

REGISTER ALLOCATION

Baishakhi Ray

Fall 2019



The Register Allocation Problem

- Intermediate code uses unlimited temporaries
 - Simplifies code generation and optimization
 - Complicates final translation to assembly
- Typical intermediate code uses too many temporaries
- The problem:
 - Rewrite the intermediate code to use no more temporaries than there are machine registers
- Method:
 - Assign multiple temporaries to each register – But without changing the program behavior

An Example

- Consider the program

a := c + d

e := a + b

f := e - 1

- Assume a and e dead after use
 - Temporary a can be “reused” after e := a + b
 - So can temporary e
- Can allocate a, e, and f all to one register (r1):

r1 := r2 + r3

r1 := r1 + r4

r1 := r1 - 1

- A dead temporary is not needed
 - A dead temporary can be reused

The Idea

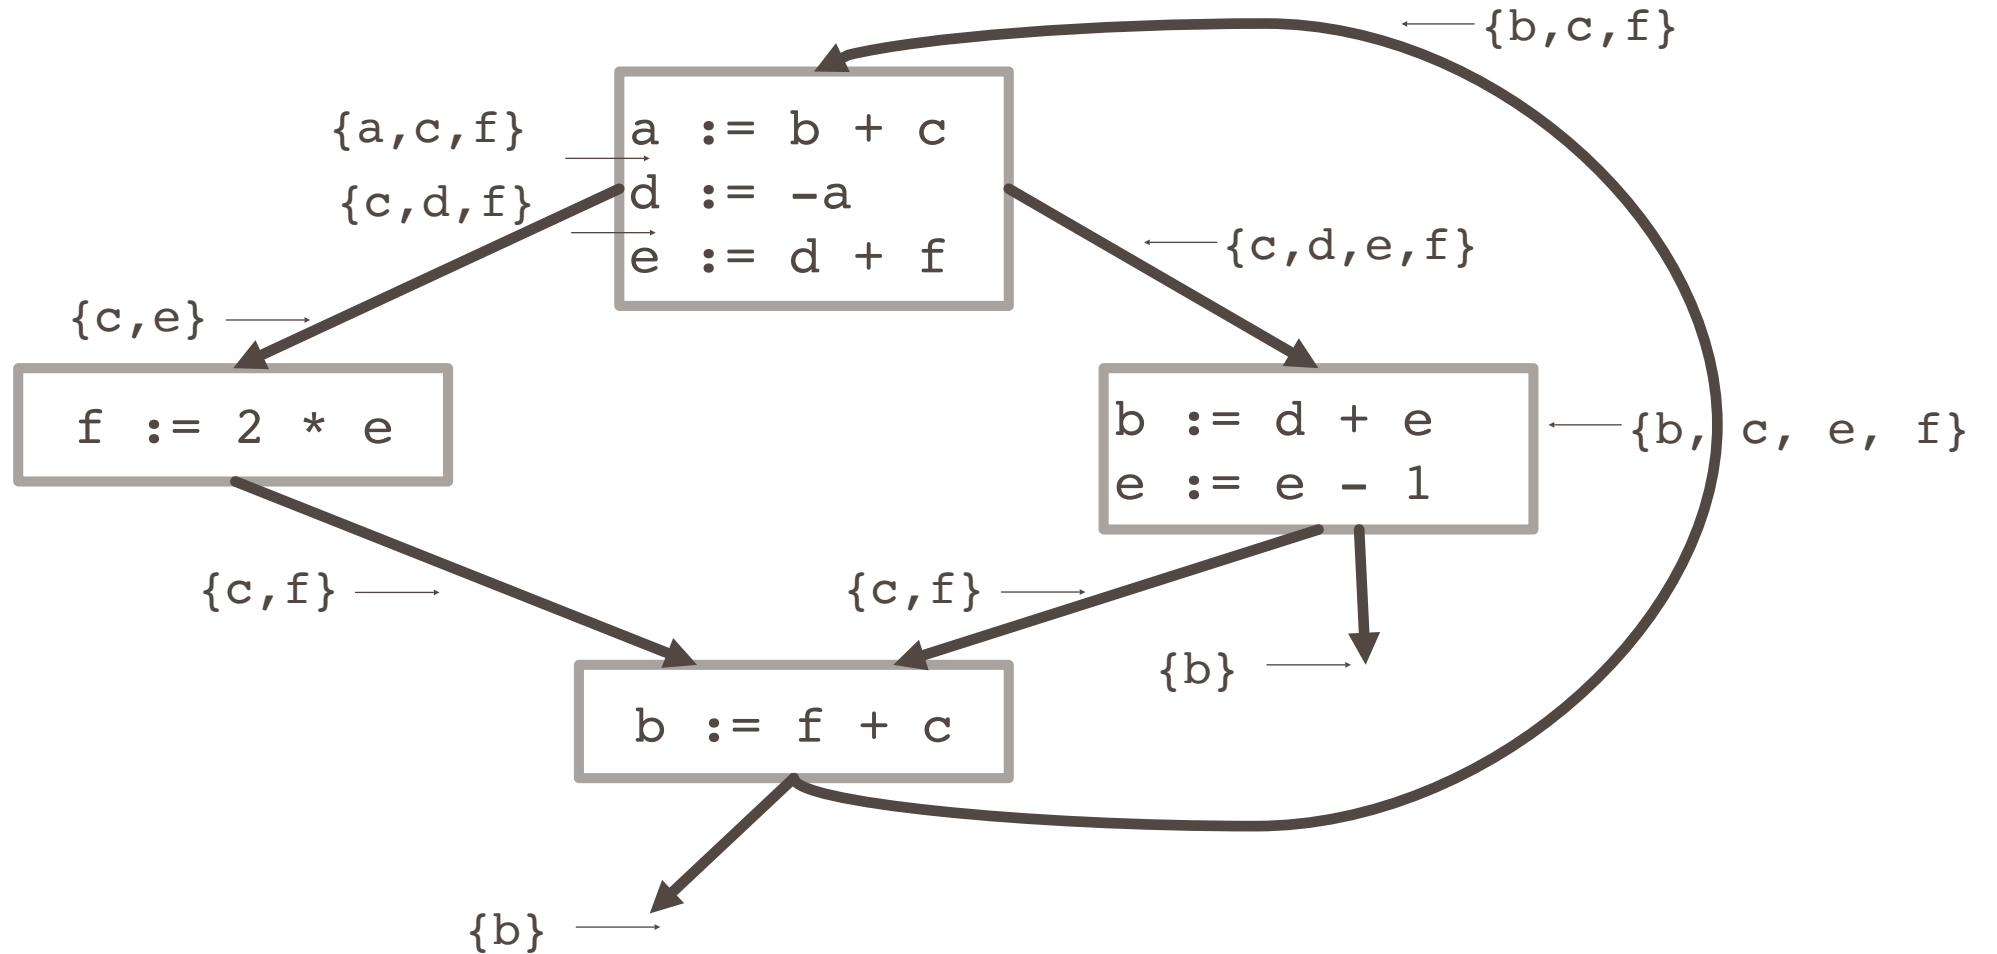
- Temporaries t_1 and t_2 can share the same register if at any point in the program at most one of t_1 or t_2 is live.

i.e.,

- If t_1 and t_2 are live at the same time, they cannot share a register

Algorithm: Part I


- Compute live variables for each point:



The Register Interference Graph

- Construct an undirected graph
 - A node for each temporary
 - An edge between t_1 and t_2 if they are live simultaneously at some point in the program
- This is the register interference graph (RIG)
 - Two temporaries can be allocated to the same register if there is no edge connecting them


Example



- E.g., b and c cannot be in the same register
- E.g., b and d could be in the same register

Definitions


- A coloring of a graph is an assignment of colors to nodes, such that nodes connected by an edge have different colors
- A graph is k-colorable if it has a coloring with k colors



- There is no coloring with less than 4 colors
- There are 4-colorings of this graph

Example After Register Allocation

- Compute live variables for each point:



Computing Graph Colorings

- How do we compute graph colorings?
- It isn't easy:
 - This problem is very hard (NP-hard).
 - No efficient algorithms are known.
 - Solution: use heuristics
 - A coloring might not exist for a given number of registers
 - Solution: later

Graph Coloring Heuristic

- **Observation:**

- Pick a node t with fewer than k neighbors in RIG
- Eliminate t and its edges from RIG
- If resulting graph is k -colorable, then so is the original graph

- **Why?**


- Let c_1, \dots, c_n be the colors assigned to the neighbors of t in the reduced graph
- Since $n < k$ we can pick some color for t that is different from those of its neighbors

Graph Coloring Heuristic

- The following works well in practice:
 - Pick a node t with fewer than k neighbors
 - Put t on a stack and remove it from the RIG
 - Repeat until the graph has one node
- Assign colors to nodes on the stack
 - Start with the last node added
 - At each step pick a color different from those assigned to already colored neighbors

Graph Coloring Example (1)


- Start with the RIG and with $k = 4$:



Stack: {}

- Remove a


Graph Coloring Example (2)



Stack: {a}

- Remove d


Graph Coloring Example (3)



Stack: {d, a}

- Remove c

Graph Coloring Example (4)



Stack: {c, d, a}

- Remove b

Graph Coloring Example (5)



Stack: {b, c, d, a}

- Remove e

Graph Coloring Example (6)

f ●

Stack: {e, b, c, d, a}

- Remove f

Graph Coloring Example (7)


Stack: {f, e, b, c, d, a}

Graph Coloring Example (8)

r1 f ●

Stack: {e, b, c, d, a}


Graph Coloring Example (9)



Stack: {b, c, d, a}


- e must be in a different register from f

Graph Coloring Example (10)




Stack: {c, d, a}

Graph Coloring Example (11)




Stack: {d, a}

Graph Coloring Example (12)



Stack: {a}


Graph Coloring Example (13)



Stack: {}


What if the Heuristic Fails?

- What if all nodes have k or more neighbors ?
- Example: Try to find a 3-coloring of the RIG:




What if the Heuristic Fails?

- Remove a and get stuck (as shown below)
- Pick a node as a candidate for spilling
 - A spilled temporary “lives” in memory
 - Assume that f is picked as a candidate




What if the Heuristic Fails?

- Remove f and continue the simplification
 - Simplification now succeeds: b, d, e, c



What if the Heuristic Fails?

- Eventually we must assign a color to f
- We hope that among the 4 neighbors of f we use less than 3 colors \Rightarrow optimistic coloring




Spilling

- If optimistic coloring fails, we spill f
 - Allocate a memory location for f
 - Typically in the current stack frame
 - Call this address fa
- Before each operation that reads f , insert
 $f := \text{load } fa$
- After each operation that writes f , insert
 $\text{store } f, fa$

Spilling Example

- This is the new code after spilling f

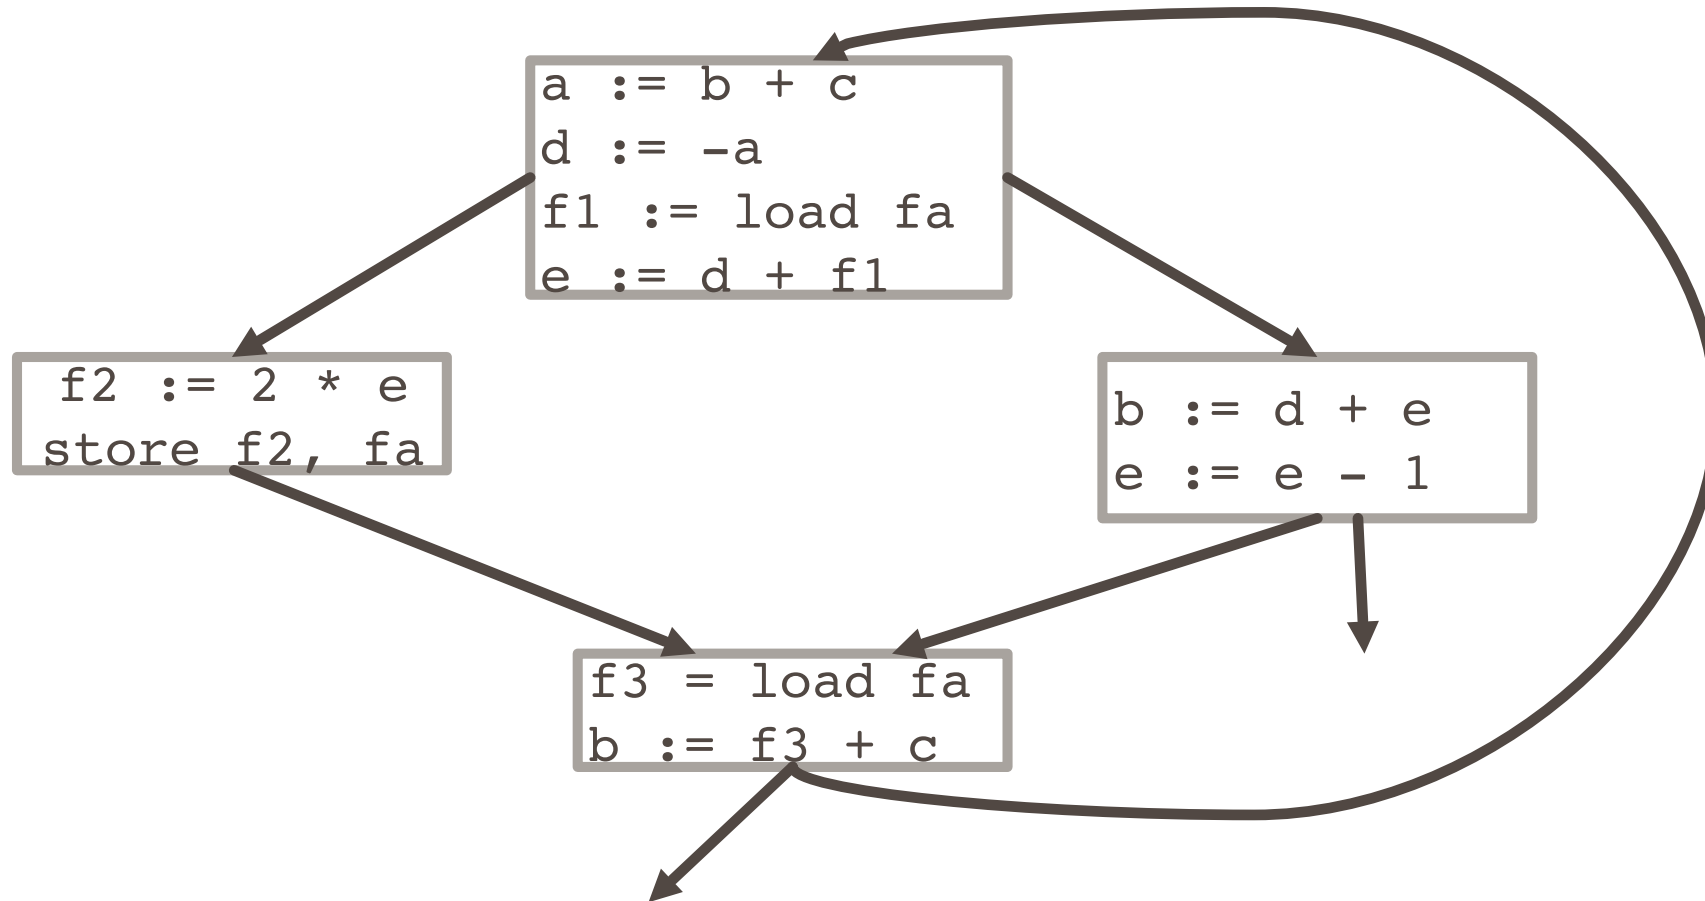


A Problem

- This code reuses the register name f
- Correct, but suboptimal
 - Should use distinct register names whenever possible
 - Allows different uses to have different colors

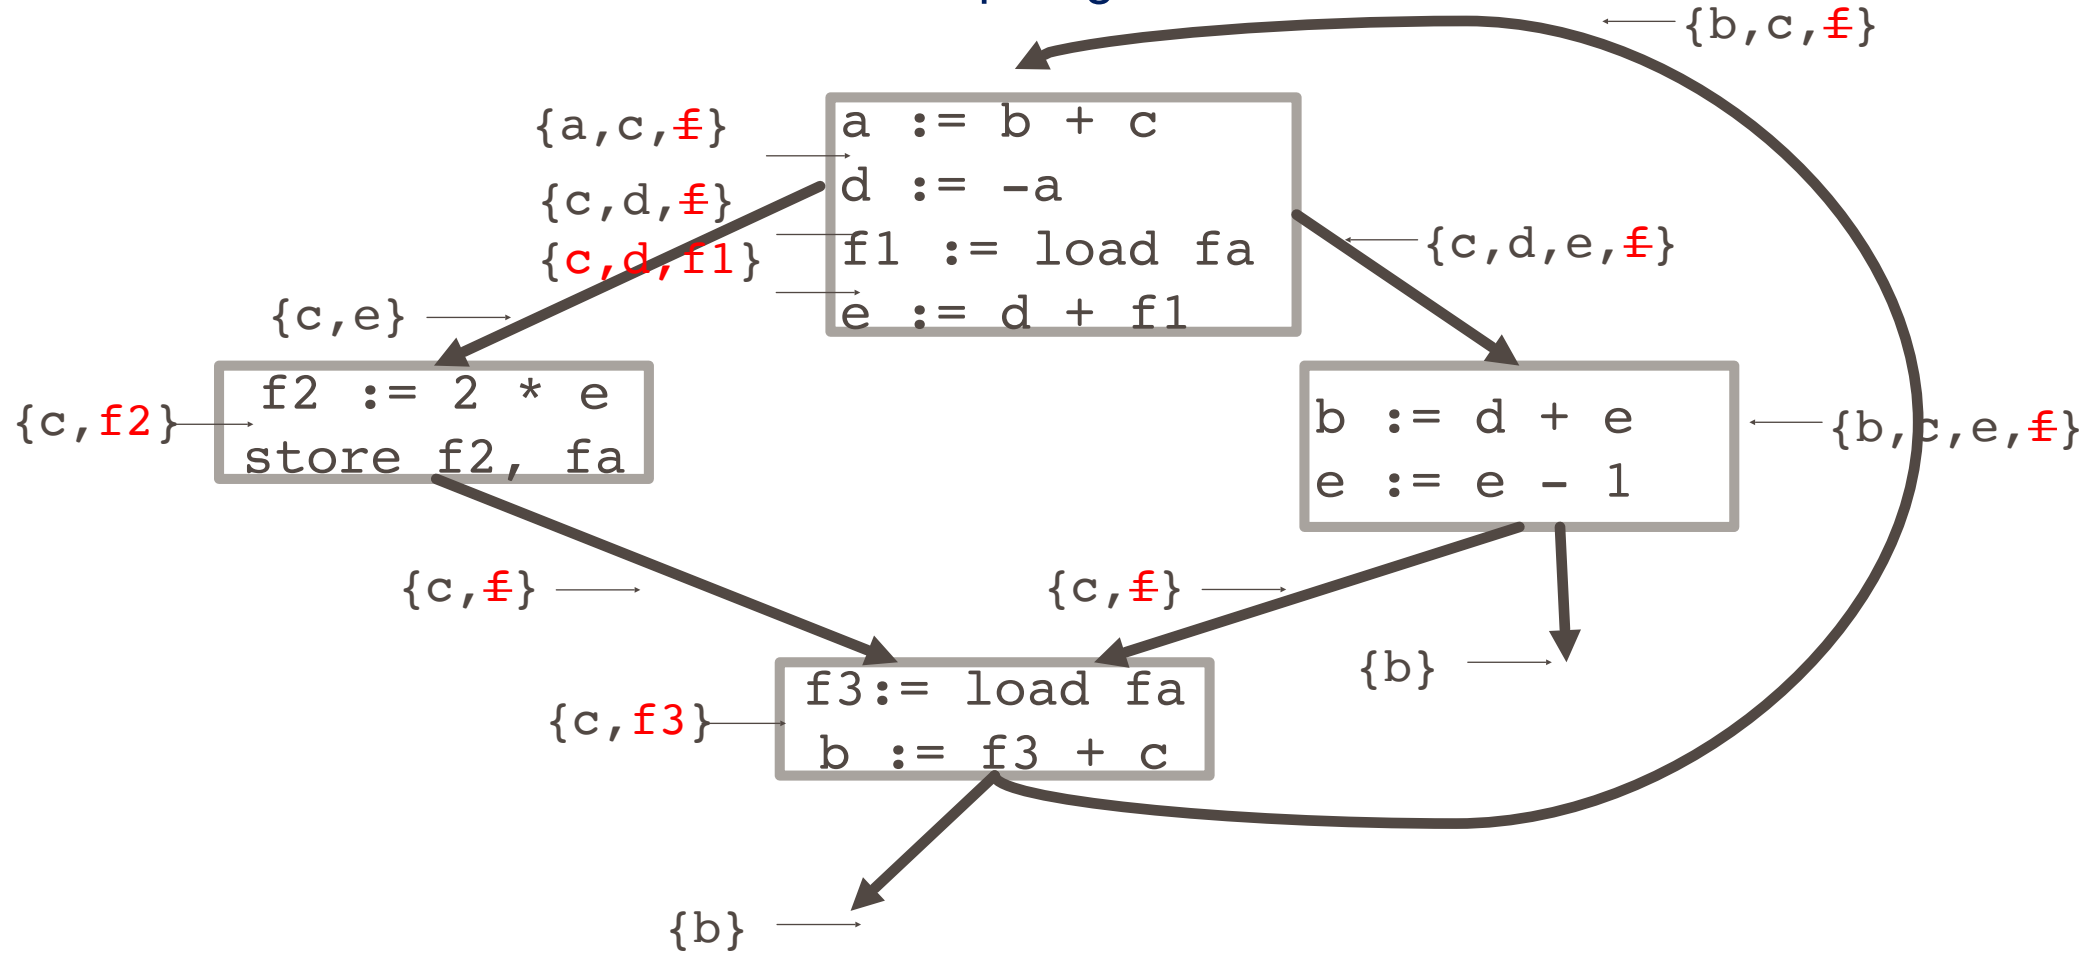
Spilling Example

- This is the new code after spilling f



Recomputing Liveness Information

- The new liveness information after spilling:




Recomputing Liveness Information

- **New liveness information is almost as before**
 - Note f has been split into three temporaries
- **f_i is live only**
 - Between a $f_i := \text{load } f_a$ and the next instruction
 - Between a $\text{store } f_i, f_a$ and the preceding instr.
- **Spilling reduces the live range of f**
 - And thus reduces its interferences
 - Which results in fewer RIG neighbors

Recompute RIG After Spilling

- Some edges of the spilled node are removed
- In our case f still interferes only with c and d
- And the resulting RIG is 3-colorable



Spilling Notes

- Additional spills might be required before a coloring is found
- The tricky part is deciding what to spill
 - But any choice is correct
- Possible heuristics:
 - Spill temporaries with most conflicts
 - Spill temporaries with few definitions and uses
 - Avoid spilling in inner loops

Caches

- **Compilers are very good at managing registers**
 - Much better than a programmer could be
- **Compilers are not good at managing caches**
 - This problem is still left to programmers
 - It is still an open question how much a compiler can do to improve cache performance
- **Compilers can, and a few do, perform some cache optimizations**

Cache Optimization

- Consider the loop

```
for(j := 1; j < 10; j++)  
  for(i=1; i<1000;i++)  
    a[i] *= b[i]
```

- This program has terrible cache performance
 - Why?

Cache Optimization

- Consider the program

```
for(i=1; i<1000; i++)  
    for(j := 1; j < 10; j++)  
        a[i] *= b[i]
```

- Computes the same thing
 - But with much better cache behavior
 - Might actually be more than 10x faster
- A compiler can perform this optimization
 - called loop interchange

Conclusions

- Register allocation is a “must have” in compilers:
 - Because intermediate code uses too many temporaries
 - Because it makes a big difference in performance
- Register allocation is more complicated for CISC machines