

## CS143 Spring 2023 – Written Assignment 4 – Solutions

1. Consider the following program in Cool, representing a “slightly” over-engineered implementation which calculates the factorial of 3 using an operator class and a reduce() method:

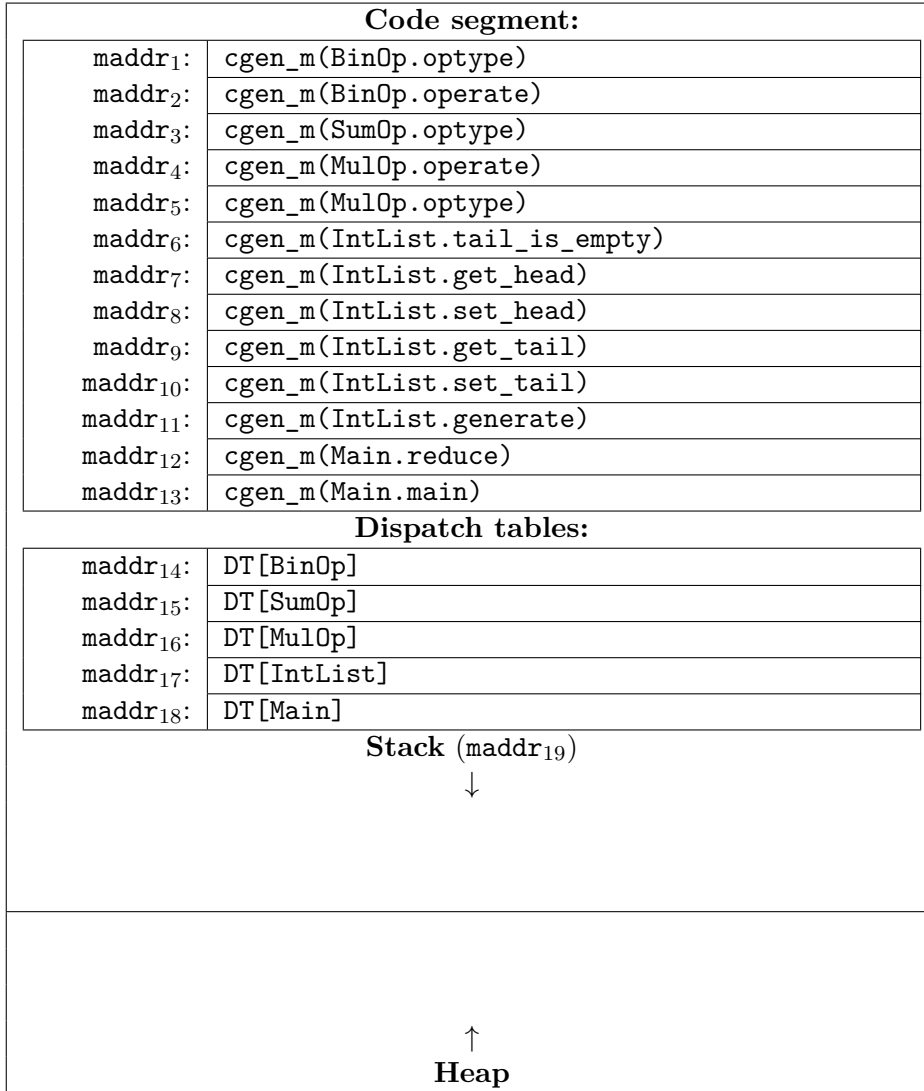
```
1 class BinOp {
2     optype(): String {
3         "BinOp"
4     };
5
6     operate(a: Int, b: Int): Int {
7         a + b
8     };
9 };
10
11 class SumOp inherits BinOp {
12     optype(): String {
13         "SumOp"
14     };
15 };
16
17 class MulOp inherits BinOp {
18     optype(): String {
19         "MulOp"
20     };
21
22     operate(a: Int, b: Int): Int {
23         a * b
24     };
25 };
26
27 class IntList {
28     head: Int;
29     tail: IntList;
30     empty_tail: IntList; -- Do not assign.
31
32     tail_is_empty(): Bool {
33         tail = empty_tail
34     };
35
36     get_head(): Int { head };
37
38     set_head(n: Int): Int {
39         head <- n
40     };
41
42     get_tail(): IntList { tail };
43
44     set_tail(t: IntList): IntList {
45         tail <- t
46     };
47
48     -- Create a list [n, n-1, ..., 2, 1] for n ≥ 1
```

```

49     generate(n: Int): IntList {
50         let l: IntList <- New IntList in {
51             l.set_head(n); -- Point A
52             if (n = 1) then
53                 l.set_tail(empty_tail)
54             else
55                 l.set_tail(generate(n-1))
56             fi;
57             l;
58         }
59     };
60 };
61
62 class Main {
63     reduce(result: Int, op: BinOp, l: IntList): Int {{
64         result <- op.operate(result, l.get_head());
65         if (l.tail_is_empty()) then
66             result -- Point B
67         else
68             reduce(result, op, l.get_tail())
69         fi;
70     }};
71
72     main(): Object {
73         let op: BinOp <- New MulOp, l: IntList <- New IntList, io: IO
74         <- New IO in {
75             l <- l.generate(3);
76             io.out_int(self.reduce(1, op, l));
77         }
78     };

```

The following is an abstracted representation of a memory layout of the program generated by a hypothetical Cool compiler for the above code (note that this might or might not correspond to the layout generated by your compiler or the reference coolc):



In the above, **maddr<sub>i</sub>** represents the memory address at which the corresponding method's code or dispatch table starts. You should assume that the above layout is contiguous in memory.

- (a) The following is a representation of the dispatch table for class Main:

| Method Idx | Method Name         | Address                         |
|------------|---------------------|---------------------------------|
| 0          | <code>reduce</code> | <code>maddr<sub>12</sub></code> |
| 1          | <code>main</code>   | <code>maddr<sub>13</sub></code> |

Provide equivalent representations for the dispatch tables of BinOp, SumOp and IntList. You can ignore built-in methods like `abort`, `type_name`, and `copy`.

BinOp:

| Method Idx | Method Name          | Address                        |
|------------|----------------------|--------------------------------|
| 0          | <code>optype</code>  | <code>maddr<sub>1</sub></code> |
| 1          | <code>operate</code> | <code>maddr<sub>2</sub></code> |

SumOp:

| Method Idx | Method Name          | Address                        |
|------------|----------------------|--------------------------------|
| 0          | <code>optype</code>  | <code>maddr<sub>3</sub></code> |
| 1          | <code>operate</code> | <code>maddr<sub>2</sub></code> |

MulOp:

| Method Idx | Method Name          | Address                        |
|------------|----------------------|--------------------------------|
| 0          | <code>optype</code>  | <code>maddr<sub>5</sub></code> |
| 1          | <code>operate</code> | <code>maddr<sub>4</sub></code> |

IntList:

| Method Idx | Method Name                | Address                         |
|------------|----------------------------|---------------------------------|
| 0          | <code>tail_is_empty</code> | <code>maddr<sub>6</sub></code>  |
| 1          | <code>get_head</code>      | <code>maddr<sub>7</sub></code>  |
| 2          | <code>set_head</code>      | <code>maddr<sub>8</sub></code>  |
| 3          | <code>get_tail</code>      | <code>maddr<sub>9</sub></code>  |
| 4          | <code>set_tail</code>      | <code>maddr<sub>10</sub></code> |
| 5          | <code>generate</code>      | <code>maddr<sub>11</sub></code> |

- (b) Consider the state of the program at runtime when reaching (for the first time) the beginning of the line marked with the comment “Point A”. Give the object layout (as per Lecture 12) of every object currently on the heap which is of a class defined by the program (i.e. ignoring Cool base classes such as IO or Int). For attributes, you can directly represent Int values by integers and an unassigned pointer by **void**. However, note that in a real Cool program, Int is an object and would have its own object layout, omitted here for simplicity. Finally, you can assume class tags are numbers from 1 to 5 given in the same order as the one in which classes appear in the layout above.

**Answer:**

Main

|                     |
|---------------------|
| 5                   |
| 3                   |
| maddr <sub>18</sub> |

MulOp

|                     |
|---------------------|
| 3                   |
| 3                   |
| maddr <sub>16</sub> |

IntList (in Main.main)

|                     |
|---------------------|
| 4                   |
| 6                   |
| maddr <sub>17</sub> |
| 0                   |
| void                |
| void                |

IntList (in IntList.generate)

|                     |                     |
|---------------------|---------------------|
| 4                   | 4                   |
| 6                   | 6                   |
| maddr <sub>17</sub> | maddr <sub>17</sub> |
| 0                   | 3                   |
| void                | void                |
| void                | void                |

or

- (c) The following table represents an abstract view of the layout of the stack at runtime when reaching (for the first time) the beginning of the line marked with the comment “Point A”.

| Address                              | Method                        | Contents  | Description                  |
|--------------------------------------|-------------------------------|---|------------------------------|
| <code>maddr<sub>19</sub></code>      | <code>Main.main</code>        | <code>self</code>                                     | <code>arg<sub>0</sub></code> |
| <code>maddr<sub>19</sub> - 4</code>  | <code>Main.main</code>        | <code>...</code>                                      | Return                       |
| <code>maddr<sub>19</sub> - 8</code>  | <code>Main.main</code>        | <code>op</code>                                       | local                        |
| <code>maddr<sub>19</sub> - 12</code> | <code>Main.main</code>        | <code>l</code>  | local                        |
| <code>maddr<sub>19</sub> - 16</code> | <code>Main.main</code>        | <code>io</code>                                       | local                        |
| <code>maddr<sub>19</sub> - 20</code> | <code>IntList.generate</code> | <code>maddr<sub>19</sub> - 4</code>                   | FP                           |
| <code>maddr<sub>19</sub> - 24</code> | <code>IntList.generate</code> | <code>3</code>  | <code>arg<sub>1</sub></code> |
| <code>maddr<sub>19</sub> - 28</code> | <code>IntList.generate</code> | <code>self</code>                                     | <code>arg<sub>0</sub></code> |
| <code>maddr<sub>19</sub> - 32</code> | <code>IntList.generate</code> | <code>maddr<sub>13</sub> + <math>\delta</math></code> | Return                       |
| <code>maddr<sub>19</sub> - 36</code> | <code>IntList.generate</code> | <code>l</code>  | local                        |

Note that we are assuming there are no stack frames above `Main.main(...)`. This doesn’t necessarily match a real implementation of the Cool runtime system, where `main` must return control to the OS or the Cool runtime on exit. For the purposes of this exercise, feel free to ignore this issue. Also, since you don’t have the generated code for every method above, you cannot directly calculate the return address to be stored on the stack. You should however give it as `maddri +  $\delta$` , denoting an unknown address between `maddri` and `maddri+1`. This notation is used in the example above. For locals, you should use the variable name, but remember that in practice it is the heap address that gets stored in memory for objects.

Give a similar view of the stack at runtime when reaching (for the first time) the beginning of the line marked with the comment “Point B”.

| Address                  | Method      | Contents                       | Description    |
|--------------------------|-------------|--------------------------------|----------------|
| $\text{maddr}_{19}$      | Main.main   | self                           | $\text{arg}_0$ |
| $\text{maddr}_{19} - 4$  | Main.main   | ...                            | Return         |
| $\text{maddr}_{19} - 8$  | Main.main   | op                             | local          |
| $\text{maddr}_{19} - 12$ | Main.main   | 1                              | local          |
| $\text{maddr}_{19} - 16$ | Main.main   | io                             | local          |
| $\text{maddr}_{19} - 20$ | Main.reduce | $\text{maddr}_{19} - 4$        | FP             |
| $\text{maddr}_{19} - 24$ | Main.reduce | ptr to [3,2,1]                 | $\text{arg}_3$ |
| $\text{maddr}_{19} - 28$ | Main.reduce | ptr to MulOp                   | $\text{arg}_2$ |
| $\text{maddr}_{19} - 32$ | Main.reduce | 3                              | $\text{arg}_1$ |
| $\text{maddr}_{19} - 36$ | Main.reduce | self                           | $\text{arg}_0$ |
| $\text{maddr}_{19} - 40$ | Main.reduce | $\text{maddr}_{13} + \delta_1$ | Return         |
| $\text{maddr}_{19} - 44$ | Main.reduce | $\text{maddr}_{19} - 40$       | FP             |
| $\text{maddr}_{19} - 48$ | Main.reduce | ptr to [2,1]                   | $\text{arg}_3$ |
| $\text{maddr}_{19} - 52$ | Main.reduce | ptr to MulOp                   | $\text{arg}_2$ |
| $\text{maddr}_{19} - 56$ | Main.reduce | 6                              | $\text{arg}_1$ |
| $\text{maddr}_{19} - 60$ | Main.reduce | self                           | $\text{arg}_0$ |
| $\text{maddr}_{19} - 64$ | Main.reduce | $\text{maddr}_{12} + \delta_2$ | Return         |
| $\text{maddr}_{19} - 68$ | Main.reduce | $\text{maddr}_{19} - 64$       | FP             |
| $\text{maddr}_{19} - 72$ | Main.reduce | ptr to [1]                     | $\text{arg}_3$ |
| $\text{maddr}_{19} - 76$ | Main.reduce | ptr to MulOp                   | $\text{arg}_2$ |
| $\text{maddr}_{19} - 80$ | Main.reduce | 6                              | $\text{arg}_1$ |
| $\text{maddr}_{19} - 84$ | Main.reduce | self                           | $\text{arg}_0$ |
| $\text{maddr}_{19} - 88$ | Main.reduce | $\text{maddr}_{12} + \delta_2$ | Return         |

Notice that in the last three stack frames, we do not have a separate entry for “result”, since it is a formal parameter and not a local variable.

2. Consider the following assembly language used to program a stack machine (**r**, **r1**, and **r2** denote arbitrary registers):

- **push** *r*: copies the value of *r* and pushes it onto the stack.
- **top** *r*: copies the value at the top of the stack into *r*. This command does not modify the stack.
- **pop**: discards the value at the top of the stack.
- **swap**: swaps the value at top of the stack with the value right beneath it. E.g. if the stack was  $\langle \$, \dots, 5, 2 \rangle$  swap would change the stack to be  $\langle \$, \dots, 2, 5 \rangle$
- $r1+ = r2$ : adds *r1* and *r2* and saves the result in *r1*. *r1* may be the same as *r2*.
- $r1- = r2$ : subtracts *r2* from *r1* and saves the result in *r1*. *r1* may be the same as *r2*.
- **clamp** *r*: sets *r* to 0 if *r* is negative otherwise leaves *r* unchanged
- **jump** *r*: jumps to the address in *r* and resumes execution.
- **ite** *r1 r2 r3*: if *r1* is not equal to zero then jumps to the address in *r2* else jumps to the address in *r3*.
- **loadconst** *r* int: loads a constant int into *r*
- **loadlabel** *r* label: loads the address of a labeled code segment into *r*.

Provide a code generation function for each the of these instructions, except **loadlabel**, targeting MIPS. Assume that registers used in the stack language are valid MIPS registers. Use **\$sp** to hold a pointer to the top of the stack and a single temporary register **\$at** which is guaranteed to not appear in the stack language.



```

cgen(push r) =
    sw r 0($sp)
    addiu $sp $sp -4

cgen(top r) =
    lw r 4($sp)

cgen(pop) =
    addiu $sp $sp 4

cgen(swap) =
    lw $at 4($sp) // get the top of the stack
    sw $at 0($sp) // duplicate it
    lw $at 8($sp) // get the item below it
    sw $at 4($sp) // store it on top of the stack
    lw $at 0($sp) // get the duplicated item
    sw $at 8($sp) // store it in the second slot

cgen(r1 += r2) =
    addu r1 r1 r2

cgen(r1 -= r2) =
    subu r1 r1 r2

cgen(clamp r) =
    bltz r neg
    j done
neg:
    li r 0
done:

cgen(jump r) =
    jr r

cgen(ite r1 r2 r3) =
    li $at 0
    beq $at r1 false_branch
    jr r2
false_branch:
    jr r3

cgen(loadconst r int) =
    li r int

```

3. Suppose you want to add a for-loop construct to Cool, having the following syntax:

for id : Int  $\leftarrow e_1$  to  $e_2$  do  $e_3$  rof

The above for-loop expression is evaluated as follows: expressions  $e_1$  and  $e_2$  are evaluated only once, then the body of the loop ( $e_3$ ) is executed once for every value of id starting with the value of  $e_1$  and incremented by 1 thereafter until reaching the value of  $e_2$  (inclusive). Similar to the while loop, the for-loop returns void.

(a) Give the operational semantics for the for-loop construct above.

**Answer:** These are similar to the operational semantics of the let and while expressions, with some care applied to requirement that  $e_1$  and  $e_2$  be evaluated only once. First, we need a rule for when the for loop executes zero times (the bounds don't include any number):

$$\frac{so, S, E \vdash e_1 : Int(i_1), S_1 \quad so, S_1, E \vdash e_2 : Int(i_2), S_2 \quad i_1 > i_2}{so, S, E \vdash \text{for id : Int } \leftarrow e_1 \text{ to } e_2 \text{ do } e_3 \text{ rof : void}, S_2}$$

Then the general case where we do enter the loop:

$$\frac{\begin{array}{c} so, S, E \vdash e_1 : Int(i_1), S_1 \\ so, S_1, E \vdash e_2 : Int(i_2), S_2 \\ i_1 \leq i_2 \\ l_1 = newloc(S_2) \\ so, S_2[Int(i_1)/l_1], E[l_1/id] \vdash e_3 : v_3, S_3 \\ i_n = S_3(l_1) + 1 \\ so, S_3, E \vdash \text{for id : Int } \leftarrow i_n \text{ to } i_2 \text{ do } e_3 \text{ rof : void}, S_4 \end{array}}{so, S, E \vdash \text{for id : Int } \leftarrow e_1 \text{ to } e_2 \text{ do } e_3 \text{ rof : void}, S_4}$$

Where  $Int(i_k)$  represents a Cool Int object with  $i_k$  as its corresponding numerical value. When used in the program syntax,  $i_k$  refers to the corresponding integer literal.

- (b) Give the code generation function `cgen(for id : Int  $\leftarrow$   $e_1$  to  $e_2$  do  $e_3$  rof)` for this construct. Use the code generation conventions from the lecture. The result of `cgen(...)` must be MIPS code following the stack-machine with one accumulator model.

Assume that `cgen( $e_1$ )` (and similarly `cgen( $e_2$ )`) does integer unboxing, i.e., the evaluation result of  $e_1$  will be stored in `$a0` after executing `cgen( $e_1$ )`. You can use the instruction `ble r1, r2, label` in order to branch to `label` if  $r1 \leq r2$ .

**Answer:** Note that `id` is a new local introduced by the `for` construct. Thus, it gets assigned a space in the AR. In the code below we will assume that a variable  $i$  is defined while generating the code for the body of each method, and contains the offset from the frame pointer at which `id` should be stored in the AR. Also, we assume our compiler does integer unboxing, so we can operate directly with `id` (and other Cool Ints) as MIPS `int32` values directly, avoiding additional accesses to the heap.

```

1      cgen(for id: Int <- e1 to e2 do e_3 rof, nt) =
2          cgen(e1, nt)
3          sw $a0 i($fp)                # Save id=val(e1) on i-th var
4          cgen(e2, nt+4)
5          sw $a0 4($sp)                # Save v2=val(e2) on stack
6          addiu $sp $sp -4
7      for_loop:
8          lw $a0 i($fp)                # current value of id
9          lw $t1 4($sp)                # value of e2
10         bgt $a0 $t1 for_exit          # Exit if id > v2
11         cgen(e3, nt)                 # Might change id and clobber $a0
12         lw $a0 i($fp)                # Load id again
13         addiu $a0 $a0 1               # Increment it by 1
14         sw $a0 i($fp)                # Save it back
15         b for_loop                   # Jump back to the loop check
16     for_exit:
17         addiu $sp $sp 4               # Remove value of e2 from stack

```

4. Consider the following basic block, in which all variables are integers.

```
1      x := 0 * 5
2      y := a + b
3      z := x * x
4      c := y * x
5      x := x + 4
6      e := c - x
7      x := e * x
8      f := a + b
9      y := y + f
```

(a) Assume that the only variables that are live at the exit of this block are x and y, while a and b are given as inputs. In order, apply the following optimizations to this basic block. Show the result of each transformation. For each optimization, you must continue to apply it until no further applications of that transformation are possible (if any were), before writing out the result and moving on to the next.

- i. Algebraic simplification
- ii. Common sub-expression elimination
- iii. Copy propagation / Constant propagation
- iv. Algebraic simplification
- v. Dead code elimination

**Solution:**

i.

```
1      x := 0
2      y := a + b
3      z := x * x
4      c := y * x
5      x := x + 4
6      e := c - x
7      x := e * x
8      f := a + b
9      y := y + f
```

ii.

```
1      x := 0
2      y := a + b
3      z := x * x
4      c := y * x
5      x := x + 4
6      e := c - x
7      x := e * x
8      f := y
9      y := y + f
```

iii

```
1      x := 0
2      y := a + b
3      z := 0 * 0
4      c := y * 0
```

```

5      x := 0 + 4
6      e := c - x
7      x := e * x
8      f := y
9      y := y + y

```

iv

```

1      x := 0
2      y := a + b
3      z := 0
4      c := 0
5      x := 4
6      e := c - x
7      x := e * x
8      f := y
9      y := y << 1

```

v

```

1      y := a + b
2      c := 0
3      x := 4
4      e := c - x
5      x := e * x
6      y := y << 1

```

- (b) The resulting program is still not optimal. What optimizations, in what order, can you apply to fully optimize the result? Show the maximally optimized codes (with least number of instructions).

**Solution:** You can apply copy propagation, algebraic simplification, and dead code elimination to the output of the previous part to get the following:

```

1      y := a + b
2      x := -16
3      y := y << 1

```

5. Consider the following assembly-like pseudo-code, using 6 temporaries (abstract registers)  $a$  to  $f$ :

```

1      a := b + d
2      b := a + d
3      d := a - b
4      c := d + d
5      if c > 100:
6          c := c + d
7      else:
8          d := 1
9      e := d - c
10     f := e - c

```


- (a) At each program point, list the variables that are live. Note that  $b$  and  $d$  are inputs for the given code and  $f$  is a live value on exit.

**Solution:**

|    |             |        |
|----|-------------|--------|
| 1  |             | {b, d} |
| 2  | a := b + d  | {a, d} |
| 3  | b := a + d  | {a, b} |
| 4  | d := a - b  | {d}    |
| 5  | c := d + d  | {c, d} |
| 6  | if c > 100: |        |
| 7  |             | {c, d} |
| 8  | c := c + d  | {c, d} |
| 9  | else:       |        |
| 10 |             | {c}    |
| 11 | d := 1      | {c, d} |
| 12 | e := d - c  | {c, e} |
| 13 | f := e - c  | {f}    |

- (b) Draw the register interference graph between temporaries in the above program as described in class.

**Solution:**

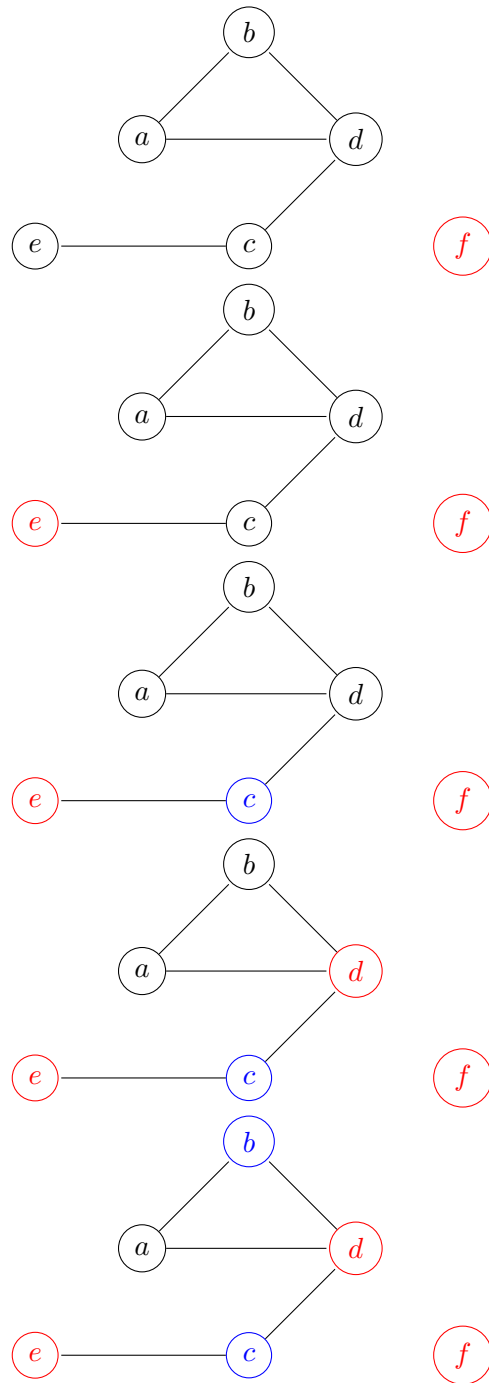



- (c) Provide a lower bound on the number of registers required by the program induced from the interference graph. Can you explain why?

**Solution:** 3. It is mainly because there is a triangle in the graph.

- (d) Using the algorithm described in class, provide a coloring of the graph in part(b). The number of colors used should be your lower bound in part (c). Provide the final k-colored graph (you may use the tikz package to typeset it or simply embed an image), along with the order in which the algorithm colors the nodes.

**Solution:**





- (e) Based on your coloring, write down a mapping from temporaries to registers (labeled r1, r2, etc.).

**Solution:**

|   |       |
|---|-------|
| 1 | a: r1 |
| 2 | b: r2 |
| 3 | c: r2 |
| 4 | d: r3 |
| 5 | e: r3 |
| 6 | f: r3 |