

# Compiler Optimisation

## 6 – Instruction Scheduling

Hugh Leather  
IF 1.18a  
hleather@inf.ed.ac.uk

Institute for Computing Systems Architecture  
School of Informatics  
University of Edinburgh

2019

# Introduction

This lecture:

- Scheduling to hide latency and exploit ILP
- Dependence graph
- Local list Scheduling + priorities
- Forward versus backward scheduling
- Software pipelining of loops

## Latency, functional units, and ILP

- Instructions take clock cycles to execute (*latency*)
- Modern machines issue several operations per cycle
- Cannot use results until ready, can do something else
- Execution time is *order-dependent*
- Latencies not always constant (cache, early exit, etc)

Operation	Cycles
load, store	3
load $\notin$ cache	100s
loadI, add, shift	1
mult	2
div	40
branch	0 – 8

# Machine types

- In order
  - Deep pipelining allows multiple instructions
- Superscalar
  - Multiple functional units, can issue  $> 1$  instruction
- Out of order
  - Large window of instructions can be reordered dynamically
- VLIW
  - Compiler statically allocates to FUs

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Simple schedule<sup>1</sup>

`a := 2*a*b*c`

Cycle	Operations	Operands waiting
loadAI	$r_{arp}, @a \Rightarrow r_1$	
add	$r_1, r_1 \Rightarrow r_1$	
loadAI	$r_{arp}, @b \Rightarrow r_2$	
mult	$r_1, r_2 \Rightarrow r_1$	
loadAI	$r_{arp}, @c \Rightarrow r_2$	
mult	$r_1, r_2 \Rightarrow r_1$	
storeAI	$r_1 \Rightarrow r_{arp}, @a$	
Done		

<sup>1</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Simple schedule<sup>1</sup>

`a := 2*a*b*c`

Cycle	Operations	Operands waiting
1	loadAI $r_{arp}, @a \Rightarrow r_1$	$r_1$
2		$r_1$
3		$r_1$
	add $r_1, r_1 \Rightarrow r_1$	
	loadAI $r_{arp}, @b \Rightarrow r_2$	
	mult $r_1, r_2 \Rightarrow r_1$	
	loadAI $r_{arp}, @c \Rightarrow r_2$	
	mult $r_1, r_2 \Rightarrow r_1$	
	storeAI $r_1 \Rightarrow r_{arp}, @a$	
	Done	

<sup>1</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Simple schedule<sup>1</sup>

$a := 2*a*b*c$

Cycle	Operations	Operands waiting
1	loadAI $r_{arp}, @a \Rightarrow r_1$	$r_1$
2		$r_1$
3		$r_1$
4	add $r_1, r_1 \Rightarrow r_1$	$r_1$
	loadAI $r_{arp}, @b \Rightarrow r_2$	
	mult $r_1, r_2 \Rightarrow r_1$	
	loadAI $r_{arp}, @c \Rightarrow r_2$	
	mult $r_1, r_2 \Rightarrow r_1$	
	storeAI $r_1 \Rightarrow r_{arp}, @a$	
	Done	

<sup>1</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Simple schedule<sup>1</sup>

$a := 2*a*b*c$

Cycle		Operations		Operands waiting
1	loadAI	$r_{arp}, @a$	$\Rightarrow r_1$	$r_1$
2				$r_1$
3				$r_1$
4	add	$r_1, r_1$	$\Rightarrow r_1$	$r_1$
5	loadAI	$r_{arp}, @b$	$\Rightarrow r_2$	$r_2$
6				$r_2$
7				$r_2$
	mult	$r_1, r_2$	$\Rightarrow r_1$	
	loadAI	$r_{arp}, @c$	$\Rightarrow r_2$	
	mult	$r_1, r_2$	$\Rightarrow r_1$	
	storeAI	$r_1$	$\Rightarrow r_{arp}, @a$	
	Done			

<sup>1</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Simple schedule<sup>1</sup>

`a := 2*a*b*c`

Cycle		Operations	Operands waiting
1	loadAI	$r_{arp}, @a \Rightarrow r_1$	$r_1$
2			$r_1$
3			$r_1$
4	add	$r_1, r_1 \Rightarrow r_1$	$r_1$
5	loadAI	$r_{arp}, @b \Rightarrow r_2$	$r_2$
6			$r_2$
7			$r_2$
8	mult	$r_1, r_2 \Rightarrow r_1$	$r_1$
9		<b>Next op does not use <math>r_1</math></b>	$r_1$
	loadAI	$r_{arp}, @c \Rightarrow r_2$	
	mult	$r_1, r_2 \Rightarrow r_1$	
	storeAI	$r_1 \Rightarrow r_{arp}, @a$	
	Done		

<sup>1</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Simple schedule<sup>1</sup>

$a := 2*a*b*c$

Cycle	Operations	Operands waiting
1	loadAI $r_{arp}, @a \Rightarrow r_1$	$r_1$
2		$r_1$
3		$r_1$
4	add $r_1, r_1 \Rightarrow r_1$	$r_1$
5	loadAI $r_{arp}, @b \Rightarrow r_2$	$r_2$
6		$r_2$
7		$r_2$
8	mult $r_1, r_2 \Rightarrow r_1$	$r_1$
9	loadAI $r_{arp}, @c \Rightarrow r_2$	$r_1, r_2$
10		$r_2$
11		$r_2$
	mult $r_1, r_2 \Rightarrow r_1$	
	storeAI $r_1 \Rightarrow r_{arp}, @a$	
	Done	

<sup>1</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Simple schedule<sup>1</sup>

$a := 2*a*b*c$

Cycle		Operations		Operands waiting
1	loadAI	$r_{arp}, @a$	$\Rightarrow r_1$	$r_1$
2				$r_1$
3				$r_1$
4	add	$r_1, r_1$	$\Rightarrow r_1$	$r_1$
5	loadAI	$r_{arp}, @b$	$\Rightarrow r_2$	$r_2$
6				$r_2$
7				$r_2$
8	mult	$r_1, r_2$	$\Rightarrow r_1$	$r_1$
9	loadAI	$r_{arp}, @c$	$\Rightarrow r_2$	$r_1, r_2$
10				$r_2$
11				$r_2$
12	mult	$r_1, r_2$	$\Rightarrow r_1$	$r_1$
13	storeAI	$r_1$	$\Rightarrow r_{arp}, @a$	$r_1$
	Done			

<sup>1</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Simple schedule<sup>1</sup>

$a := 2*a*b*c$

Cycle		Operations		Operands waiting
1	loadAI	$r_{arp}, @a$	$\Rightarrow r_1$	$r_1$
2				$r_1$
3				$r_1$
4	add	$r_1, r_1$	$\Rightarrow r_1$	$r_1$
5	loadAI	$r_{arp}, @b$	$\Rightarrow r_2$	$r_2$
6				$r_2$
7				$r_2$
8	mult	$r_1, r_2$	$\Rightarrow r_1$	$r_1$
9	loadAI	$r_{arp}, @c$	$\Rightarrow r_2$	$r_1, r_2$
10				$r_2$
11				$r_2$
12	mult	$r_1, r_2$	$\Rightarrow r_1$	$r_1$
13				$r_1$
14	storeAI	$r_1$	$\Rightarrow r_{arp}, @a$	store to complete
15				store to complete
16				<b>store to complete</b>
	Done			

<sup>1</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Schedule loads early<sup>2</sup>

`a := 2*a*b*c`

Cycle	Operations	Operands waiting
loadAI	<code>r<sub>arp</sub>, @a</code> ⇒ <code>r<sub>1</sub></code>	
loadAI	<code>r<sub>arp</sub>, @b</code> ⇒ <code>r<sub>2</sub></code>	
loadAI	<code>r<sub>arp</sub>, @c</code> ⇒ <code>r<sub>3</sub></code>	
add	<code>r<sub>1</sub>, r<sub>1</sub></code> ⇒ <code>r<sub>1</sub></code>	
mult	<code>r<sub>1</sub>, r<sub>2</sub></code> ⇒ <code>r<sub>1</sub></code>	
mult	<code>r<sub>1</sub>, r<sub>2</sub></code> ⇒ <code>r<sub>1</sub></code>	
storeAI	<code>r<sub>1</sub></code> ⇒ <code>r<sub>arp</sub>, @a</code>	
Done		

<sup>2</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Schedule loads early<sup>2</sup>

`a := 2*a*b*c`

Cycle	Operations	Operands waiting
1	loadAI $r_{arp}, @a \Rightarrow r_1$	$r_1$
	loadAI $r_{arp}, @b \Rightarrow r_2$	
	loadAI $r_{arp}, @c \Rightarrow r_3$	
	add $r_1, r_1 \Rightarrow r_1$	
	mult $r_1, r_2 \Rightarrow r_1$	
	mult $r_1, r_3 \Rightarrow r_1$	
	storeAI $r_1 \Rightarrow r_{arp}, @a$	
	Done	

<sup>2</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Schedule loads early<sup>2</sup>

`a := 2*a*b*c`

Cycle	Operations	Operands waiting
1	loadAI $r_{arp}, @a \Rightarrow r_1$	$r_1$
2	loadAI $r_{arp}, @b \Rightarrow r_2$	$r_1, r_2$
	loadAI $r_{arp}, @c \Rightarrow r_3$	
	add $r_1, r_1 \Rightarrow r_1$	
	mult $r_1, r_2 \Rightarrow r_1$	
	mult $r_1, r_3 \Rightarrow r_1$	
	storeAI $r_1 \Rightarrow r_{arp}, @a$	
	Done	

<sup>2</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Schedule loads early<sup>2</sup>

`a := 2*a*b*c`

Cycle	Operations	Operands waiting
1	loadAI $r_{arp}, @a \Rightarrow r_1$	$r_1$
2	loadAI $r_{arp}, @b \Rightarrow r_2$	$r_1, r_2$
3	loadAI $r_{arp}, @c \Rightarrow r_3$	$r_1, r_2, r_3$
	add $r_1, r_1 \Rightarrow r_1$	
	mult $r_1, r_2 \Rightarrow r_1$	
	mult $r_1, r_3 \Rightarrow r_1$	
	storeAI $r_1 \Rightarrow r_{arp}, @a$	
	Done	

<sup>2</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Schedule loads early<sup>2</sup>

`a := 2*a*b*c`

Cycle		Operations	Operands waiting
1	loadAI	$r_{arp}, @a \Rightarrow r_1$	$r_1$
2	loadAI	$r_{arp}, @b \Rightarrow r_2$	$r_1, r_2$
3	loadAI	$r_{arp}, @c \Rightarrow r_3$	$r_1, r_2, r_3$
4	add	$r_1, r_1 \Rightarrow r_1$	$r_1, r_2, r_3$
	mult	$r_1, r_2 \Rightarrow r_1$	
	mult	$r_1, r_3 \Rightarrow r_1$	
	storeAI	$r_1 \Rightarrow r_{arp}, @a$	
	Done		

<sup>2</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Schedule loads early<sup>2</sup>

`a := 2*a*b*c`

Cycle		Operations		Operands waiting
1	loadAI	$r_{arp}, @a \Rightarrow r_1$		$r_1$
2	loadAI	$r_{arp}, @b \Rightarrow r_2$		$r_1, r_2$
3	loadAI	$r_{arp}, @c \Rightarrow r_3$		$r_1, r_2, r_3$
4	add	$r_1, r_1 \Rightarrow r_1$		$r_1, r_2, r_3$
5	mult	$r_1, r_2 \Rightarrow r_1$		$r_1, r_3$
6				$r_1$
	mult	$r_1, r_3 \Rightarrow r_1$		
	storeAI	$r_1 \Rightarrow r_{arp}, @a$		
	Done			

<sup>2</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Schedule loads early<sup>2</sup>

`a := 2*a*b*c`

Cycle		Operations		Operands waiting
1	loadAI	$r_{arp}, @a$	$\Rightarrow r_1$	$r_1$
2	loadAI	$r_{arp}, @b$	$\Rightarrow r_2$	$r_1, r_2$
3	loadAI	$r_{arp}, @c$	$\Rightarrow r_3$	$r_1, r_2, r_3$
4	add	$r_1, r_1$	$\Rightarrow r_1$	$r_1, r_2, r_3$
5	mult	$r_1, r_2$	$\Rightarrow r_1$	$r_1, r_3$
6				$r_1$
7	mult	$r_1, r_3$	$\Rightarrow r_1$	$r_1$
8				$r_1$
	storeAI	$r_1$	$\Rightarrow r_{arp}, @a$	
	Done			

<sup>2</sup>loads/stores 3 cycles, mults 2, adds 1

# Effect of scheduling

Superscalar, 1 FU: New op each cycle if operands ready

Schedule loads early<sup>2</sup>

a := 2\*a\*b\*c

Cycle		Operations		Operands waiting
1	loadAI	$r_{arp}, @a \Rightarrow r_1$		$r_1$
2	loadAI	$r_{arp}, @b \Rightarrow r_2$		$r_1, r_2$
3	loadAI	$r_{arp}, @c \Rightarrow r_3$		$r_1, r_2, r_3$
4	add	$r_1, r_1 \Rightarrow r_1$		$r_1, r_2, r_3$
5	mult	$r_1, r_2 \Rightarrow r_1$		$r_1, r_3$
6				$r_1$
7	mult	$r_1, r_3 \Rightarrow r_1$		$r_1$
8				$r_1$
9	storeAI	$r_1 \Rightarrow r_{arp}, @a$		store to complete
10				store to complete
11				store to complete
	Done			

Uses one more register  
11 versus 16 cycles – 31% faster!

<sup>2</sup>loads/stores 3 cycles, mults 2, adds 1

## Scheduling problem

- Schedule maps operations to cycle;  $\forall a \in Ops, S(a) \in \mathbb{N}$
- Respect latency;  
 $\forall a, b \in Ops, a \text{ dependson } b \implies S(a) \geq S(b) + \lambda(b)$
- Respect function units; no more ops per type per cycle than FUs can handle
  
- Length of schedule,  $L(S) = \max_{a \in Ops} (S(a) + \lambda(a))$
- Schedule  $S$  is time-optimal if  $\forall S_1, L(S) \leq L(S_1)$
  
- **Problem:** Find a time-optimal schedule<sup>3</sup>
- Even local scheduling with many restrictions is *NP-complete*

---

<sup>3</sup>A schedule might also be optimal in terms of registers, power, or space

# List scheduling

Local greedy heuristic to produce schedules for single basic blocks


- 1 Rename to avoid anti-dependences
- 2 Build dependency graph
- 3 Prioritise operations
- 4 For each cycle
  - 1 Choose the highest priority ready operation & schedule it
  - 2 Update ready queue

# List scheduling

## Dependence/Precedence graph

- Schedule operation only when operands ready
- Build dependency graph of read-after-write (RAW) deps
  - Label with latency and FU requirements

Example:  $a = 2 * a * b * c$




# List scheduling

## Dependence/Precedence graph

- Schedule operation only when operands ready
- Build dependency graph of read-after-write (RAW) deps
  - Label with latency and FU requirements
- Anti-dependences (WAR) restrict movement

Example:  $a = 2 * a * b * c$




# List scheduling

## Dependence/Precedence graph

- Schedule operation only when operands ready
- Build dependency graph of read-after-write (RAW) deps
  - Label with latency and FU requirements
- Anti-dependences (WAR) restrict movement – renaming removes

Example:  $a = 2 * a * b * c$



# List scheduling

## List scheduling algorithm

```
Cycle  $\leftarrow$  1
Ready  $\leftarrow$  leaves of ( $D$ )
Active  $\leftarrow$   $\emptyset$ 
while(Ready  $\cup$  Active  $\neq$   $\emptyset$ )
     $\forall a \in$  Active where  $S(a) + \lambda(a) \leq$  Cycle
        Active  $\leftarrow$  Active - a
         $\forall b \in$  succs(a) where isready(b)
            Ready  $\leftarrow$  Ready  $\cup$  b
    if  $\exists a \in$  Ready and  $\forall b, a_{priority} \geq b_{priority}$ 
        Ready  $\leftarrow$  Ready - a
        S(op)  $\leftarrow$  Cycle
        Active  $\leftarrow$  Active  $\cup$  a
    Cycle  $\leftarrow$  Cycle + 1
```


# List scheduling

## Priorities

- Many different priorities used
  - Quality of schedules depends on good choice
- The longest latency path or critical path is a good priority
- Tie breakers
  - Last use of a value - decreases demand for register as moves it nearer def
  - Number of descendants - encourages scheduler to pursue multiple paths
  - Longer latency first - others can fit in shadow
  - Random

# List scheduling


Example: Schedule with priority by critical path length



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 1



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 2



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 3



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 4



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 5



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 6



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 7



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 8



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 9



# List scheduling

Example: Schedule with priority by critical path length


Cycle = 10



# List scheduling

Example: Schedule with priority by critical path length

Cycle = 11




# List scheduling

## Forward vs backward

- Can schedule from root to leaves (backward)
- May change schedule time
- List scheduling cheap, so try both, choose best

# List scheduling

## Forward vs backward



Opcode	loadI	lshift	add	addI	cmp	store
Latency	1	1	2	1	1	4

# List scheduling

## Forward vs backward

Forwards			
	Int	Int	Stores
1	loadI <sub>1</sub>	lshift	
2	loadI <sub>2</sub>	loadI <sub>3</sub>	
3	loadI <sub>4</sub>	add <sub>1</sub>	
4	add <sub>2</sub>	add <sub>3</sub>	
5	add <sub>4</sub>	addI	store <sub>1</sub>
6	cmp		store <sub>2</sub>
7			store <sub>3</sub>
8			store <sub>4</sub>
9			store <sub>5</sub>
10			
11			
12			
13	cbr		

Backwards			
	Int	Int	Stores
1	loadI <sub>1</sub>		
2	addI	lshift	
3	add <sub>4</sub>	loadI <sub>3</sub>	
4	add <sub>3</sub>	loadI <sub>2</sub>	store <sub>5</sub>
5	add <sub>2</sub>	loadI <sub>1</sub>	store <sub>4</sub>
6	add <sub>1</sub>		store <sub>3</sub>
7			store <sub>2</sub>
8			store <sub>1</sub>
9			
10			
11	cmp		
12	cbr		

# Scheduling Larger Regions


- Schedule extended basic blocks (EBBs)
  - Super block cloning
- Schedule traces
- Software pipelining

# Scheduling Larger Regions

## Extended basic blocks

### Extended basic block

EBB is maximal set of blocks such that  
Set has a single entry,  $B_i$   
Each block  $B_j$  other than  $B_i$  has  
exactly one predecessor




# Scheduling Larger Regions

## Extended basic blocks

### Extended basic block


EBB is maximal set of blocks such that  
Set has a single entry,  $B_i$   
Each block  $B_j$  other than  $B_i$  has  
exactly one predecessor



# Scheduling Larger Regions

## Extended basic blocks


- Schedule entire paths through EBBs
- Example has four EBB paths



# Scheduling Larger Regions

## Extended basic blocks


- Schedule entire paths through EBBs
- Example has four EBB paths



# Scheduling Larger Regions

## Extended basic blocks


- Schedule entire paths through EBBs
- Example has four EBB paths



# Scheduling Larger Regions

## Extended basic blocks


- Schedule entire paths through EBBs
- Example has four EBB paths



# Scheduling Larger Regions

## Extended basic blocks


- Schedule entire paths through EBBs
- Example has four EBB paths
- Having  $B_1$  in both causes conflicts
  - Moving an op **out of**  $B_1$  causes problems



# Scheduling Larger Regions

## Extended basic blocks


- Schedule entire paths through EBBs
- Example has four EBB paths
- Having  $B_1$  in both causes conflicts
  - Moving an op **out of**  $B_1$  causes problems
  - Must insert compensation code



# Scheduling Larger Regions

## Extended basic blocks


- Schedule entire paths through EBBs
- Example has four EBB paths
- Having  $B_1$  in both causes conflicts
- Moving an op **into**  $B_1$  causes problems



# Scheduling Larger Regions

## Superblock cloning


- Join points create context problems



# Scheduling Larger Regions

## Superblock cloning


- Join points create context problems
- Clone blocks to create more context



# Scheduling Larger Regions

## Superblock cloning


- Join points create context problems
- Clone blocks to create more context
- Merge any simple control flow



# Scheduling Larger Regions

## Superblock cloning


- Join points create context problems
- Clone blocks to create more context
- Merge any simple control flow
- Schedule EBBs



# Scheduling Larger Regions

## Trace scheduling


- Edge frequency from profile (not block frequency)



# Scheduling Larger Regions

## Trace scheduling


- Edge frequency from profile (not block frequency)
- Pick “hot” path
- Schedule with compensation code



# Scheduling Larger Regions

## Trace scheduling


- Edge frequency from profile (not block frequency)
- Pick “hot” path
- Schedule with compensation code
- Remove from CFG



# Scheduling Larger Regions

## Trace scheduling

- Edge frequency from profile (not block frequency)
- Pick “hot” path
- Schedule with compensation code
- Remove from CFG
- Repeat



# Loop scheduling

- Loop structures can dominate execution time
- Specialist technique software pipelining
- Allows application of list scheduling to loops
  
- Why not loop unrolling?


# Loop scheduling

- Loop structures can dominate execution time
- Specialist technique software pipelining
- Allows application of list scheduling to loops
  
- Why not loop unrolling?
- Allows loop effect to become arbitrarily small, but
- Code growth, cache pressure, register pressure

# Software pipelining


Consider simple loop to sum array

```
b = 0
for i = 0 to n
    b = b + A[i]
```



## Software pipelining


Schedule on 1 FU - 5 cycles



load 3 cycles, add 1 cycle, branch 1 cycle

# Software pipelining


Schedule on VLIW 3 FUs - 4 cycles



load 3 cycles, add 1 cycle, branch 1 cycle

# Software pipelining


A better steady state schedule exists



load 3 cycles, add 1 cycle, branch 1 cycle

# Software pipelining


Requires prologue and epilogue (may schedule others in epilogue)



load 3 cycles, add 1 cycle, branch 1 cycle

# Software pipelining


Respect dependences and latency – including loop carries



load 3 cycles, add 1 cycle, branch 1 cycle

# Software pipelining

## Complete code



load 3 cycles, add 1 cycle, branch 1 cycle

# Software pipelining

## Some definitions

### Initiation interval ( $ii$ )

Number of cycles between initiating loop iterations

- Original loop had  $ii$  of 5 cycles
- Final loop had  $ii$  of 2 cycles

### Recurrence

Loop-based computation whose value is used in later loop iteration

- Might be several iterations later
- Has dependency chain(s) on itself
- Recurrence latency is latency of dependency chain

# Software pipelining

## Algorithm

- Choose an initiation interval,  $ii$ 
  - Compute lower bounds on  $ii$
  - Shorter  $ii$  means faster overall execution
- Generate a loop body that takes  $ii$  cycles
  - Try to schedule into  $ii$  cycles, using modulo scheduler
  - If it fails, increase  $ii$  by one and try again
- Generate the needed prologue and epilogue code
  - For prologue, work backward from upward exposed uses in the scheduled loop body
  - For epilogue, work forward from downward exposed definitions in the scheduled loop body

# Software pipelining

## Initial initiation interval ( $ii$ )

Starting value for  $ii$  based on minimum resource and recurrence constraints

### Resource constraint

- $ii$  must be large enough to issue every operation
- Let  $N_u$  = number of FUs of type  $u$
- Let  $I_u$  = number of operations of type  $u$
- $\lceil I_u/N_u \rceil$  is lower bound on  $ii$  for type  $u$
- $\max_u(\lceil I_u/N_u \rceil)$  is lower bound on  $ii$

# Software pipelining

## Initial initiation interval ( $ii$ )

Starting value for  $ii$  based on minimum resource and recurrence constraints

### Recurrence constraint

- $ii$  cannot be smaller than longest recurrence latency
- Recurrence  $r$  is over  $k_r$  iterations with latency  $\lambda_r$
- $\lceil \lambda_r / k_u \rceil$  is lower bound on  $ii$  for type  $r$
- $\max_r(\lceil \lambda_r / k_u \rceil)$  is lower bound on  $ii$

# Software pipelining

Initial initiation interval ( $ii$ )

Starting value for  $ii$  based on minimum resource and recurrence constraints

$$\text{Start value} = \max(\max_u(\lceil I_u/N_u \rceil), \max_r(\lceil \lambda_r/k_r \rceil))$$

For simple loop

```
a = A[ i ]  
b = b + a  
i = i + 1  
if i < n goto  
end
```

## Resource constraint

	Memory	Integer	Branch
$I_u$	1	2	1
$N_u$	1	1	1
$\lceil I_u/N_u \rceil$	1	2	1

## Recurrence constraint




	b	i
$k_r$	1	1
$\lambda_r$	2	1
$\lceil I_u/N_u \rceil$	2	1

# Software pipelining

## Modulo scheduling

### Modulo scheduling

Schedule with cycle modulo initiation interval






# Software pipelining

## Modulo scheduling

### Modulo scheduling

Schedule with cycle modulo initiation interval






# Software pipelining

## Modulo scheduling

### Modulo scheduling

Schedule with cycle modulo initiation interval






# Software pipelining

## Modulo scheduling

### Modulo scheduling

Schedule with cycle modulo initiation interval






# Software pipelining

## Modulo scheduling

### Modulo scheduling

Schedule with cycle modulo initiation interval



# Software pipelining

## Current research

- Much research in different software pipelining techniques
- Difficult when there is general control flow in the loop
- Predication in IA64 for example really helps here
- Some recent work in exhaustive scheduling -i.e. solve the NP-complete problem for basic blocks

# Summary

- Scheduling to hide latency and exploit ILP
- Dependence graph - dependences between instructions + latency
- Local list Scheduling + priorities
- Forward versus backward scheduling
- Scheduling EBBs, superblock cloning, trace scheduling
- Software pipelining of loops

# PPar CDT Advert

## EPSRC Centre for Doctoral Training in Pervasive Parallelism

- 4-year programme:  
MSc by Research + PhD
- Research-focused:  
Work on your thesis topic  
from the start
- Collaboration between:
  - ▶ University of Edinburgh's  
School of Informatics
    - \* Ranked top in the UK by  
2014 REF
  - ▶ Edinburgh Parallel Computing  
Centre
    - \* UK's largest supercomputing  
centre
- Research topics in software,  
hardware, theory and  
application of:
  - ▶ Parallelism
  - ▶ Concurrency
  - ▶ Distribution
- Full funding available
- Industrial engagement  
programme includes  
internships at leading  
companies

The biggest revolution  
in the technological  
landscape for fifty years

Now accepting applications!  
Find out more and apply at:

[pervasiveparallelism.inf.ed.ac.uk](http://pervasiveparallelism.inf.ed.ac.uk)



THE UNIVERSITY OF EDINBURGH  
**informatics**

**EPSRC**

Engineering and Physical Sciences  
Research Council