

**Question 1.** (10 points) For each of the following tasks, indicate the earliest stage of the compiler that can always perform that task or detect the situation. Assume the compiler is a conventional one that generates native code (x86, MIPS, etc.) for a language like C++ or Java. Use the following abbreviations to identify the stages:

scan – scanner

parse – parser

sem – semantics/type check

opt – optimization

instr – instruction selection & scheduling

reg – register allocation

run – runtime (i.e., when the compiled code is executed)

can't – can't be done during compilation or execution

scan A comment beginning with `/*` terminates with `*/` before the end of the file

sem A variable is declared at most once in a scope (function, class, etc.)

reg Add code to temporarily store a value when there aren't enough registers to hold all live values at a given point in the program.

opt Replace array subscripting (`a[i]=0; i++`) and a test (`i<n`) with pointer arithmetic (`*p+=0`) and a pointer test (`p<&a[n]`).

sem Warn the programmer that the declaration of a variable in a subclass hides (shadows) the declaration of the same variable in a superclass.

parse Detect that there is an extra `else stmt` clause in the program that is not associated with a previous `if`

run Detect that in an array reference `a[k]`, the subscript `k` is out of bounds (i.e., in Java where bounds must be checked)

can't Guarantee that a particular pointer (reference) variable `p` will never be `NULL`.

sem Guarantee that in a pointer reference `p.x` (`p->x` in C/C++) that the value referenced by `p` will have a field `x`, assuming that `p` is not `NULL`.

instr Use x86 addressing modes to calculate `x*5` with `leal (%eax,%eax,4),%eax`.

**Question 2.** (34 points) Runtime data structures and code. Suppose we have the following two Java classes representing characters in a simulation and a third class with a small main program:

```
public class Character {
    int x, y;    // coordinates
    int age;

    void move(int dx, int dy) {
        x = x + dx;  y = y + dy;
    }
    int getX()    { return x; }
    int getY()    { return y; }
    int getAge() { return age; }
}

public class Elf extends Character {
    int weight;
    int age;

    void setAge(int years) { age = years; }
    int  getAge() { return age; }
}

public class Main {
    public static void main(String[] useless) {
        Character santa = new Character();
        Character bernie = new Elf();

        santa.move(12, 24);
        ((Elf)Bernie).setAge(42);
    }
}
```

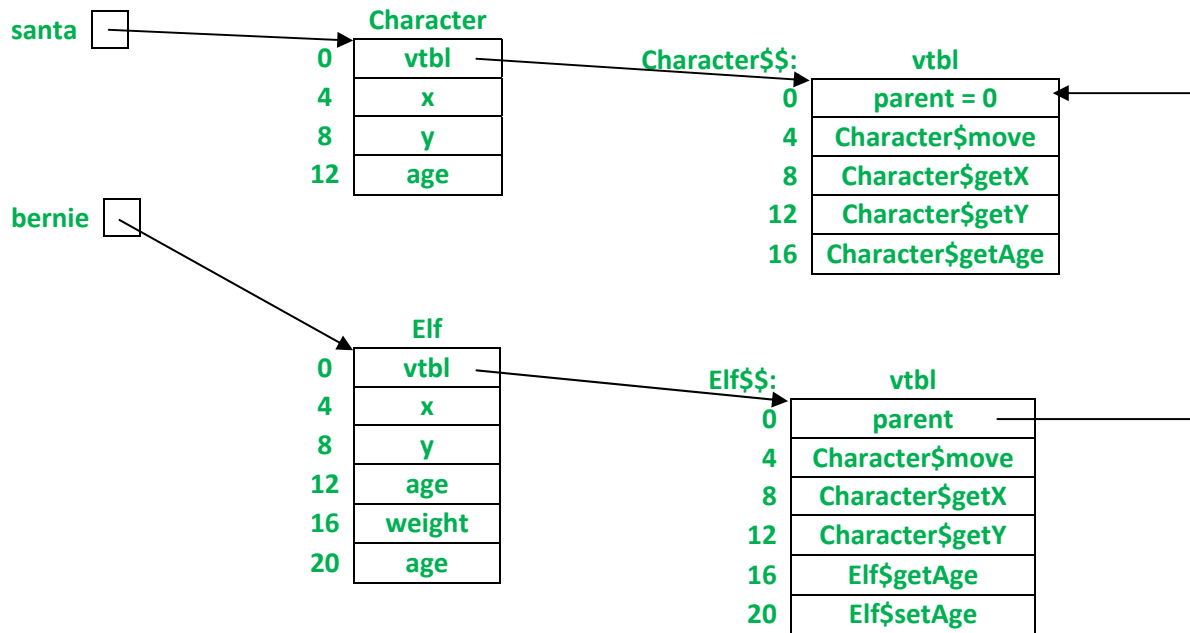
(a) (8 points) Runtime data structures. On the next page draw a diagram of the runtime data structures for this program after the variable declarations in method `main` have been executed.

Show the variables in `main`, the objects they refer to, and how the objects and their fields are organized in memory. Your diagram should show all of the fields in the objects and should assign numeric offsets to each field (0, 4, 8, ...) for use in later parts of this question. You do not need to show the details of `main`'s stack frame. Just show pointer variables that refer to the objects.

Then add to your diagram any mechanisms needed to support dynamic method dispatch in Java (e.g., vtables). You should ignore constructors. The vtable diagram(s) should also show the numeric offsets of the various pointers in the table(s) (0, 4, 8, ...).

(You may remove this page from the exam for reference if you wish.)

**Question 2. (cont).** (a) Draw your diagram below, including variable names, objects, vtbls, field names, and field offsets in objects and vtbls.

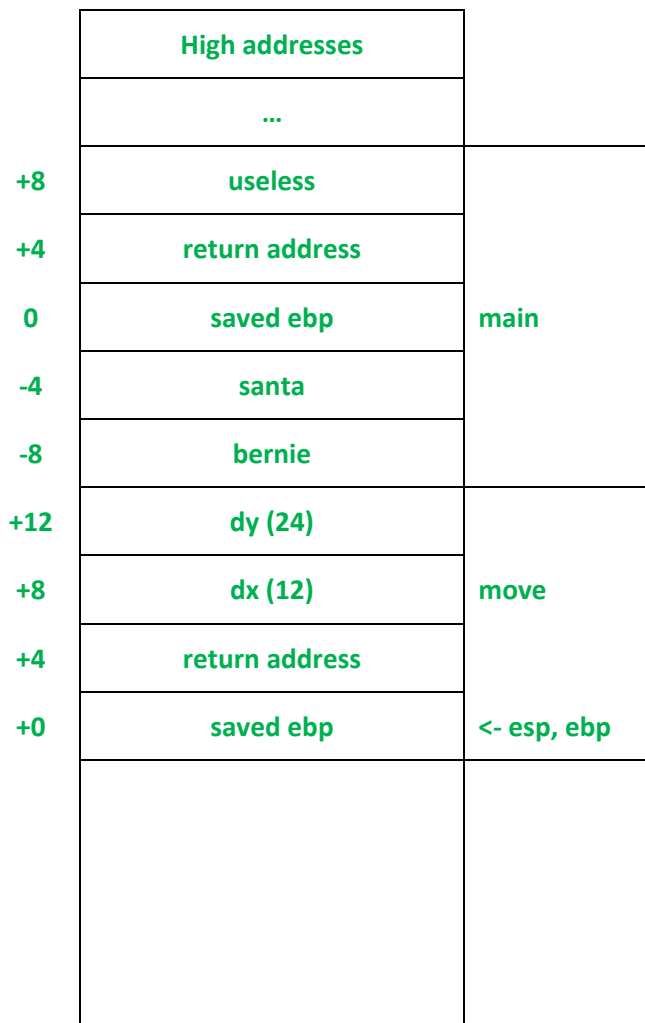


Offsets in the objects and vtables could be different, although most people wound up with these layouts. However, the layout in the common parts of the Character and Elf objects and vtables must match.

Strictly speaking, the vtable entries are pointers to function code stored elsewhere. That's awful tedious to draw using a computer, so labels are used in the above solution, and were ok on the exam as well.

(continued on next page)

**Question 2. (cont.)** (b) (7 points) Stack frames. Suppose now that we start executing method `main`, and continue until we reach the body of method `move` after executing `santa.move(12, 24)`. Draw a picture that shows the stack frames of methods `main` and `move` right after entering `move` and after executing the function prologue (allocating the stack frame and adjusting the base pointer), but before executing the first assignment statement in the body of `move`. Draw labeled arrows showing where the registers `ebp` (frame pointer) and `esp` (stack pointer) point after executing the prologue code in `move`. Your picture should show the variables and parameters of both methods and their numeric offsets in each stack frame. (Notice that there will be two sets of offset numbers – one for the stack frame of `main` and one for the stack frame of `move`.)



**Note:** The parameters to function `move` are not anonymous – they have names as well as values. Many people omitted the names from their solution. We let that go while grading and didn't deduct for it.

**Question 2. (cont.)** (c) (9 points) Method call. Now suppose we continue execution and return to the `main` method after method `move` finishes. The next line of code is `((Elf)bernie).setAge(42)`. Translate this line of code into x86 assembly language as it would be compiled in `main`. You should assume that the stack frame for method `main` is organized as you've drawn it in the previous part of this question, and that the object and vtable layouts are as drawn in part (a). You should use the standard calling conventions described in class, and use register `ecx` as the "this" pointer. You may assume that no code is generated to check the `(Elf)` cast at runtime. Any reasonable x86 code that obeys the regular calling conventions is ok – you don't need to reproduce the code that would be generated by your compiler. You should use the GNU/AT&T x86 assembly language for your code.

```

movl    -8(%ebp),%ecx    move obj ptr to ecx
pushl   $42             push argument
movl    (%ecx),%eax     copy vtbl ptr to eax
call    *20(%eax)      indirect call to setAge
addl    $4,%esp        pop arguments (popl ok also)

```

(d) (10 points) Method implementation. Last part(!). Give a x86 translation of method `setAge` from class `Elf`. Your code should include all of the code for method `setAge`, including the function prologue and return. Since this is a void method there is no return value required to be placed in `eax` when the function exits. Include the usual method prologue and epilogue code to save/restore the frame pointer even though there may be other ways to write the code without it.

`Elf$setAge:`

```

pushl   %ebp           usual prologue
movl    %esp,%ebp     (no locals needed)
movl    8(%ebp),%eax   copy argument to eax
movl    %eax,20(%ecx) copy to object
movl    %ebp,%esp     return
popl    %ebp
ret

```

The final `movl/popl` could be replaced by a `leave` instruction.

Several solutions tried to use a single `movl 8(%ebp),20(%ecx)` instruction to copy the data. That can't be done since a single x86 instruction can only have one address specification.

**Question 3.** (32 points) Compiler hacking: another question of many parts.

We would like to add a new kind of counting loop to our MiniJava compiler. The syntax is

```
for id in exp1 to exp2 do stmt
```

The idea is that *stmt* is to be executed repeatedly, with variable *id* having successive values *exp1*, *exp1+1*, *exp1+2*, ..., *exp2* each time that *stmt* is executed. More specifically the loop is executed as follows. First, *id* is assigned the value *exp1*. Then the value of *id* is tested to see if it is less than or equal to *exp2*. If it is, the loop body, *stmt*, is executed and then the value of *id* is incremented by 1. This test-execute-increment sequence is repeated until the test becomes false because the value of *id* exceeds *exp2*.

The variable *id* must be declared previously in the function and must have type **int**. It must be a local variable in the function containing the loop; it may not be a global or class instance variable. The two expressions *exp1* and *exp2* are evaluated only once before the initial assignment to *id* and are not reevaluated during execution of the loop. Further, the loop body *stmt* may not contain any assignments to the variable *id*. (In other words, having evaluated *exp1* and *exp2* at the start of the loop, we can tell exactly how many times the loop will execute and the sequence of values that *id* will assume. Execution of the loop body is not allowed to change that sequence – although, of course, it could return from the function, throw an exception, or otherwise terminate abnormally.)

Answer the rest of this question on the next few pages. You can remove this page and use it for reference as you work on the question. You may also find the other pages handed out along with the test containing the MiniJava grammar and some of the AST class declarations to be useful.

Write

Your

Answers

On

The

Next

Few

Pages

(continued next page)

**Question 3. (cont.)** (a) (3 points) What new tokens would need to be added to the scanner and parser of our MiniJava compiler to add the new `for` statement to the original MiniJava grammar? Just list the tokens; you don't need to give a JFlex or CUP specification for them.

**FOR IN TO DO**

**(FOR is not a token in regular minijava)**

(b) (7 points) Complete the following class to define a new AST node class for this new `for` statement. You only need to define instance variables and the constructor. Assume that all appropriate package and import declarations are supplied, and don't worry about visitor code.

```
public class ForLoop extends Statement {  
    // add instance variables below
```

```
    Identifier id;  
    Exp exp1, exp2  
    Statement stmt;
```

```
// constructor - add parameters and body
```

```
public ForLoop ( Identifier id, Exp exp1, Exp exp2, Statement stmt,  
                int line_nbr_ ) {
```

```
    super(line_nbr);  
  
    this.id = id;  
    this.exp1 = exp1;  
    this.exp2 = exp2;  
    this.stmt = stmt;
```

```
}
```

```
}
```

**Question 3. (cont.)** (c) (6 points) Complete the CUP specification below to define a new production for the `for` statement and the associated semantic action(s) needed to parse `for` and insert an appropriate `ForLoop` node (as defined in part (b)) into the AST. We have added the necessary additional code to the parser rule for `Statement` as shown below.

```
Statement ::= ...
           | ForStatement:s { : RESULT = s; : }
           ...
           ;
```

```
ForStatement ::= FOR Identifier:id IN Exp:exp1 TO Exp:exp2
                DO Statement:stmt

                { : RESULT = new ForLoop(id,exp1,exp2,stmt,idleft); : }
```

(d) (6 points) Describe the checks that would be needed in the semantics/type-checking part of the compiler to verify that a `for` statement is legal. You do not need to give code for a visitor method or anything like that – just describe what rules (if any) need to be checked.

- Verify `id` has type `int`
- Verify `id` is declared locally in the function scope
- Verify that `exp1` and `exp2` have type `int`
- Verify that `stmt` contains no assignments to `id`



**Question 3. (cont.)** (e) (10 points) Describe the code that would be generated for the new loop statement. You need to show the instructions, labels, and any other assembly language code that need to be generated for the `loop` statement itself, and show where the generated code for `exp1`, `exp2`, the assignment(s) to `id`, and `stmt` would appear in the code sequence for this new loop. If you need to create any temporary variables or otherwise save values somewhere, be sure it is clear what you are doing and where the values are stored.

**There are obviously many solutions. This one is simple with no optimizations.**

**We need to allocate temporary space for `exp1` while evaluating `exp2`, and we need to allocate a temporary to hold `exp2` throughout the loop. Here we push `exp1` on the stack and assume the compiler has allocated an anonymous local variable to hold `exp2`.**


**(Note that `exp1` and `exp2` need to be evaluated in order and before the assignment to the variable to get something like `for i in i+1 to i+10 <stmt>` to execute properly. We weren't quite so strict in the grading if you didn't get this exactly right. Also note that the identifier is a regular variable and its current value needs to be stored properly since it may be referenced by code in the loop body. It cannot remain in a register throughout the loop. Also, remember that the code in the loop body may use registers, so the values of the variable or `exp2` cannot be stored in, for example, `edx`.)**

```
<code for exp1>
pushl %eax                save exp1 on stack
<code for exp2>
movl %eax,offset_tmp(%ebp) store exp2 in compiler temporary
popl %eax                 assign exp1 to variable
movl %eax,offset_id(%ebp)
loop_label:
movl offset_id(%ebp),%eax  compare id to exp2
cmpl %eax,offset_tmp(%ebp)
jl loop_exit              exit if tmp2 less than id
<code for stmt>
incl offset_id(%ebp)      id++
j loop_label
loop_exit:
```

**Question 4.** (12 points) Register allocation. Considering the following C function:

```
int f(int a, int b) {  
    int k, n;  
    k = a;  
    n = 0;  
    while (k < b) {  
        n = n + k;  
        k = k + 1;  
    }  
    return n;  
}
```

(a) Draw the interference graph for the variables and parameters of this function. You are not required to draw the control flow graph, but it could be useful to sketch it out to help with the solution and to leave clues about what might have happened if the graph is not quite correct.



(b) Give an assignment of (groups of) variables to registers using the minimum number of registers possible, based on the information in the interference graph. You do not need to go through the steps of the graph coloring algorithm explicitly, although it may be helpful as a guide to assigning registers. If there is more than one possible answer that uses the minimum number of registers, any of them will be fine. Use R1, R2, R3, ... for the register names.

**Three registers are needed. b, n, and k each are in separate registers. a can be allocated in the same register as either k or n.**

**Question 5.** (12 points) A little optimization. For this question we'd like to perform local constant propagation and folding, followed by dead code elimination. We have a C function with the following array declaration:

```
double A[10][10];    // 10x10 array of 8-byte doubles
```

The first column of the table below gives the three-address code generated for this assignment statement:

```
a[3][5] = a[3][5] + x;
```

(a) Fill in the second column with the statements from the first column after they have been modified by local constant propagation and folding (compile-time arithmetic).

(b) In the third column, check the box "deleted" if the statement would be deleted by dead code elimination after performing the constant propagation/folding optimizations from part (a).

Original Code	After constant propagation & folding	"X" if deleted as dead code
t1 = *(fp+xoffset) // x	t1 = *(fp+xoffset)	
t2 = 3	t2 = 3	X
t3 = t2 * 10	t3 = 30	X
t4 = t3 + 5	t4 = 35	X
t5 = t4 * 8	t5 = 280	X
t6 = fp + t5	t6 = fp + 280	
t7 = *(t6+aoffset) // A[3][5]	t7 = *(t6+aoffset)	
t8 = t1 + t7 // A[3][5]+x	t8 = t1 + t7	
t9 = 3	t9 = 3	X
t10 = t9 * 10	t10 = 30	X
t11 = t10 + 5	t11 = 35	X
t12 = t11 * 8	t12 = 280	X
t13 = fp + t12	t13 = fp + 280	
*(t13+aoffset) = t8	*(t13+aoffset) = t8	