

VU2 185.324

Compilation Techniques for VLIW Architectures

Dietmar Ebner ebner@complang.tuwien.ac.at
Florian Brandner brandner@complang.tuwien.ac.at

<http://complang.tuwien.ac.at/cd/vliw>

Last Lectures (1)

- Instruction scheduling
 - List scheduling
 - Classification
 - Regime, search strategy, region shapes
- Region formation
 - Mutual most likely strategy

Last Lectures (2)


- Region Types
 - Linear
 - Traces / Superblocks
 - Non linear
 - Hyperblocks / Treeregion / Trace 2
 - SEME vs MEME
- Details on Superblocks
 - Moderate compile time penalty
 - Major performance improvements
 - Sometimes severe code size increase

Last Lectures (3)

- Region enlargement
 - Loop unrolling
 - Loop peeling
 - Tail duplication
 - If-conversion (Hyperblocks)
- Dependence elimination
 - Renaming
 - Induction variable / accumulator expansion
 - Operation combining and migration

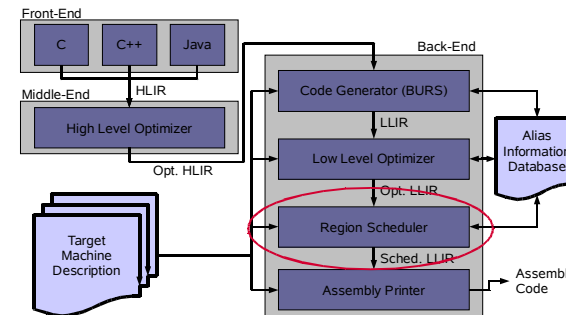
In Today's Lecture

- Cyclic Scheduling
 - Software Pipelining
 - Modulo Scheduling
- Predicated Execution
 - If-conversion



05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #5

Phases of an ILP Oriented Compiler



05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #6

Cyclic Scheduling

- Programs spend most of the time in loops
- Apply specialized scheduling techniques
 - Enlarge the loop body by unrolling
 - Overlap the execution of several loop iterations
 - Execute (parts of) different iterations in parallel
 - When applicable leads to good improvements
 - Not all loops can be handled

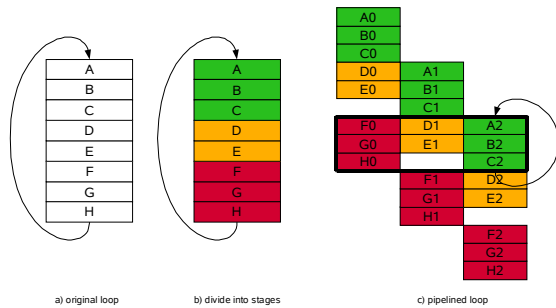
05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #7

Software Pipelining

- Family of cyclic schedulers
 - Divide the loop into stages
 - Execute stages of different iterations in parallel
 - Effectively pipelines the loop (similar to pipelining for computer architectures)
 - Dominating approach: Modulo Scheduling
- Optimize for throughput
 - The latency of a single iteration is irrelevant

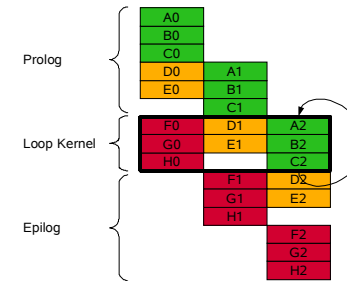
05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #8

Example: Software Pipelining (1)



05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #9

Example: Software Pipelining (2)



05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #10

Modulo Scheduling (1)

- Important SW Pipelining Technique
 - Explores the space of possible loop kernels
 - Initiation interval (II)
 - Constant interval between the start of successive kernel iterations
 - Lower bound of II (MinII)
 - Consider available hardware resources (ResII)
 - Data dependencies and recurrences (RecII)
 - Upper bound of II (MaxII)
 - Length of a linear schedule

05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #11

Modulo Scheduling (2)

- Search for an II starting from MinII
 - If a valid schedule could be found stop
 - Otherwise decide whether to
 - Backtrack - revert some scheduling decision
 - Increase the II
 - Abort if II is larger than MaxII
- High computational complexity
 - Backtracking and searching the II

05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #12

Scheduling Heuristics

- Modified version of list scheduling
 - Employ the modulo reservation table (MRT)
 - Similar to regular reservation tables
 - Ensure that a resource is not used at the same cycle, or at following cycles that modulo the II collide with it
 - Allows backtracking to revert scheduling decisions
 - Limited by a backtracking budget
 - May cause significant overhead


Example: Modulo Scheduling

```
int i, x ;
short a[];
for (i=0; i < 100; ++i)
    a[i+3] = a[i]*x + 7;
```

(a) Original C Code

	ALU0	ALU1	MEM0	MEM1	BR
0	---	---	mpyl	ldh	---
1	---	---	---	sth	---
2	add	br	---	---	tr(1)

(c) Modulo Reservation Table



(b) Data Dependence Graph

Cycle mod II	Cycle				
0	0	ldh			
1	1				
2	2				
0	3	mpyl	ldh		
1	4				
2	5	add			
0	6		mpyl	sth	
1	7	sth			
2	8	add	add		
0	9	cmp	sth	mpyl	ldh
1	10				
2	11	tr	add	add	
0	12	cmp	sth	mpyl	
1	13				
2	14	tr	add	add	
0	15		cmp	sth	
1	16				

(d) Modulo Schedule


Prolog / Epilog

- Partial copies of the loop kernel
 - May cause some code size increase
 - Strongly depending on the number of stages
- Several versions required in the presence of multiple loop exits and uncertain trip counts
- May harm runtime on small trip counts
 - Loop kernel is never reached (steady state)

Variable Names


- MS increases the register pressure
 - Preserve multiple copies of the same variable for different iterations
 - Life ranges exceeding the II
 - Spilling may disrupt the compact schedule
- Solutions
 - Modulo variable expansion (loop unrolling)
 - Copy operations
 - Hardware support (rotating register files)

Example: Variable Names (1)



- Minor change in the schedule
- The value of *ldh* exceeds the II
 - *mpyl* still requires the value of the last iteration

Example: Variable Names (2)



- **Solution**
 - Shown here: A extra copy operation
 - Alternative: Modulo Variable Expansion


Modulo Variable Expansion

- Unrolling increases the II
 - It is always possible to find a suitable unrolling factor to prevent copy operations
 - Assume v is the longest lifetime of a variable
 - The unrolling factor is given by $k = \lceil \frac{v}{II} \rceil$
- Larger kernel, prolog, and epilog
- Complicates handling of loop exits

Limitations of Modulo Scheduling

- Internal control flow
 - No internal control flow permitted
 - Applying if-conversion helps
 - Early loop exits are complex to handle
- Nested Loops
 - Recursively invoke the modulo scheduler starting with the innermost loop
- Low trip counts
 - Possibly never reach the steady state
 - Two versions unrolling & pipelined loop (code bloat)

Software Pipelining on a MIPSPro



- SPECfp92 on a MIPSPro (R8000, 75Mhz)
- SPECmarks with pipelining enabled/disabled

Side note: Papers on software pipelining typically do not show complete benchmarks, but concentrate on single loops only.

Software Pipelining on a VLIW

Program	L_{loop}	L_{loop}	Δv	P_{no}	P_{swp}
adpcm	2800	1596	1.75	2.1	3.8
bitstring	1000	798	1.25	2.2	2.8
blowfish	4000	2092	1.91	1.5	2.9
chain	800	292	2.74	1.9	5.0
chain2	900	295	3.05	0.8	5.0
crc32	1000	403	2.48	1.5	3.0
ctfir	1200	1200	1.00	3.3	3.3
dijkstra	2100	1194	1.76	1.8	3.2
films	1000	1000	1.00	3.6	3.6
imatt1x3	2900	799	3.13	1.0	3.1
imattred	4000	1985	2.02	1.0	3.1
FFT	5400	2386	2.26	1.2	2.8
gsm	2000	1692	1.18	2.9	3.4
isqrt	1100	796	1.38	2.6	3.6
patricia	1800	701	2.58	1.5	3.4
p.p.p	1400	698	2.01	1.7	2.7
shp	800	800	1.00	2.4	2.4
\emptyset			1.81		3.3

- TriMedia TM1000
 - 5 issue VLIW
 - 100 Mhz
- Benchmarks
 - DSPStone
 - MiBench

Side note: Papers on software pipelining typically do not show complete benchmarks, but concentrate on single loops only.

Source: D. Kästner, Absint


Predicated Execution

- Predicated operations
 - Conditionally nullify the result of the operation
 - Partial vs. full predication
 - All operations can be predicated
 - Conditional move or select operations
- Enable elimination of branches
 - Increase available ILP

If-Conversion

- Restructure the CFG
 - Augment blocks with predicates
 - Merge blocks
 - Eliminate branches
- Additional benefits
 - Increase scope for schedulers (Hyperblocks)
 - Enable software pipelining (internal control flow)
 - Other optimizations may benefit as well

Example: If-Conversion



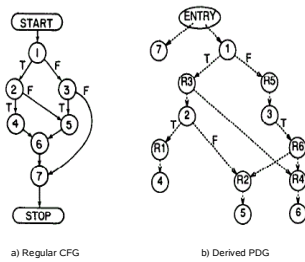
05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #25

If-Conversion

- Simple approaches
 - Select a candidate basic block following a heuristic
 - Often following simple patterns (e.g., if-then-else, etc.)
 - Calculate a predicate
 - Predicate all instructions of the block
- More Sophisticated
 - Use the Program Dependence Graph (PDG)
 - Try to place predicate definitions optimally
 - Try to derive a minimal set of predicates

05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #26

Program Dependence Graph



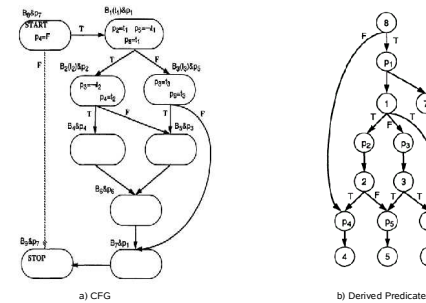
a) Regular CFG

b) Derived PDG

Source: J. Ferrante, HP

05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #27

If-conversion using RK




a) CFG

b) Derived Predicates

* R and K are functions on the control dependence graph, introduced by Park et al. Source: J. Park, HP

05/30/08 Ebner, Brandner | Compilation Techniques for VLIWs | SS08 Slide #28


If-conversion on Itanium (1)



* SPECINT2000, running on a near-production silicon

Source: Y. Choi, Intel


If-conversion on Itanium (2)



* SPECINT2000, running on a near-production silicon

Source: Y. Choi, Intel

If-conversion on Itanium (3)



* SPECINT2000, running on a near-production silicon

Source: Y. Choi, Intel

Limitation of If-Conversion

- Relies on hardware support
 - Non predictable instructions cause problems
- Aggressive application
 - Causes slow-down, increase of code size, and power consumption
 - Estimating profitability is hard
 - Reverse-if-conversion
 - Undo if-conversion if unprofitable

Outlook

- Second Assignments
 - Presentation of your results
 - Minimal execution time of the mpeg2decoder
 - Best scheduler implementation
 - The winner gets a small price