EECS 583 – Class 14 Modulo Scheduling Reloaded

University of Michigan

October 31, 2011



Announcements + Reading Material

- Project proposal Due Friday Nov 4
 - » 1 email from each group: names, paragraph summarizing what you plan to do
- Today's class reading
 - "Code Generation Schema for Modulo Scheduled Loops", B.
 Rau, M. Schlansker, and P. Tirumalai, MICRO-25, Dec. 1992.
- Next reading Last class before research stuff!
 - » "Register Allocation and Spilling Via Graph Coloring," G. Chaitin, Proc. 1982 SIGPLAN Symposium on Compiler Construction, 1982.

Review: A Software Pipeline



Loop Prolog and Epilog



Only the kernel involves executing full width of operations

Prolog and epilog execute a subset (ramp-up and ramp-down)

Separate Code for Prolog and Epilog



Generate special code before the loop (preheader) to fill the pipe and special code after the loop to drain the pipe.

Peel off II-1 iterations for the prolog. Complete II-1 iterations in epilog

Removing Prolog/Epilog



Execute loop kernel on every iteration, but for prolog and epilog selectively disable the appropriate operations to fill/drain the pipeline

Kernel-only Code Using Rotating Predicates



Modulo Scheduling Architectural Support

- Loop requiring N iterations
 - » Will take N + (S 1) where S is the number of stages
- ✤ 2 special registers created
 - » LC: loop counter (holds N)
 - » ESC: epilog stage counter (holds S)
- Software pipeline branch operations
 - » Initialize LC = N, ESC = S in loop preheader
 - » All rotating predicates are cleared
 - » BRF.B.B.F
 - While LC > 0, decrement LC and RRB, P[0] = 1, branch to top of loop
 - This occurs for prolog and kernel
 - If LC = 0, then while ESC > 0, decrement RRB and write a 0 into P[0], and branch to the top of the loop
 - This occurs for the epilog

Execution History With LC/ESC

L	C =	3.	ESC =	3	/*	Remember	0	re	lative		*/
	\sim	-,		-	·		\sim			•	

Clear all rotating predicates

P[0] = 1

A if P[0]; B if P[1]; C if P[2]; D if P[3]; P[0] = BRF.B.B.F;

LC	ESC	P[0]	P[1]	P[2]	P[3]				
3	3	1	0	0	0	А			
2	3	1	1	0	0	А	В		
1	3	1	1	1	0	А	В	С	
0	3	1	1	1	1	А	В	С	D
0	2	0	1	1	1	-	В	С	D
0	1	0	0	1	1	-	-	С	D
0	0	0	0	0	1	-	-	-	D

4 iterations, 4 stages, II = 1, Note 4 + 4 - 1 iterations of kernel executed

Review: Modulo Scheduling Process

- Use list scheduling but we need a few twists
 - » II is predetermined starts at MII, then is incremented
 - » Cyclic dependences complicate matters
 - Estart/Priority/etc.
 - Consumer scheduled before producer is considered
 - There is a window where something can be scheduled!
 - » Guarantee the repeating pattern
- ✤ 2 constraints enforced on the schedule
 - » Each iteration begin exactly II cycles after the previous one
 - » Each time an operation is scheduled in 1 iteration, it is tentatively scheduled in subsequent iterations at intervals of II
 - MRT used for this

Review: ResMII Example

Concept: If there were no dependences between the operations, what is the the shortest possible schedule?

resources: 4 issue, 2 alu, 1 mem, 1 br latencies: add=1, mpy=3, ld = 2, st = 1, br = 1

1: r3 = load(r1) 2: r4 = r3 * 26 3: store (r2, r4) 4: r1 = r1 + 4 5: r2 = r2 + 4 6: p1 = cmpp (r1 < r9) 7: brct p1 Loop ALU: used by 2, 4, 5, 6 \rightarrow 4 ops / 2 units = 2 Mem: used by 1, 3 \rightarrow 2 ops / 1 unit = 2 Br: used by 7 \rightarrow 1 op / 1 unit = 1

ResMII = MAX(2,2,1) = 2

Review: RecMII Example



<delay, distance>

Review: Priority Function

Height-based priority worked well for acyclic scheduling, makes sense that it will work for loops as well



EffDelay(X,Y) = Delay(X,Y) - II*Distance(X,Y)

Calculating Height

- 1. Insert pseudo edges from all nodes to branch with latency = 0, distance = 0 (dotted edges)
- 2. Compute II, For this example assume II = 2
- 3. Height R(4) = 0
- 4. HeightR(3) = 0 H(4) + EffDelay(3,4) = 0 + 0 - 0*II = 0 H(2) + EffDelay(3,2) = 2 + 2 - 2*II = 0 MAX(0,0) = 0
- 5. HeightR(2) = 2 H(3) + EffDelay(2,3) = 0 + 2 - 0 * II = 2H(4) + EffDelay(2,4) = 0 + 0 - 0 * II = 0MAX(2,0) = 0

6. HeightR(1) = 5 H(2) + EffDelay(1,2) = 2 + 3 - 0 * II = 5 H(4) + EffDelay(1,4) = 0 + 0 - 0 * II = 0 MAX(5,0) = 5



The Scheduling Window

With cyclic scheduling, not all the predecessors may be scheduled, so a more flexible <u>earliest schedule time</u> is:

E(Y) = MAX
for all X = pred(Y)
$$MAX (0, SchedTime(X) + EffDelay(X,Y)),$$
otherwise

where EffDelay(X,Y) = Delay(X,Y) - II*Distance(X,Y)

Every II cycles a new loop iteration will be initialized, thus every II cycles the pattern will repeat. Thus, you only have to look in a window of size II, if the operation cannot be scheduled there, then it cannot be scheduled.

Latest schedule time(Y) = L(Y) = E(Y) + II - 1

Implementing Modulo Scheduling - Driver

- compute MII
- $\bullet II = MII$
- budget = BUDGET_RATIO * number of ops
- while (schedule is not found) do
 - » iterative_schedule(II, budget)
 - » II++

 Budget_ratio is a measure of the amount of backtracking that can be performed before giving up and trying a higher II

Modulo Scheduling – Iterative Scheduler

- iterative_schedule(II, budget)
 - » compute op priorities
 - » while (there are unscheduled ops and budget > 0) do
 - op = unscheduled op with the highest priority
 - min = early time for op (E(Y))
 - $\max = \min + II 1$
 - t = find_slot(op, min, max)
 - schedule op at time t
 - ◆ /* Backtracking phase undo previous scheduling decisions */
 - Unschedule all previously scheduled ops that conflict with op
 - budget--

Modulo Scheduling – Find_slot

- find_slot(op, min, max)
 - » /* Successively try each time in the range */
 - » for $(t = \min to max)$ do
 - if (op has no resource conflicts in MRT at t)
 - return t
 - » /* Op cannot be scheduled in its specified range */
 - » /* So schedule this op and displace all conflicting ops */
 - » if (op has never been scheduled or min > previous scheduled time of op)
 - return min
 - » else
 - return MIN(1 + prev scheduled time of op, max)

Modulo Scheduling Example

resources: 4 issue, 2 alu, 1 mem, 1 br latencies: add=1, mpy=3, ld = 2, st = 1, br = 1

for (j=0; j<100; j++) b[j] = a[j] * 26 Step1: Compute to loop into form that uses LC

LC = 99



resources: 4 issue, 2 alu, 1 mem, 1 br latencies: add=1, mpy=3, ld = 2, st = 1, br = 1





resources: 4 issue, 2 alu, 1 mem, 1 br latencies: add=1, mpy=3, ld = 2, st = 1, br = 1

Step3: Draw dependence graph Calculate MII





Step 4 – Calculate priorities (MAX height to pseudo stop node)

Iter1	Iter2
1: H = 5	1: H = 5
2: H = 3	2: H = 3
3: H = 0	3: $H = 0$
4: H = 0	4: H = 4
5: H = 0	5: H = 0
7: H = 0	7: H = 0

resources: 4 issue, 2 alu, 1 mem, 1 br latencies: add=1, mpy=3, ld = 2, st = 1, br = 1 Schedule brlc at time II - 1



Step6: Schedule the highest priority op

Op1: E = 0, L = 1 Place at time 0 (0 % 2)

Loop:



Step7: Schedule the highest priority op

Op4: E = 0, L = 1 Place at time 0 (0 % 2)

Loop:



Step8: Schedule the highest priority op

Op2: E = 2, L = 3 Place at time 2 (2 % 2)

Loop:



Step9: Schedule the highest priority op

Op3: E = 5, L = 6 Place at time 5 (5 % 2)

Loop:



Step10: Schedule the highest priority op

Op5: E = 0, L = 1 Place at time 1 (1 % 2)

Loop:



Step11: calculate ESC, SC = max unrolled sched length / ii unrolled sched time of branch = rolled sched time of br + (ii*esc)



Finishing touches - Sort ops, initialize ESC, insert BRF and staging predicate, initialize staging predicate outside loop

LC = 99ESC = 2 p1[0] = 1

Loop:

1: r3[-1] = load(r1[0]) if p1[0] 2: r4[-1] = r3[-1] * 26 if p1[1] 4: r1[-1] = r1[0] + 4 if p1[0] 3: store (r2[0], r4[-1]) if p1[2] 5: r2[-1] = r2[0] + 4 if p1[0] 7: brlc Loop if p1[2] Staging predicate, each successive stage increment the index of the staging predicate by 1, stage 1 gets px[0]

> Unrolled Schedule



Example – Dynamic Execution of the Code

LC = 99	time: ops executed
ESC = 2	0: 1, 4
p1[0] = 1	1:5
	2: 1,2,4
oop: 1: r3[-1] = load(r1[0]) if p1[0]	3: 5
2: r4[-1] = r3[-1] * 26 if p1[1]	4: 1,2,4
4: r1[-1] = r1[0] + 4 if p1[0]	5: 3,5,7
3: store (r2[0], r4[-1]) if $p1[2]$	6: 1,2,4
5: $r_2[-1] = r_2[0] + 4 \text{ II } p_1[0]$ 7: $brle L con if p_1[2]$	<u>7: 3,5,7</u>
7. one 200p ii p1[2]	
	98: 1,2,4
	99: 3,5,7
	100: 2
	101: 3,7
	102: -
	103 3,7

Homework Problem

latencies: add=1, mpy=3, ld = 2, st = 1, br = 1

for
$$(j=0; j<100; j++)$$

b[j] = a[j] * 26

LC = 99

1

Loop:	1: $r3 = load(r1)$
	2: r4 = r3 * 26
	3: store (r2, r4)
	4: $r1 = r1 + 4$
	5: $r^2 = r^2 + 4$
	7: brlc Loop
	_

How many resources of each type are required to achieve an II=1 schedule?

If the resources are non-pipelined, how many resources of each type are required to achieve II=1

Assuming pipelined resources, generate the II=1 modulo schedule.

What if We Don't Have Hardware Support?

No predicates

- » Predicates enable kernel-only code by selectively enabling/disabling operations to create prolog/epilog
- » Now must create explicit prolog/epilog code segments
- No rotating registers
 - » Register names not automatically changed each iteration
 - » Must unroll the body of the software pipeline, explicitly rename
 - Consider each register lifetime i in the loop
 - Kmin = min unroll factor = MAXi (ceiling((Endi Starti) / II))
 - Create Kmin static names to handle maximum register lifetime
 - » Apply modulo variable expansion

No Predicates



Without predicates, must create explicit prolog and epilogs, but no explicit renaming is needed as rotating registers take care of this

No Predicates and No Rotating Registers

Assume Kmin = 4 for this example

