Global reservation table, resource usage map**




Instruction Scheduling (2)

- Data dependences imply **latency constraints**
→ target-level data flow graph / data dependence graph




Superscalar processor

- Run-time scheduling by instruction dispatcher
 - convenient (sequential instruction stream – as usual)
 - limited look-ahead buffer to analyze dependences, reorder instr.
 - high silicon overhead, high energy consumption
- Example: Motorola MC 88110
2-way, in-order issue superscalar




Dual-Issue (w=2)

- Example (1):
Linear code “mul R1,...; add ...,R2” expands to




Dual-Issue (w=2)

- Example (2):
Linear code “mul R1,...; add ...,R1” expands to




VLIW (Very Long Instruction Word) architecture

- Multiple slots for instructions in long instruction-word
 - Direct control of functional units and resources – low decoding OH
- Compiler (or assembler-level programmer) must determine the schedule statically
 - independence, unit availability, packing into long instruction words
 - Challenging! But the compiler has more information on the program than an on-line scheduler with a limited lookahead window.
 - Silicon- and energy-efficient



Clustered VLIW processor

- E.g., TI C62x, C64x DSP processors
- Register classes
- Parallel execution constrained by operand residence



C. Kessler, IDA, Linköping universitet.

13

EPIC architectures


- Based on VLIW
- Compiler groups instructions to LIW's (**bundles**, fetch units)
- Compiler takes care of resource and latency constraints
- Compiler marks sequences of independent instructions as **instruction groups** by inserting delimiters (**stop bits**)
- Dynamic scheduler assigns resources and reloads new bundles as required

C. Kessler, IDA, Linköping universitet.

14

EPIC Example: Instruction format for TI 'C62x

- Texas Instruments DSP processor series TMS320C62xx
- 1 **fetch packet** (a very long instruction word) has 8 slots
 - may contain up to 8 **instruction groups** (**issue packets**) to be executed in sequence
 - Instruction groups are marked by chaining bits.
 - ▶ Up to 8 instructions in an instruction group
 - Instructions within an instruction group must use disjoint resources (basically, different functional units)
- Example: 3 issue groups { A||B||C }; { D||E||F }; { G||H }




C. Kessler, IDA, Linköping universitet.

15

EPIC Example: Intel IA-64 (Itanium)

- Constraints on bundle contents and placement of delimiters for instruction groups: 24 templates



Instruction types:

M = Memory (on M unit)
I = Integer (on I unit)
F = Floatingpoint (on F unit)
B = Branch (on B unit)
A = supertype of I and M
LX = uses 2 bundle slots, uses I and B units
NOP can replace anything, uses no unit

Functional units:





An implementation of the IA-64 instruction set interface is the Itanium processor series.


C. Kes

Example: Local scheduling for IA-64

- A DAG with a greedy and an optimal schedule



used templates:
MI;I / MMI; / MFI;



used templates:
M;MI / MMF;

Adapted from: S. Haga, R. Barua: EPIC Instruction Scheduling based on Optimal Approaches.

C. Kessler, IDA, Linköping universitet.

17

A Generic ILP-Architecture Model for Retargetable Code Optimization

- Issue width w
 w -way in-order superscalar or size of longest instruction group
- w issue units
- f resources
functional units, internal buses, ...
- Instruction set I
for each instruction y in I , specify its
 - syntax: mnemonic, parameters, types
 - semantics: (tree)pattern in terms of IR operations, latency
 - resource requirements: reservation table, issue unit(s)
- Formal specification in xADML [Bednarski'06]
(register sets etc. not considered here)

C. Kessler, IDA, Linköping universitet.

18



Instruction Scheduling

Overview

Christoph Kessler, IDA,
Linköpings universitet, 2014.

Instruction Scheduling



Generic Resource model: Reservation table

Optimize: time, space, energy

Local Scheduling
(f. Basic blocks / DAGs)

- Data dependences
→ Topological sorting
 - List Scheduling (diverse heuristics)
- Optimal Scheduling
(Exhaustive search, DP, B&B, CLP, ILP)


Global Scheduling

- Code motion, Branch delay slot filling
- Trace scheduling, Region scheduling, ...
- Cyclic scheduling for loops (Software pipelining)

There exist **retargetable schedulers**
and **scheduler generators**, e.g. for GCC since 2003

C. Kessler, IDA, Linköpings universitet.

20



Local Instruction Scheduling

Christoph Kessler, IDA,
Linköpings universitet, 2014.

Optimization problems in local scheduling



MRIS – minimum register need instruction scheduling

- + Spilling (store/reload) takes additional time
- + Power consumption in embedded procs. increases with # mem. accesses
- + Superscalar processors with shadow registers and register renaming
→ compiler-generated spill code cannot be eliminated at run time
- NP-complete [Sethi'75]

MTIS – minimum time instruction scheduling

- + hiding pipeline delays
- + exploiting instruction-level parallelism (for superscalar/VLIW)
- NP-complete [Garey/Johnson'79, Gross'83, Lawler et al.'87]

RCMTIS – register-constrained minimum time instruction scheduling

SMRTIS – simultaneous minimization of space and time

C. Kessler, IDA, Linköpings universitet.

22

MRIS: Space-optimal scheduling



(a) based on postorder traversal of the DAG

Special case: **tree**: space-opt. schedule in linear time [Sethi, Ullman '70]

Special case: **vector tree** (node size attribute): space-opt. $O(n \log n)$ [Raubert'90]

Special case: **series-par. DAG**: space-opt. schedule in pol. time [Güttler '81]

General DAG, **contiguous schedules** ($\leq 2^n$)

- Random dfs [K., Paul, Rauber '91]
- Enumeration with DC strategy [K., Rauber '93/95]



(b) based on topological sorting of the DAG → **general schedules** ($\leq n!$)

- space-optimal (enumeration + dynamic programming) [K. '96]

(c) based on finding instruction lineages in the DAG

- heuristic method by [Govindarajan et al. '00]

→ see separate lecture on MRIS

C. Kessler, IDA, Linköpings universitet.


23

Example: Topological Sorting (0)



- Not yet considered
- Data ready (zero-indegree set)
- Already scheduled, still alive
- Already scheduled, no longer referenced

Given:
Data flow graph of a basic block
(a directed acyclic graph, DAG)




C. Kessler, IDA, Linköpings universitet.

24

Example: Topological Sorting (1)

- Not yet considered
- Data ready (zero-indegree set)
- Already scheduled, still alive
- Already scheduled, no longer referenced




C. Kessler, IDA, Linköpings universitet.

25

Example: Topological Sorting (2)

- Not yet considered
- Data ready (zero-indegree set)
- Already scheduled, still alive
- Already scheduled, no longer referenced




C. Kessler, IDA, Linköpings universitet.

26

Example: Topological Sorting (3)

- Not yet considered
- Data ready (zero-indegree set)
- Already scheduled, still alive
- Already scheduled, no longer referenced




C. Kessler, IDA, Linköpings universitet.

27

Example: Topological Sorting (4)


- Not yet considered
- Data ready (zero-indegree set)
- Already scheduled, still alive
- Already scheduled, no longer referenced



C. Kessler, IDA, Linköpings universitet.

28

List scheduling = Topological sorting



```

top_sort( Set z, int[] INDEG, int t )
if z ≠ ∅ // (t ≤ n)
    select arbitrary node v ∈ z;
    // implicitly remove all edges (v,u) ∀u:
    INDEG'(u) = { INDEG(u) - 1 where ∃(v,u)
                  INDEG(u) elsewhere
    // update zero-indegree set:
    z' ← z - {v} ∪ {new leaves}
      = {u : INDEG(u) = 0}
    S[t] ← v;
    top_sort( z', INDEG', t + 1 );
else output S[1 : n] fi



Call top_sort( z₀, INDEG₀, 1 )
produces a schedule in S[1 : n]
    
```

C. Kessler, IDA, Linköpings universitet.

29

Topological Sorting and Scheduling

- Construct schedule incrementally in topological (= causal) order
 - "Appending" instructions to partial code sequence: close up in target schedule reservation table (as in "Tetris")
 - Idea: Find optimal target-schedule by enumerating all topological sortings ...
 - Beware of scheduling anomalies with complex reservation tables!




C. Kessler, IDA, Linköpings universitet.

30

Greedy List Scheduling for VLIW (1)

A **greedy** heuristic for list scheduling
fills in one step as many slots in a VLIW word as possible
with ready instructions of the zeroindegree set.




C. Kessler, IDA, Linköpings universitet.

31

Greedy List Scheduling for VLIW (2)

A **greedy** heuristic for list scheduling
fills in one step as many slots in a VLIW word as possible
with ready instructions of the zeroindegree set.



Not optimal!

(optimal) greedy schedule:

t	IntFit-Unit	IntMem-Unit
1	FitOp	FitLd
2	IntOp	IntLd

(non-optimal) greedy schedule:

t	IntFit-Unit	IntMem-Unit
1	FitOp	IntOp
2	—	IntLd
3	—	FitLd

C. Kessler, IDA, Linköpings universitet.

32

Local Scheduling Heuristics

List Scheduling Heuristics

- **Deepest Level First** (a.k.a. highest level first etc.)
 - ▶ Select, among ready instructions, one with longest accumulated latency on a path towards any dependence sink (root node)
 - ▶ Forward vs Backward scheduling

Critical Path Scheduling

- Detect a *critical path* (longest accumulated latency) in the DAG, schedule its nodes → partial schedule, and remove them from the DAG.
- Repeat until DAG is empty, splicing in new nodes between scheduled ones as appropriate, or inserting fresh time slots where needed

C. Kessler, IDA, Linköpings universitet.

33

DF00100 Advanced Compiler Construction

TDDC86 Compiler Optimizations and Code Generation

Global Instruction Scheduling

Christoph Kessler, IDA,
Linköpings universitet, 2014.

Scheduling Branch Instructions

Delayed Branch

- Effect of conditional branch on program counter is delayed
- 1 or more instructions after a branch instruction are always executed, independent of the condition outcome
 - ▶ SPARC, HP PA-RISC: 1 delay slot
 - ▶ TI 'C62x: 5 delay slots

- Scheduling: Fill delay slots with useful instruction if there is one, otherwise with NOP

Heuristic for finding candidate instructions:

1. Instructions from same basic block that are not control dependent on the branch and that the condition is not data dependent of
 2. Instructions from most likely branch target basic block for speculative execution
- See e.g. [Muchnick Ch. 17.1.1] for further details

C. Kessler, IDA, Linköpings universitet.

35

Remark: Hardware Support for Branch Delay Slot Filling

Nullification feature (SPARC) for conditional branch delay slots
in one branch direction

- activated by setting a bit in the instruction opcode
- delay slot instruction executed only if a conditional branch is **taken**
treated as NOP if not taken.
- not applicable to "branch always": fixed delay slot

C. Kessler, IDA, Linköpings universitet.

36

Trace Scheduling

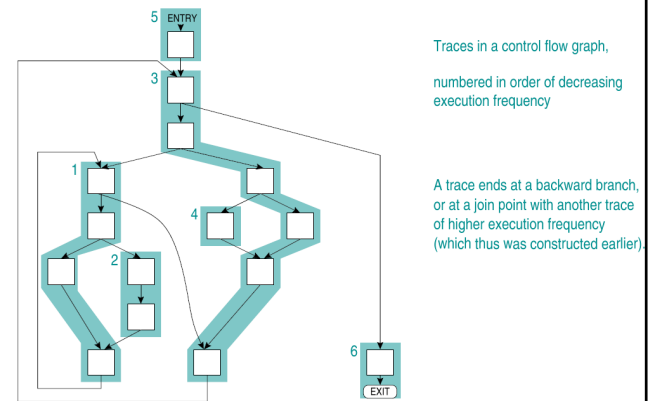
developed for VLIW architectures [Fisher'81] [Ellis'85]

- idea: enlarge the scope of local scheduling to **traces**
trace = acyclic path of basic blocks in the CFG
 track execution frequencies for BB's/traces (e.g., profiling)
- idea: make the most frequent trace fast:
 - + virtually merge BB's in the most frequent trace
 schedule trace as one BB, e.g. by greedy VLIW list scheduling
 - + insert compensation code in less frequent side traces for correctness
 → accept slowdown for side traces
 → program length may grow (worst case: exponentially)
 - + continue same procedure with next frequent trace

C. Kessler, IDA, Linköping universitet.

37

Trace Scheduling (2)




C. Kessler, IDA, Linköping universitet.

38

Trace Scheduling (3)

- Insertion of compensation code
 - Case:** When moving an instruction $i2$ to a predecessor block B in the trace T (e.g., to fill a branch delay slot)




C. Kessler, IDA, Linköping universitet.

39

Trace Scheduling (4)

- Insertion of compensation code
 - Case:** When moving an instruction $i1$ to a successor block of B in the trace T



C. Kessler, IDA, Linköping universitet.

40

Trace Scheduling (5)

- Summary of cases:**
 - Code reordering with insertion of compensation code
 - Hoisting an assignment
 - Interchange assignment and label
 - Moving assignments across conditional branches
 - Moving a branch
 - interchange branches

C. Kessler, IDA, Linköping universitet.

41

Region Scheduling

[Gupta/Soffa'90]

Idea: avoid idle cycles caused by regions with insufficient parallelism

Program region = one or several BB's that require the same control condition

Repeatedly apply a set of local code transformations:

- loop unrolling
- moving instructions from BB's with excessive parallelism into BB's with insufficient parallelism
- merging of regions to balance the degree of parallelism

Heuristic measure for average degree of parallelism in a region:

$\frac{\# \text{ instructions(region)}}{\text{length of critical path(region)}}$

C. Kessler, IDA, Linköping universitet.

42

Program regions for global scheduling



- **Trace** (see above)
 - A path of basic blocks
- **Superblock**
 - A trace with the restriction that there may be no branches into any of its basic blocks, except the first one
- **Treeregions = Extended Basic Blocks**
 - An out-tree of basic blocks – no branch into any of its basic blocks, except the first one
- **Hyperblock**
 - A single-entry, multiple exit region with internal control flow. As superblocks, but allow *hammocks* resolved by predication.
- All these regions are acyclic (but may be part of a cycle around)
- Traces and superblocks are "linear regions", while treeregions and hyperblocks are "nonlinear regions"

C. Kessler, IDA, Linköping universitet.

43

Summary + Outlook: Instruction Scheduling



- usually, optimize for time (other important metrics: space, energy)
 - see also lecture on energy-aware code generation
- local methods
 - postorder traversals, forward/backward list scheduling, optimal methods
 - see also lecture on space-optimal scheduling (MRIS)
- global methods
 - trace scheduling, percolation scheduling, region scheduling
 - see also lecture on software pipelining
- interferences with instruction selection, register allocation,
 - phase-ordering problems
 - see also lecture on integrated code generation
- interferences with data layout, exploit advanced addressing units, ...
 - see also lecture on code generation for DSPs

C. Kessler, IDA, Linköping universitet.

44

Further scheduling issues, not covered



- creating and scheduling predicated code
- speculation (with and without hardware support)
 - prefetching (load speculation), branch speculation, value speculation ...
- run-time scheduling, profile-driven scheduling
- automatic generation of instruction schedulers: finite state automata
 - [Proebsting/Fraser: *Detecting Pipeline Hazards Quickly*, POPL'94],
 - [Bala/Rubin MICRO-28, 1995]
 - e.g. the new GCC scheduler [Makarov, GCC Dev. Summit 2003]

C. Kessler, IDA, Linköping universitet.

45

Generation of Instruction Schedulers



- Given: Instruction set with
 - reservation table for each instruction
- Set of resource-valid schedules = regular language over the alphabet of instructions
- Scheduling instr. A after B leads to a certain **pipeline state** (functional unit reservations and pending latencies of recently issued instructions)
- Scheduling A in pipeline state q leads to new pipeline state q'
- → **Finite automaton** ("Müller automaton") of all possible pipeline states and (appending) scheduling transitions
 - Or finite transducer → gives also the time offset for next instruction
- Precompute possible states + transitions → Scheduling much faster (table lookup instead of interpreting reservation table composition)
- Reversed automaton to allow insertions at any location
- Automata become huge! But can be optimized.

C. Kessler, IDA, Linköping universitet.

46

Recommended Reading (global scheduling)



- J. Fisher. Trace scheduling: A technique for global microcode compaction. *IEEE Trans. Computers*, 30(7):478–490, 1981.
- Paolo Faraboschi, Joseph A. Fisher, Cliff Young: Instruction Scheduling for Instruction Level Parallel Processors *Proceedings of the IEEE*, vol. 89 no. 11, Nov. 2001
- Daniel Kästner, Sebastian Winkel: ILP-based Instruction Scheduling for IA-64. *Proc. ACM SIGPLAN LCTES-2001*, June 2001
- Sebastian Winkel. *Optimal Global Instruction Scheduling for the Itanium® Processor Architecture*. Ph.D. thesis. Saarland University, Saarbrücken, Germany, 2004. ISBN 3-937436-01-6

C. Kessler, IDA, Linköping universitet.

47

Recommended Reading (Generating Schedulers from Reservation Tables)



- T. Müller: Employing finite automata for resource scheduling. Proc. MICRO-26, 1993
- Proebsting, Fraser: Detecting pipeline structural hazards quickly. Proc. ACM POPL-1994
- Bala, Rubin: Efficient instruction scheduling using finite state automata. Proc. MICRO-28, 1995
- Eichenberger, Davidson: A reduced multi-pipeline machine description that preserves scheduling constraints. Proc. ACM PLDI-1996

C. Kessler, IDA, Linköping universitet.

48