



Code Generation for RISC and Instruction-Level Parallel Processors

RISC/ILP Processor Architecture Issues
Instruction Scheduling
Register Allocation
Phase Ordering Problems
Integrated Code Generation

Christoph Kessler, IDA,
Linköpings universitet, 2007.



1. RISC and Instruction-Level Parallel Target Architectures

Christoph Kessler, IDA,
Linköpings universitet, 2007.

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

TDDb44: Code Generation for RISC and ILP Processors

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. ARM9E, 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

TDDb44: Code Generation for RISC and ILP Processors

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 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	write
	sic1	sic2	0	1	0	1	2	3
Time	opnd	opnd						result
0								
1								
2								
3								

mul:

	read	read	stage	stage	stage	stage	stage	write
	sic1	sic2	0	1	0	1	2	3
Time	opnd	opnd						result
0								
1								
2								
3								
4								
5								


Reservation table
specifies required resource
occupations

[Davidson 1975]

t: mul ...
t+1: ...
t+2: add ...
... structural hazard at t=5

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



C. Kessler, IDA, Linköping universitet

VLIW (Very Long Instruction Word) architectures

- 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



C. Kessler, IDA, Linköping universitet

EPIC architectures

- Based on VLIW
- Compiler groups instructions to LIW's (bundles)
- Compiler takes care of resource and latency constraints
- Compiler marks sequences of independent instructions
- Dynamic scheduler assigns resources and reloads new bundles as required

C. Kessler, IDA, Linköping universitet

9

TDD844: Code Generation for RISC and ILP Processors


TDD829 Compilers and Interpreters (opt.)
TDD844 Compiler Construction

2. Instruction Scheduling

Christoph Kessler, IDA,
Linköping universitet, 2007.

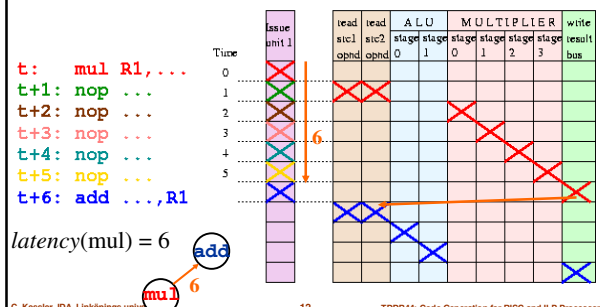
Instruction Scheduling (1)

- 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



C. Kessler, IDA, Linköping universitet

12

TDD844: Code Generation for RISC and ILP Processors

Instruction Scheduling

Generic Resource model

Reservation table

```
t: mul R1, ...
t+1: nop
t+2: nop
t+3: nop
t+4: nop
t+5: nop
t+6: add ..., R1
```

Local Scheduling

(f. Basic blocks / DAGs)

Data dependences


→ Topological sorting

- List Scheduling (diverse heuristics)

Global Scheduling


- Trace scheduling, Region scheduling, ...
- Cyclic scheduling (Software pipelining)

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




Example: Topological Sorting (0)

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




Example: Topological Sorting (1)

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




Example: Topological Sorting (2)

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




Example: Topological Sorting (3)

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



Example: Topological Sorting (4)


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



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!

[K. / Bednarski / Eriksson 2007]



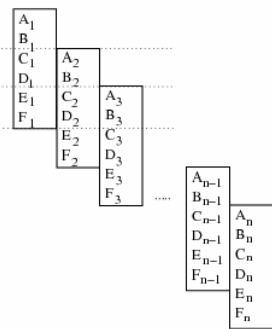
C. Kessler, IDA, Linköping universitet. 19 TDDB44: Code Generation for RISC and ILP Processors

Software Pipelining of Loops (1)



C. Kessler, IDA, Linköping universitet. 20 TDDB44: Code Generation for RISC and ILP Processors

Software Pipelining of Loops (2)




Unit1 Unit2 Unit3
 Prologue: A₁ B₁ C₁ A₂ D₁ B₂
 DO i=1, n-2
 Pattern: E_i C_{i+1} A_{i+2}
 F_i D_{i+1} B_{i+2}
 Epilogue: E_{n-1} C_n F_{n-1} D_n E_n F_n

→ More about Software Pipelining in TDDC86 Compiler Optimizations and Code Generation

C. Kessler, IDA, Linköping universitet. 21 TDDB44: Code Generation for RISC and ILP Processors

Software Pipelining of Loops (3) Modulo Scheduling



C. Kessler, IDA, Linköping universitet. 22 TDDB44: Code Generation for RISC and ILP Processors

TDDB29 Compilers and Interpreters (opt.)
 TDDB44 Compiler Construction

3. Register Allocation

Christoph Kessler, IDA, Linköping universitet, 2007.

Register Allocation

- Register Allocation:** Determines values (variables, temporaries, constants) to be kept when in registers
- Register Assignment:** Determine in which physical register such a value should reside.
- Essential for Load-Store Architectures
- Reduce memory traffic (→ memory / cache latency, energy)
- Limited resource
- Values that are alive simultaneously cannot be kept in the same register
- Strong interdependence with instruction scheduling
 - scheduling determines live ranges
 - spill code needs to be scheduled
- Local register allocation** (for a single basic block) can be done in linear time (see function getreg() above).
- Global register allocation** (with minimal spill code) is NP-complete. Can be modeled as a graph coloring problem [Ershov'62] [Cocke'71].

C. Kessler, IDA, Linköping universitet. 24 TDDB44: Code Generation for RISC and ILP Processors

Local Register Allocation

For variable v and basic block B_i :

$$netsave(v, i) = \#uses_i \cdot usesave + \#defs_i \cdot defsave$$

- $l \cdot ldcost$ ($l = 1$ if $Load(v)$ needed at beg. of B_i , 0 otherw.)
- $s \cdot stcost$ ($s = 1$ iff $Store(v)$ needed at end of B_i , 0 otherw.)

For loop L estimate benefit of keeping v in a register:

$$benefit(v, L) = 10^{depth(L)} \cdot \sum_{i \in blocks(L)} netsave(v, i)$$

with R registers available:

allocate the R objects v with greatest benefit in L .

moves may be necessary instead of $Load(v)$ / $Store(v)$

if v could reside in (different) registers in $Pred(B_i)$, B_i , $Succ(B_i)$

add worst-case terms $|Pred(v)| \cdot mvcost$, $|Succ(v)| \cdot mvcost$

C. Kessler, IDA, Linköping universitet. 25 TDD844: Code Generation for RISC and ILP Processors

Live range

(Here, variable = program variable or temporary)

- A variable is being **defined** at a program point if it is written (given a value) there.
- A variable is **used** at a program point if it is read (referenced in an expression) there.
- A variable is **alive** at a point if it is referenced there or at some following point that has not (may not have) been preceded by any definition.
- A variable is **reaching** a point if an (arbitrary) definition of it, or usage (because a variable can be used before it is defined) reaches the point.
- A variable's **live range** is the area of code (set of instructions) where the variable is both alive and reaching.
 - does not need to be consecutive in program text.


C. Kessler, IDA, Linköping universitet. 26 TDD844: Code Generation for RISC and ILP Processors

Register Allocation for Loops


Example:

```
x3 = 7
for i = 1 to 100 {
  x1 = x3 + 2
  x2 = x1 + x3
  x3 = x2 + x1
}
y = x3 + 42
```

Control flow graph



Live ranges (loop only):
cyclic intervals
e.g. for i: [0, 6), [6, 7)



All variables interfere with each other – need 4 regs?

At most 3 values alive at a time
→ 3 registers sufficient?

C. Kessler, IDA, Linköping universitet. 27 TDD844: Code Generation for RISC and ILP Processors

Register Allocation by Graph Coloring

- **Step 1:** Given a program with symbolic registers $s1, s2, \dots$
 - Determine live ranges of all variables


```
i = c+4; load 8(fp), s1 ! c
nop
addi s1, #4, s2
store s2, 4(fp) ! i
d = c-2; subi s1, #2, s3
store s3, 12(fp) ! d
c = c*i; muli s1, s2, s4
store s4, 8(fp) ! c
```



C. Kessler, IDA, Linköping universitet. 28 TDD844: Code Generation for RISC and ILP Processors

Register Allocation by Graph Coloring


- **Step 2:** Build the **Register Interference Graph**
 - Undirected edge connects two symbolic registers (s_i, s_j) if live ranges of s_i and s_j overlap in time
 - Reserved registers (e.g. fp) interfere with all s_i



C. Kessler, IDA, Linköping universitet. 29 TDD844: Code Generation for RISC and ILP Processors

- **Step 3:** Color the register interference graph with k colors, where $k = \#$ available registers.
 - If not possible: pick a victim s_i to spill, generate spill code (store after def., reload before use)
 - This may remove some interferences.

Rebuild the register interference graph + repeat Step 3...



This register interference graph cannot be colored with less than 4 colors, as it contains a 4-clique

C. Kessler, IDA, Linköping universitet. 30 TDD844: Code Generation for RISC and ILP Processors

Coloring a graph with k colors

- NP-complete for $k > 3$
- **Chromatic number** $\chi(G)$ = minimum number of colors to color a graph G
- $\chi(G) \geq c$ if the graph contains a c -clique
 - A **c -clique** is a completely connected subgraph of c nodes

■ Chaitin's heuristic (1981):

```

S ← { s1, s2, ... } // set of spill candidates
while ( S not empty )
  choose some s in S.
  if s has less than k neighbors in the graph
    then // there will be some color left for s:
      delete s (and incident edges) from the graph
    else
      modify the graph (spill, split, coalesce ... nodes)
      and restart.
// once we arrive here, the graph is empty:
color the nodes greedily in reverse order of removal.
    
```


C. Kessler, IDA, Linköping universitet.

31

TDD844: Code Generation for RISC and ILP Processors

Live range splitting

- Instead of spilling completely (reload before each use), it may be sufficient to split a live range at one position where register pressure is highest
- save, and reload once




C. Kessler, IDA, Linköping universitet.

32

TDD844: Code Generation for RISC and ILP Processors

Live range coalescing

- For a copy instruction $s_j \leftarrow s_i$
 - where s_i and s_j do not interfere
 - and s_i and s_j are not rewritten after the copy operation
- Merge s_i and s_j :
 - patch (rename) all occurrences of s_i to s_j
 - update the register interference graph
- and remove the copy operation.

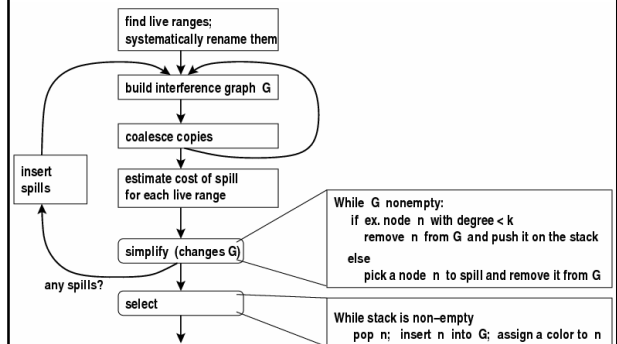


C. Kessler, IDA, Linköping universitet.

33

TDD844: Code Generation for RISC and ILP Processors

Chaitin's Register Allocator (1981)



C. Kessler, IDA, Linköping universitet.

34


TDD844: Code Generation for RISC and ILP Processors

TDD829 Compilers and Interpreters (opt.)
TDD844 Compiler Construction

4. Phase Ordering Problems and Integrated Code Generation

Christoph Kessler, IDA,
Linköping universitet, 2007.

Phase ordering problems



C. Kessler, IDA, Linköping universitet.


36

TDD844: Code Generation for RISC and ILP Processors


Phase ordering problems (1)

Instruction scheduling vs. register allocation

(a) **Scheduling first:**
determines Live-Ranges
→ Register need,
possibly spill-code to be
inserted afterwards



(b) **Register allocation first:**
Reuse of same register by different
values introduces "artificial"
data dependencies
→ constrains scheduler



C. Kessler, IDA, Linköping universitet

37

TDOB44: Code Generation for RISC and ILP Processors

Phase ordering problems (2)

Conflicts Instruction selection ↔ Scheduling / Reg. alloc.

■ **Selection first:** Cost attribute of a pattern covering rule is only a coarse estimate of the real cost (effect on e.g. overall time)

Real cost based on


- currently free functional units
- other instructions ready to execute simultaneously
- pending latencies of already issued but unfinished instructions

→ Integration with instruction scheduling desirable

■ **Mutations** with different resource requirements

- $a = 2 * b$ or $a = b < 1$ or $a = b + b$?

■ Different instructions with different register need




C. Kessler, IDA, Linköping universitet

38

TDOB44: Code Generation for RISC and ILP Processors

Clustered VLIW processor

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




C. Kessler, IDA, Linköping universitet

39

TDOB44: Code Generation for RISC and ILP Processors

More phase ordering problems



■ In parallel e.g.: load on A || load on B || move A→B

■ Mapping instructions → cluster

- should preferably know already beforehand about free move-slots in schedule...

■ Instruction scheduling

- must know mapping to generate moves where needed

■ Heuristic [Leupers'00]

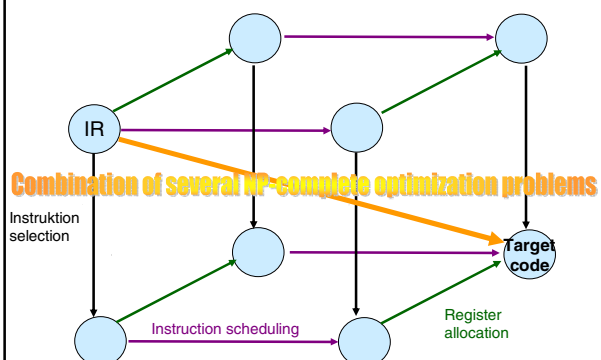
- iterative optimization by simulated annealing

C. Kessler, IDA, Linköping universitet

40

TDOB44: Code Generation for RISC and ILP Processors

5. Integrated Code generation



C. Kessler, IDA, Linköping universitet

41

TDOB44: Code Generation for RISC and ILP Processors

Integrated Code Generation Related Work

■ Heuristic Phase Interleaving

[Goodman/Hsu '88], [Leupers'00], ...

■ Combination of 2 phases

- Instruction selection and register allocation
 - ▶ DP for space optimization, e.g. [Aho/Johnsson '77]
 - ▶ DP for space- or time optimization, e.g. [Fraser et al. '92]
- Instruction scheduling and register allocation
 - ▶ ILP e.g. [Kästner'00]

■ Combination of 3 phases

- ILP [Wilson et al. '94]
 - ▶ only for simple, non-pipelined RISC processor

C. Kessler, IDA, Linköping universitet


42

TDOB44: Code Generation for RISC and ILP Processors

Results (2) – DP

- Clustered 8-issue VLIW processor TI C6201

Leupers' example




TI-C comp.	Schedule by Leupers' heuristic	Optimal schedule by OPTIMIST
LD *A4,B4	LD *A0,A8 MV A1,B8	LD *A0,A8 MV A1,B8
LD *A1,A8	LD *B8,B1 LD *A2,A9 MV A3,B10	LD *B8,B1 LD *A2,A9 MV A3,B10
LD *A3,A9	LD *B10,B3 LD *A4,A10 MV A5,B12	LD *B10,B3 LD *A4,A10 MV A5,B12
LD *A0,B0	LD *B12,B5 LD *A6,A11 MV A7,B14	LD *B12,B5 LD *A6,A11 MV A7,B14
LD *A2,B2	LD *B14,B7	LD *B14,B7 MV A8,B0
LD *A5,B5	MV A8,B0	MV A9,B2
LD *A7,A4	MV A9,B2	MV A10,B4
LD *A6,B6	MV A10,B4	MV A11,B6
NOP	MV A11,B6	NOP
MV A8,B1		
MV A9,B3		
MV A4,B7		

C. Kessler, IDA 12 cycles 9 cycles 9 cycles (15min)

Results (4) – Optimal integrated code generation, OPTIMIST – DP algorithm

- Single-issue vs. Single-cluster vs. Double-cluster Architecture



benchmark	BB	size	ARM9E-ARM mode t[s] #merged #ESnode	TI-C62x single-cluster t[s] #merged #ESnode	TI-C62x t[s] #merged #ESnode
analyzer	51	34	98.9 546250 62965	12623.2 2547538 630132	— —
analyzer	203	31	56.5 309673 49833	14.1 34559 29780	— —
bmmse	4	14	0.1 489 165	0.1 523 375	10.9 15215 704
bmmse	10	17	0.3 2558 582	0.8 5349 2413	6404.2 860455 204054
bmmse	11	14	0.1 768 234	0.2 825 562	56.8 38938 1555
codebk_srch	12	15	0.1 595 236	0.2 658 678	163.4 83220 3366
codebk_srch	16	21	0.7 6880 1512	8.3 47665 19613	— —
codebk_srch	20	23	2.2 20380 3723	45.0 240079 76005	— —
codebk_srch	24	42	178.5 783854 98580	— —	— —
fir_vslp	6	23	0.5 4479 1215	4.4 20483 10331	— —
fourinarow	6	29	0.2 862 561	1.2 4976 2214	10.0 23599 920
irr	4	21	0.4 4526 1016	5.8 42477 12036	— —
irr	7	38	2.1 8753 3259	38.0 87666 47233	— —

C. Kessler, IDA, Linköping universitet. 44 TDD84: Code Generation for RISC and ILP Processors


Our project: OPTIMIST

Retargetable integrated code generator

Open Source:

www.ida.liu.se/~chrke/optimist

architecture description



Available specifications:
 - TI C6201
 - ARM 9E
 - Motorola MC56K

C. Kessler, IDA, Linköping universitet. 45 TDD84: Code Generation for RISC and ILP Processors

Processor specification language xADML

Specify reservation tables by

```

<architecture omega="8">
  <registers> ... </registers>
  <residenceclasses> ... </residenceclasses>
  <funits> ... </funits>
  <patterns> ... </patterns>
  <instruction_set>
    <instruction id="ADDP4" op="4407">
      <target id="ADD .L1" op0="A" op1="A" op2="A" use_fu="L1"/>
      <target id="ADD .L2" op0="B" op1="B" op2="B" use_fu="L2"/>
    </instruction>
    ...
  </instruction_set>
  <transfer>
    <target id="MOVE" op0="A" op1="B">
      <use_fu="X2"/>
      <use_fu="L1"/>
    </target>
    ...
  </transfer>
</architecture>
  
```

Program cache/Program memory
 Program cache/Program memory

Project Literature (Selection)

- Christoph Kessler, Andrzej Bednarski:
Optimal integrated code generation for VLIW architectures.
Concurrency and Computation: Practice and Experience 18: 1353-1390, 2006.
- Andrzej Bednarski, Christoph Kessler:
Optimal Integrated VLIW Code Generation with Integer Linear Programming.
Proc. Euro-Par 2006 conference, Springer LNCS 4128, pp. 461-472, Aug. 2006.
- Andrzej Bednarski:
Optimal Integrated Code Generation for Digital Signal Processors.
PhD thesis, Linköping university - Institute of Technology, Linköping, Sweden, June 2006.
- Christoph Kessler, Andrzej Bednarski, Mattias Eriksson:
Classification and generation of schedules for VLIW processors.
Concurrency and Computation: Practice and Experience 19: 2369-2389, 2007.
- www.ida.liu.se/~chrke/optimist

C. Kessler, IDA, Linköping universitet. 47 TDD84: Code Generation for RISC and ILP Processors