

CSE 401/M501 – Compilers

Compiler Backend Survey

Hal Perkins

Autumn 2024

Administrivia (1)

- CSE 401 compiler codegen: due tomorrow night, 11:59 pm (usual group late days apply)
 - All previous 🐛🕷️🐜🦋🐞 must be fixed. After codegen eval, there is an overall evaluation of the complete project – what does the entire compiler look like at the end – and that is a major part of the final grade
 - Chance to fix/recover from problems in earlier parts
 - Check your work (does it compile after a clone and ant clean? does the generated code run on attu?)
- Project reports: description posted now
 - CSE 401 report must be pushed to repos by next Monday 11:59 pm – no late days

Administrivia (2)

- All: HW4 out now, due next Thur. 11:59 pm.
Usual individual late days apply
- CSE M 501 projects due end of next week:
 - Codegen+CSE M 501 extras due Sat. 12/7, 11:59 pm
 - Report (basics+extensions) due Sun. 12/8 11:59 pm
 - No late anything accepted after those deadlines
 - This is for M501 **only** – 401 project is a week earlier!
- What should we do about class this Wednesday (afternoon before thanksgiving)?

Administrivia (added Mon. 12/2)

- CSE401 project reports due tonight 11:59 pm – no late submissions accepted
- Hw4 due Thur. night 12/5; if you have late days you can use them
- Lectures this week: back end today; GC and wrapup Wed.; no class Fri.
- Sections this week: last-minute hw4 Q&A, instruction scheduling and register allocation, quarter wrapup, and whatever you and your instructors want to do 😊
- CSE M 501 compiler project due Saturday 11:59 pm
 - Report due Sunday night 11:59 pm
 - No late submissions on either of these (see m501 additions writeup)
- Please finish course evals while they are still available
- Please nominate great TAs for the Bander award.
 - Both for CSE 401/M501 and for other courses


Final exam (added Mon. 12/2)

- Next Tuesday 12/10 2:30-4:20 here
- Closed book like midterm, but ok to have two 5x8 hand-written note cards (can re-use midterm card or replace it)
 - Blank cards available end of class today & Wednesday & sections this week & review session
- Topic list posted on website now
- Review Q&A Monday 12/9 4:30, also here (G10)
- HW4 sample solutions will be available outside Perkins' office (CSE 548) by late Sunday (will announce on discussion board when ready)
 - Copies also available at the Monday review session.

Agenda

- Survey major pieces of a compiler back end
 - Instruction selection
 - Instruction scheduling
 - Register allocation
- And three particularly neat algorithms
 - Instruction selection by tree pattern matching
 - Instruction list scheduling
 - Register allocation by graph coloring

Compiler Organization



Big Picture

- Compiler consists of lots of fast stuff followed by hard problems
 - Scanner: $O(n)$
 - Parser: $O(n)$
 - Analysis & Optimization: $\sim O(n \log n)$
 - Instruction selection: fast or NP-Complete (or worse)
 - Instruction scheduling: NP-Complete
 - Register allocation: NP-Complete

IR for Code Generation

- Assume a (very) low-level IR
 - 3 address, register-register instructions plus load/store
$$r1 \leftarrow r2 \text{ op } r3$$
 - Could be tree structure or linear
 - Expose as much detail as possible
- Assume “enough” (i.e., ∞) registers
 - Invent new temporaries for intermediate results
 - Map to actual registers towards the end

Overview: Instruction Selection

- Map IR into assembly code
- Assume known storage layout and code shape
 - i.e., the optimization phases have already done their thing
- Combine low-level IR operations into machine instructions (take advantage of addressing modes, etc.)

Overview: Instruction Scheduling

- Reorder instructions to minimize execution time
 - hide latencies – processor function units, memory/cache stalls
 - Originally invented for supercomputers (60s)
 - Required to get reasonable (or correct!) code on early RISC architectures (~3-address code hardware ISAs)
 - Still important on most machines
 - Even non-RISC machines, e.g., x86 family
 - Even if processor reorders on the fly
- Good schedules help processor do a better job
- Assume fixed program at this point

Overview: Register Allocation

- Map values to actual registers
 - Previous phases change need for registers
- Add code to spill values to temporaries in memory and reload as needed, etc.
- Usually worth doing another instruction scheduling pass afterwards if spill code inserted

Conventional Wisdom

- We typically lose little by solving these independently
 - But not always, of course (iterating phases on x86-64 can help because of limited registers and use of memory operands)
- Instruction selection
 - Use some form of pattern matching
 - ∞ virtual registers – create as needed
- Instruction scheduling
 - Within a block, list scheduling is close to optimal
 - Across blocks: extended basic blocks or trace scheduling if list scheduling is not good enough
- Register allocation
 - Start with unlimited virtual registers and map to some subset of K real registers

Agenda

- Survey major pieces of a compiler back end
 - Instruction selection
 - Instruction scheduling
 - Register allocation
- And three particularly neat algorithms
 - Instruction selection by tree pattern matching
 - Instruction list scheduling
 - Register allocation by graph coloring

Instruction Selection

- Map IR into assembly code
- Assume known storage layout and code shape

A Simple Low-Level IR (1)

- This example is from Appel, but details aren't really important. What matters is to get a feel for the level of detail involved.
- Expressions:
 - $\text{CONST}(i)$ – integer constant i
 - $\text{TEMP}(t)$ – temporary t (i.e., register)
 - $\text{BINOP}(op, e1, e2)$ – application of op to $e1, e2$
 - $\text{MEM}(e)$ – contents of memory at address e
 - Means value when used in an expression
 - Means address when used as target of assignment
 - $\text{CALL}(f, args)$ – apply function f to argument list $args$

Simple Low-Level IR (2)

- Statements
 - MOVE(TEMP t, e) – evaluate e and store in temporary t
 - MOVE(MEM(e1), e2) – evaluate e1 to yield address a; evaluate e2 and store at a
 - EXP(e) – evaluate expressions e and discard result
 - SEQ(s1,s2) – execute s1 followed by s2
 - NAME(n) – assembly language label n
 - JUMP(e) – jump to e, which can be a NAME label, or more complex (e.g., switch)
 - CJUMP(op,e1,e2,t,f) – evaluate e1 op e2; if true jump to label t, otherwise jump to f
 - LABEL(n) – defines location of label n in the code


Low-Level IR Example (1)

- Access a local variable at a known offset k from the frame pointer fp

- Linear


MEM(BINOP(PLUS, TEMP fp , CONST k))

- Tree



Low-Level IR Example (2)

- Access an array element $e[k]$, where each element takes up w storage locations



Instruction Selection Issues

- Given the low-level IR, there are many possible code sequences that implement it correctly
 - e.g. set %rax to 0 on x86-64 (did we miss any?)

movq \$0,%rax	salq 64,%rax
subq %rax,%rax	shrq 64,%rax
xorq %rax,%rax	imulq \$0,%rax
 - Many machine instructions do several things at once – e.g., register arithmetic and effective address calculation:
movq offset(%rbase, %rindex, scale), %rdest

Instruction Selection Criteria

- Several possibilities
 - Fastest
 - Smallest
 - Minimize power consumption (ex: don't use a function unit if leaving it powered-down is a win)
 - Reduce memory traffic
 - etc. etc.
- Sometimes not obvious
 - e.g., if one of the function units in the processor is idle and we can select an instruction that uses that unit, it effectively executes for free, even if that instruction wouldn't be chosen normally
 - (Some interaction with scheduling here...)
 - (and it might consume extra power, so bad if that matters)

Tree Pattern Matching

- Goal: find a sequence of machine instructions that perform the computation described by the program IR code
 - Describe machine instructions using same low-level IR used for program, then
 - Use tree pattern matching to pick instructions that match fragments of the program IR tree; use a combination of these to cover the whole IR tree

An Example Target Machine (1)


- Arithmetic Instructions

- (unnamed) r_i
- $\text{ADD } r_i \leftarrow r_j + r_k$

- $\text{MUL } r_i \leftarrow r_j * r_k$


- SUB and DIV are similar

- For some examples, we'll assume there is at least one register (R_0) hardwired to be 0 always




An Example Target Machine (2)

- Immediate Instructions
 - ADDI $ri \leftarrow rj + c$






- SUBI $ri \leftarrow rj - c$




An Example Target Machine (3)

- Load
 - `LOAD ri <- M[rj + c]`



An Example Target Machine (4)

- Store
 - STORE $M[rj + c] \leftarrow r_i$



Tree Pattern Matching (1)



- Goal: Tile the low-level IR tree with operation (instruction) trees
- A *tiling* is a collection of $\langle \text{node}, \text{op} \rangle$ pairs
 - node is a node in the tree
 - op is an operation tree
 - $\langle \text{node}, \text{op} \rangle$ means that op could implement the subtree at node

Tree Pattern Matching (2)

- A tiling “implements” a tree if it covers every node in the tree and the overlap between any two tiles (trees) is limited to a single node
 - If $\langle \text{node}, \text{op} \rangle$ is in the tiling, then node is also covered by a leaf of another operation tree in the tiling – unless it is the root
 - Where two operation trees meet, they must be compatible (i.e., expect the same value in the same location)

Generating Tilings

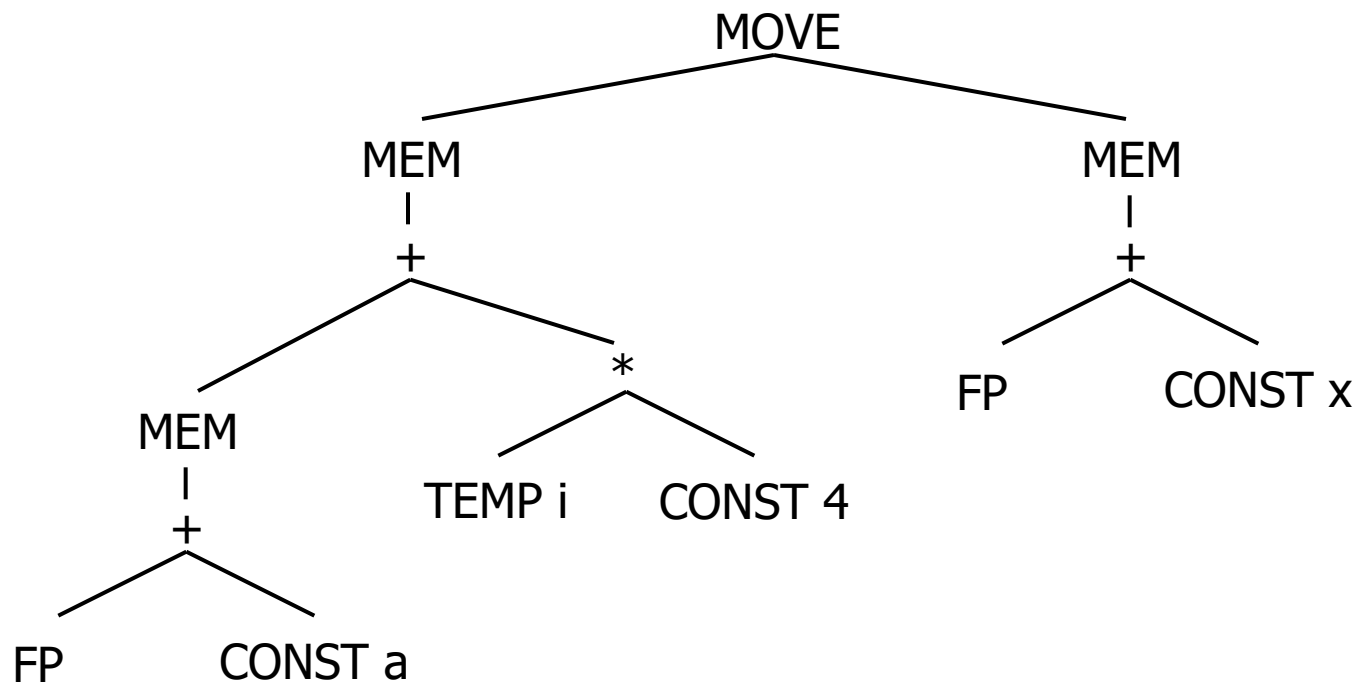
Two common algorithms

- Maximal munch: 
 - Top-down tree walk.
 - Find largest tile that fits each node
- Dynamic programming: 
 - Assign costs to each node in the tree
 $\text{cost} = \text{cost of individual node} + \text{subtree costs}$
 - Try all possible combinations bottom-up and pick cheapest


Heuristic: One instruction that “does more” is likely cheaper than several that do less

Slower, but optimal for a given cost model

Example – Tree for $a[i] := x$



Example – Tree for $a[i] := x$




Generating Code

Given a tiled tree, to generate code

- Do a postorder treewalk with node-dependant order for children
- Each tile corresponds to a code sequence; emit code sequences in order
- Connect tiles by using same register (or temporary) name to tie boundaries together

Example – Tree for $a[i] := x$



2. `LOAD r1 <- M[fp+aoff]`
4. `ADDI r2 <- 4 + r0`
5. `MUL r2 <- r2 * ri`
6. `ADD r1 <- r1 + r2`
8. `LOAD r2 <- M[fp+xoff]`
9. `STORE M[r1+0] <- r2`

Agenda

- Survey major pieces of a compiler back end
 - Instruction selection
 - Instruction scheduling
 - Register allocation
- And three particularly neat algorithms
 - Instruction selection by tree pattern matching
 - Instruction list scheduling
 - Register allocation by graph coloring

Instruction Scheduling

- Reorder instructions to minimize execution time given instruction and operand latencies
- Assume fixed program code at this point

Some Scheduling Issues (1)

- Many operations have non-zero latencies
- Modern machines can issue several operations per cycle
 - Want to take advantage of multiple function units on chip
- Loads & Stores may or may not block
 - may be (many) cycles after load/store starts to do other useful work before load/store result ready

Some Scheduling Issues (2)

- Branch costs vary
- Branches on some processors have delay slots
 - (relatively rare on newer processors)
- Modern processors have good heuristics to predict whether branches are taken and try to keep pipelines full, but good code from compiler makes these more effective

GOAL: Scheduler should reorder instructions to hide latencies, take advantage of multiple function units and delay slots, and help the processor effectively pipeline execution

Latencies for a Simple Example Machine

Operation	Cycles
LOAD	3
STORE	3
ADD	1
MULT	2
SHIFT	1
BRANCH	0 TO 8

Example: $w = w * 2 * x * y * z;$

Simple schedule

```
1  LOAD    r1 <- w
4  ADD     r1 <- r1,r1
5  LOAD    r2 <- x
8  MULT    r1 <- r1,r2
9  LOAD    r2 <- y
12 MULT    r1 <- r1,r2
13 LOAD    r2 <- z
16 MULT    r1 <- r1,r2
18 STORE   w <- r1
21 r1 free
```

2 registers, 20 cycles

Loads early

```
1  LOAD    r1 <- w
2  LOAD    r2 <- x
3  LOAD    r3 <- y
4  ADD     r1 <- r1,r1
5  MULT    r1 <- r1,r2
6  LOAD    r2 <- z
7  MULT    r1 <- r1,r3
9  MULT    r1 <- r1,r2
11 STORE   w <- r1
14 r1 is free
```

3 registers, 13 cycles

List Scheduling Algorithm Overview

- Build a precedence graph P of instructions, labeled with priorities (usually number of cycles on critical path to the end)
- Use list scheduling to construct a schedule, one cycle at a time
- Rename registers to avoid false dependencies and conflicts


Precedence Graph

- Nodes n are operations
- Attributes of each node
 - type – kind of operation
 - delay – latency until end of graph
- If node n_2 uses the result of node n_1 , there is an edge $e = (n_1, n_2)$ in the graph

Example

- Code

a	LOAD	r1 <- w
b	ADD	r1 <- r1,r1
c	LOAD	r2 <- x
d	MULT	r1 <- r1,r2
e	LOAD	r2 <- y
f	MULT	r1 <- r1,r2
g	LOAD	r2 <- z
h	MULT	r1 <- r1,r2
i	STORE	w <- r1



List Scheduling

- Construct a schedule, one cycle at a time
 - Keep a list of operations that are ready to execute
 - At each cycle, chose a ready operation and schedule it
 - Best pick: one that is on the “critical path” – i.e., an instruction that has longest latency from end of graph
 - Update ready list, deleting scheduled op and add ones that will be ready on next cycle

List Scheduling Algorithm


```
Cycle = 1; Ready = leaves of P; Active = empty;
while (Ready and/or Active are not empty)
    if (Ready is not empty)
        remove an op from Ready;
        S(op) = Cycle;
        Active = Active  $\cup$  op;
    Cycle++;
    for each op in Active
        if (S(op) + delay(op) <= Cycle)
            remove op from Active;
            for each successor s of op in P
                if (s is ready – i.e., all operands available)
                    add s to Ready
```

Example

- Code

a	LOAD	r1 <- w
b	ADD	r1 <- r1,r1
c	LOAD	r2 <- x
d	MULT	r1 <- r1,r2
e	LOAD	r2 <- y
f	MULT	r1 <- r1,r2
g	LOAD	r2 <- z
h	MULT	r1 <- r1,r2
i	STORE	w <- r1

cycle: 1
ready: a c e g
active: --



#	instr	done
1	a LOAD	4

Example


- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```

```

cycle:    ± 2
ready:    a c e g
active:    a
    
```



#	instr	done
1	a LOAD	4
2	c LOAD	5


Example

- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```

cycle: ~~1~~ 2 3
 ready: a e e g
 active: a c



#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6

Example


- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```

```

cycle:   1 2 3 4
ready:   a c e g b
active:  a c e
    
```




#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6
4	b ADD	5

Example

- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```



#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6
4	b ADD	5
5	d MULT	7

```


cycle: 1 2 3 4 5
ready: a c e g b d
active: a c e b
    
```

Example

- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```



#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6
4	b ADD	5
5	d MULT	7
6	g LOAD	9

```


cycle:  1 2 3 4 5 6
ready:  a c e g b d
active: a c e b d
    
```

Example

- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```



#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6
4	b ADD	5
5	d MULT	7
6	g LOAD	9
7	f MULT	9

```


cycle:  1 2 3 4 5 6 7
ready:  a c e g b d f
active: a c e b d g
    
```

Example

- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```



#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6
4	b ADD	5
5	d MULT	7
6	g LOAD	9
7	f MULT	9
8	---	

```


cycle:  1 2 3 4 5 6 7 8
ready:  a c e g b d f
active: a c e b d g f
    
```

Example

- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
  
```



#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6
4	b ADD	5
5	d MULT	7
6	g LOAD	9
7	f MULT	9
8	---	
9	h MULT	11

```


cycle:  1 2 3 4 5 6 7 8 9
ready:  a c e g b d f h
active: a c e b d g f
  
```

Example

- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```



#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6
4	b ADD	5
5	d MULT	7
6	g LOAD	9
7	f MULT	9
8	---	
9	h MULT	11
10	---	

```

cycle:  1 2 3 4 5 6 7 8 9 10
ready:  a c e g b d f h
active: a c e b d g f h
    
```

Example


- Code

```

a  LOAD    r1 <- w
b  ADD     r1 <- r1,r1
c  LOAD    r2 <- x
d  MULT    r1 <- r1,r2
e  LOAD    r2 <- y
f  MULT    r1 <- r1,r2
g  LOAD    r2 <- z
h  MULT    r1 <- r1,r2
i  STORE   w <- r1
    
```

```

cycle:  1 2 3 4 5 6 7 8 9 10 11
ready:  a c e g b d f h i
active:  a c e b d g f h
    
```



#	instr	done
1	a LOAD	4
2	c LOAD	5
3	e LOAD	6
4	b ADD	5
5	d MULT	7
6	g LOAD	9
7	f MULT	9
8	---	
9	h MULT	11
10	---	
11	i STORE	14

Forward vs Backwards

- Alternative: backward list scheduling
 - Work from the root to the leaves
 - Schedules instructions from end to beginning of the block
- In practice, compilers try both and pick the result that minimizes costs
 - Little extra expense since the precedence graph and other information can be reused
 - Different directions win in different cases

Beyond Basic Blocks

- List scheduling dominates, but moving beyond basic blocks can improve quality of the code.
Some possibilities:
 - Schedule extended basic blocks (sequences of basic blocks)
 - Watch for exit points – limits reordering or requires compensating
 - Trace scheduling
 - Use profiling information to select regions for scheduling using traces (paths) through code
 - Optimize schedules for high-frequency paths

Agenda

- Survey major pieces of a compiler back end
 - Instruction selection
 - Instruction scheduling
 - Register allocation
- And three particularly neat algorithms
 - Instruction selection by tree pattern matching
 - Instruction list scheduling
 - Register allocation by graph coloring

k

- Intermediate code typically assumes infinite number of registers
- Real machine has k registers available
- Goals
 - Produce correct code that uses k or fewer registers
 - Minimize added loads and stores
 - Minimize space needed for spilled values
 - Do this efficiently – $O(n)$, $O(n \log n)$, maybe $O(n^2)$

Register Allocation

- Task
 - At each point in the code, pick the values to keep in registers
 - Insert code to move values between registers and memory if needed
 - No additional transformations – scheduling should have done its job
 - But we will usually rerun scheduling if we insert spill code
 - Minimize inserted code, both dynamically and statically

Allocation vs Assignment

- Allocation: deciding which values to keep in registers
- Assignment: choosing specific registers for values
- Compiler must do both

Local Register Allocation

- Apply to basic blocks
- Produces decent register usage inside a block
 - But can have inefficiencies at boundaries between blocks
- Two variations: top-down, bottom-up

Top-down Local Allocation

- Principle: keep most heavily used values in registers
 - Priority = # of times register referenced in block
- If more virtual registers than physical,
 - Reserve some registers for values allocated to memory
 - Need enough to address and load two operands and store result
 - Other registers dedicated to “hot” values
 - But are tied up for entire block with particular value, even if only needed for part of the block

Bottom-up Local Allocation (1)

- Keep a list of available registers (initially all registers at beginning of block)
- Scan the code
- Allocate a register when one is needed
- Free register as soon as possible
 - In $x := y \text{ op } z$, free y and z if they are no longer needed before allocating x

Bottom-up Local Allocation (2)

- If no registers are free when one is needed for allocation:
 - Look at values assigned to registers – find the one not needed for longest forward stretch in the code
 - Insert code to spill the value to memory and insert code to reload it when needed later
 - If a copy already exists in memory, no need to spill

Local "bottom-up" Register Allocation, -1

1. ; load v2 from memory
2. ; load v3 from memory
3. $v1 = v2 + v3$
4. ; load v5, v6 from memory
5. $v4 = v5 - v6$
6. $v7 = v2 - 29$
7. ; load v9 from memory
8. $v8 = -v9$
9. $v10 = v6 * v4$
10. $v11 = v10 - v3$

- Still in LIR. So lots (too many!) virtual registers required (v2, etc).
- Grey instructions (1,2,4,7) load operands from memory into virtual registers.
- We will ignore these going forward. Focus on mapping virtual to physical.

Local "bottom-up" Register Allocation, 0

1. $v1 = v2 + v3$
2. $v4 = v5 - v6$
3. $v7 = v2 - 29$
4. $v8 = -v9$
5. $v10 = v6 * v4$
6. $v11 = v10 - v3$

pReg	vReg
R1	-
R2	-
R3	-
R4	-



vReg	NextRef
v1	1
v2	1
v3	1
v4	2
v5	2
v6	2
v7	3
v8	4
v9	4
v10	5
v11	6

Local "bottom-up" Register Allocation, 1

1. $v1 = v2 + v3$
2. $v4 = v5 - v6$
3. $v7 = v2 - 29$
4. $v8 = -v9$
5. $v10 = v6 * v4$
6. $v11 = v10 - v3$

pReg	vReg
R1	v2
R2	v3
R3	v1
R4	-

$R3 = R1 + R2$

vReg	NextRef
v1	1 ∞
v2	1 3
v3	1 6
v4	2
v5	2
v6	2
v7	3
v8	4
v9	4
v10	5
v11	6

Local "bottom-up" Register Allocation, 2

1. $v1 = v2 + v3$
2. $v4 = v5 - v6$
3. $v7 = v2 - 29$
4. $v8 = -v9$
5. $v10 = v6 * v4$
6. $v11 = v10 - v3$

pReg	vReg
R1	v2
R2	v3 v4
R3	v1 v6
R4	v5

$R3 = R1 + R2$
 ; spill R3
 ; spill R2? - no - still *clean*
 $R2 = R4 - R3$

vReg	NextRef
v1	∞
v2	3
v3	6
v4	2 5
v5	2 ∞
v6	2 5
v7	3
v8	4
v9	4
v10	5
v11	6

Local "bottom-up" Register Allocation, 3

1. $v1 = v2 + v3$
2. $v4 = v5 - v6$
3. $v7 = v2 - 29$
4. $v8 = -v9$
5. $v10 = v6 * v4$
6. $v11 = v10 - v3$

pReg	vReg
R1	v2
R2	v4
R3	v6
R4	v5 v7

$R3 = R1 + R2$
 ; spill R3
 ; spill R2? - no!
 $R2 = R4 - R3$
 ; spill R4? - no!
 $R4 = R1 - 29$

And so on . . .

vReg	NextRef
v1	∞
v2	3 ∞
v3	6
v4	5
v5	∞
v6	5
v7	3 ∞
v8	4
v9	4
v10	5
v11	6

Bottom-Up Allocator

- Invented about once per decade
 - Sheldon Best, 1955, for Fortran I
 - Laslo Belady, 1965, for analyzing paging algorithms
 - William Harrison, 1975, ECS compiler work
 - Chris Fraser, 1989, LCC compiler
 - Vincenzo Liberatore, 1997, Rutgers
- Will be reinvented again, no doubt
- Many arguments for optimality of this

Global Register Allocation by Graph Coloring

- Convert the (seemingly) infinite sequence of temporary data references, t_1, t_2, \dots into assignments to finite number of actual registers
- Goal: Use available registers with minimum spilling
- Problem: Minimizing the number of registers is NP-complete ... it is equivalent to chromatic number – minimum colors needed to color nodes of a graph so no edge connects the same color

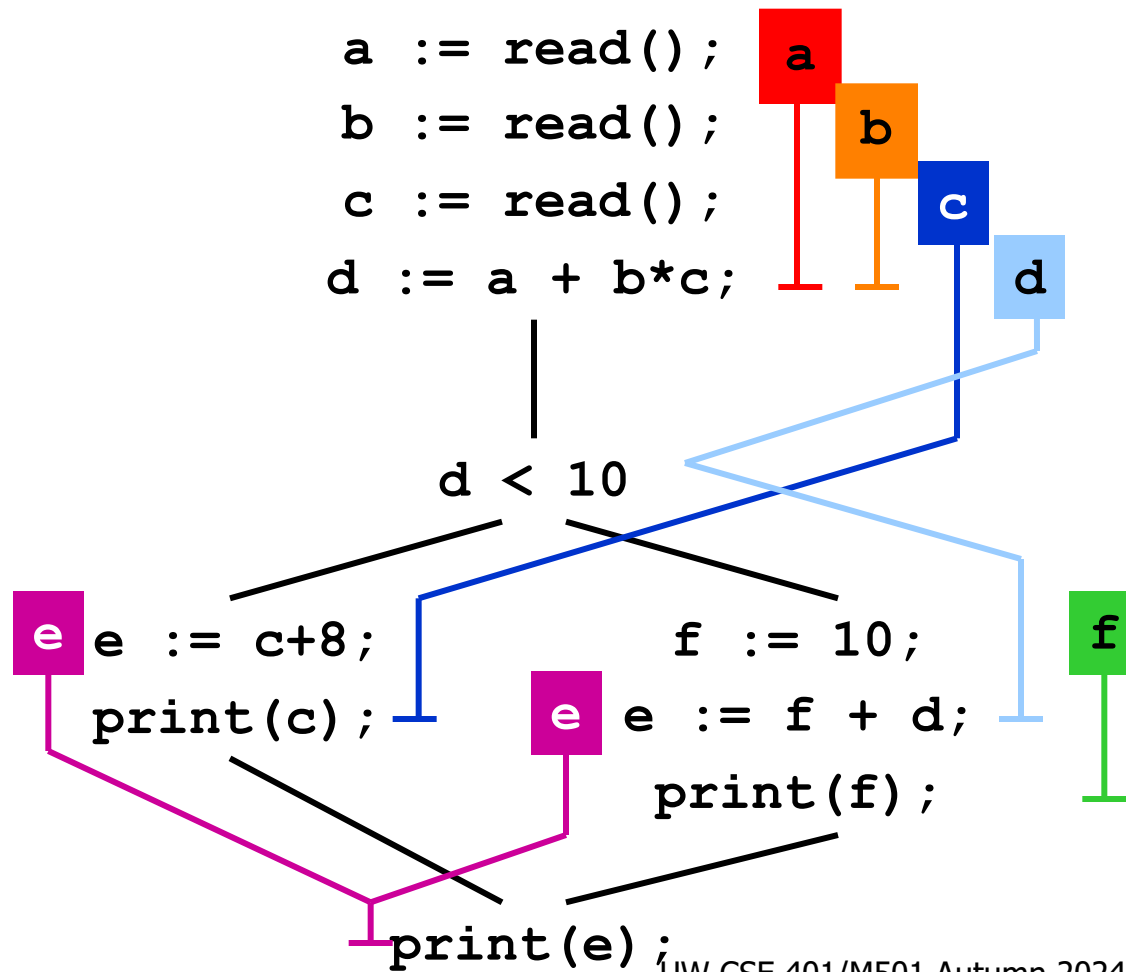
Begin With Data Flow Graph

- procedure-wide register allocation
- only live variables require register storage

dataflow analysis: a variable is **live** at node N if *the value* it holds is used on some path further down the control-flow graph; otherwise it is **dead**

- two variables(values) interfere when their live ranges overlap


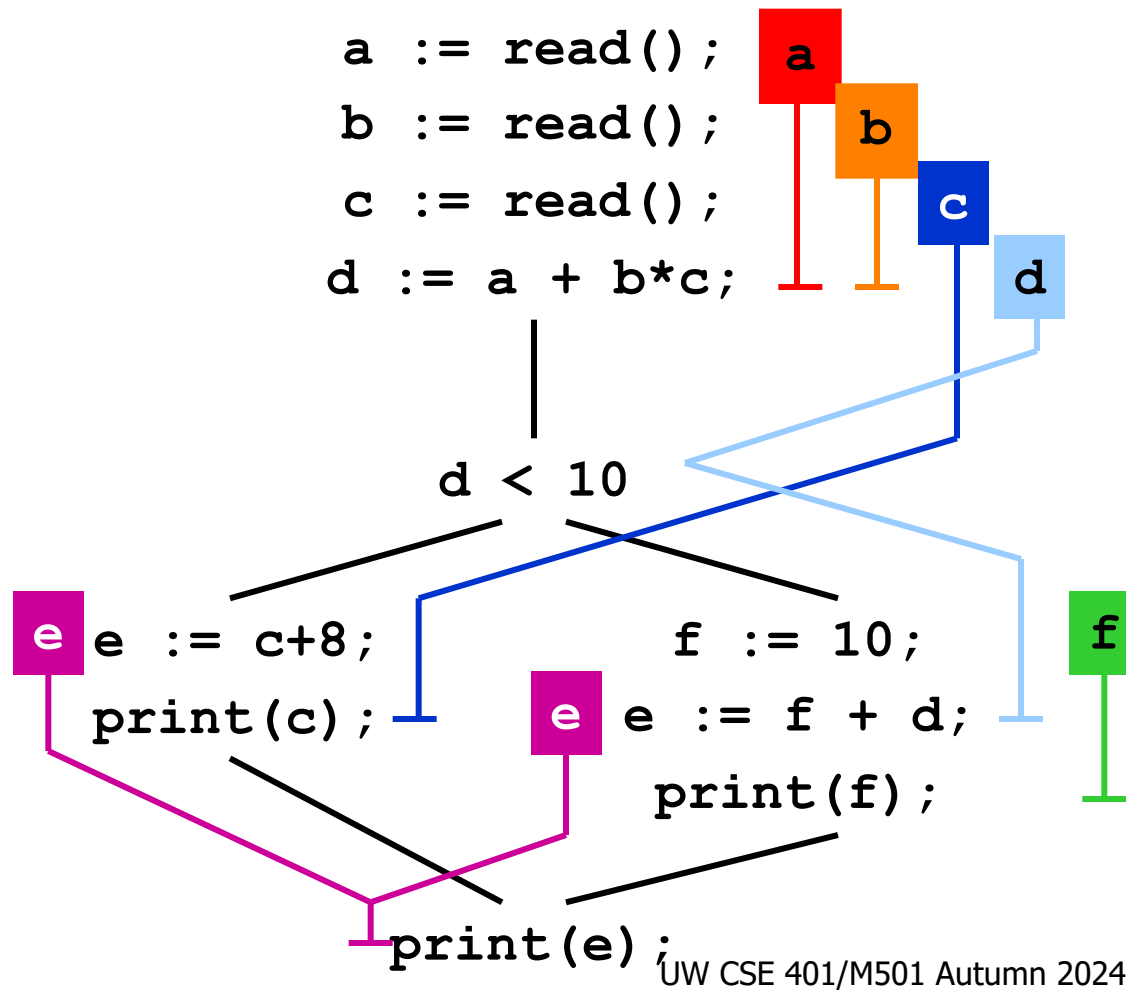
Live Variable Analysis



```


a := read();
b := read();
c := read();
d := a + b*c;
if (d < 10 ) then
    e := c+8;
    print(c);
else
    f := 10;
    e := f + d;
    print(f);
fi
print(e);
  
```

Register Interference Graph




Graph Coloring


- NP complete problem
- Heuristic: color easy nodes last
 - find node N with lowest degree
 - remove N from the graph
 - color the simplified graph
 - set color of N to the first color that is not used by any of N 's neighbors
- Basics due to Chaitin (1982), refined by Briggs (1992)




Apply Heuristic




Apply Heuristic




Apply Heuristic




Apply Heuristic




Continued




Continued




Continued




Continued




Continued




Continued




Continued




Continued



Continued



Final Assignment



```
a := read();  
b := read();  
c := read();  
d := a + b*c;  
if (d < 10) then  
    e := c+8;  
    print(c);  
else  
    f := 10;  
    e := f + d;  
    print(f);  
fi  
print(e);
```

Some Graph Coloring Issues

- May run out of registers
 - Solution: insert spill code and reallocate
- Special-purpose and dedicated registers
 - Examples: function return register, function argument registers, registers required for particular instructions
 - Solution: “pre-color” some nodes to force allocation to a particular register

Live Ranges

- Real graph-coloring register allocators don't allocate temp registers – they allocate *live ranges*
- A *live range*
 - In a basic block is the interval between a defn and last use
 - In a CFG, similar but more complex. Result is a coherent set of definitions and uses.
 - Every definition can reach every use
 - Every use that a definition can reach is in the same live range
- Idea: disjoint uses of a variable in different parts of the program don't actually interfere so they should be in separate live ranges
 - So we build a SSA form of the IR to construct the interference graph!

Live Ranges: Example

1. `loadi ... → rfp`
2. `loadai rfp, 0 → rw`
3. `loadi 2 → r2`
4. `loadai rfp,xoffset → rx`
5. `loadai rfp,yoffset → ry`
6. `loadai rfp,zoffset → rz`
7. `mult rw, r2 → rw`
8. `mult rw, rx → rw`
9. `mult rw, ry → rw`
10. `mult rw, rz → rw`
11. `storeai rw → rfp, 0`

Register	Interval
rfp	[1,11]
rw	[2,7]
rw	[7,8]
rw	[8,9]
rw	[9,10]
rw	[10,11]
r2	[3,7]
rx	[4,8]
ry	[5,9]
rz	[6,10]

Coalescing Live Ranges

- Idea: if two live ranges are connected by a copy operation ($\text{MOV } r_i \rightarrow r_j$) but do not otherwise interfere, then the live ranges can be coalesced (combined)
 - Rewrite all references to r_j to use r_i
 - Remove the copy instruction
- Then need to fix up interference graph


Advantages?

- Makes the code smaller, faster (no copy operation)
- Shrinks set of live ranges
- Reduces the degree of any live range that interfered with both live ranges r_i, r_j
- But: coalescing two live ranges can prevent coalescing of others, so ordering matters
 - Best: Coalesce most frequently executed ranges first (e.g., inner loops)
- Can have a substantial payoff – do it!

Graph Representation

- The interference graph representation drives the time and space requirements for the allocator (and maybe the compiler)
- Not unknown to have $O(5K)$ nodes and $O(1M)$ edges
- Dual representation works best
 - Triangular bit matrix for efficient access to interference information
 - Vector of adjacency vectors for efficient access to node neighbors

Overall Structure



- Then you may want to iterate with additional instruction selection and scheduling passes, particularly on a complex machine where operations can have both memory or register operands (e.g., x86)

And that's it!

Modulo all the picky details, that is...