**Ouestion 1. (10 points)** Compiler Passes & Optimization – An array of questions. (a) (6 points) Consider the following code snippets, each of which defines and assigns an array. For each snippet, identify *which compiler pass/phase* will detect an error (if any) Assume that all code is MiniJava! (not standard Java, nor some other language) Please use the following abbreviations: (n.b. not all of these will be used) scan – scanner parse – parser check - checker mid – compiler middle (translating from ASTs to 3AC, SSA, optimization passes, dataflow) back – compiler backend (instruction selection, scheduling & register allocation) run – runtime (i.e., when the compiled code is executed) none – there is no error check int[] arr; arr = new int[false]; int[] arr; arr = new int [1 % 5]; scan int[] arr; arr = new int [1]; none parse int[] arr; arr = new boolean[1]; run int[] arr; arr = new int[0-1]; parse int[] arr; arr = new int[1.5];

(b) (4 points) Consider the following MiniJava class in isolation.

```
class Foo {
    int[] arr;
    public int get_init() {
        arr = new int[1];
        arr = new int[2];
        return arr;
    }
}
```

Which optimization can a compiler safely make to this code? And which analysis enables that optimization?

The first (`arr = new int[1];`) assignment can be eliminated via Dead Code Elimination. There are two valid answers for which analysis enables the optimization. The first answer is "liveness analysis." Liveness analysis will reveal that `arr` is dead in-between the two assignments. The second answer is SSA. Converting this code into SSA can reveal that  $arr_1$  (unlike  $arr_2$ ) is never used, and hence the first assignment is dead code.

Question 2. (14 points) Scoping, and VTables — I Object! It's not as easy as one, two, ...

(a) (6 points) Scope & Dispatch. Here is a peculiar Java (not MiniJava) program consisting of a main class and two subclasses.

```
public class Main {
public class one {
                                    static void main(String[] args) {
    int one;
                                      one one = new one();
    int two;
                                      one.init();
    void init() {
                                      System.out.println(one.one(1));
        one = 1;
                                      System.out.println(one.two(2));
        two = 2;
    }
                                      two two = new two();
    int one(int one) {
                                      one = two;
        one = two;
                                      one.init();
        two = one;
                                      System.out.println(one.one(1));
        return 1;
                                      System.out.println(one.two(2));
    }
                                      System.out.println(two.three(3));
    int two(int two) {
                                    }
        this.one = this.two;
                                  }
        this.one(two);
        return one;
    }
}
public class two extends one {
    int two;
    int one(int one) {
        this.two = one;
        this.one = one + two;
        return one;
    }
    int three(int three) {
        return one + two + three;
    }
```

What output does this program produce when we compile it and then execute the main method in class Main? (The program does compile and execute without errors.) (If you wish, you may supply additional information about the state of the objects after each println call. We may award partial credit on that basis, and it may help you keep everything straight.)

```
1 (one -- { one = 1; two = 2; })
2 (one -- { one = 2; two = 2; })
1 (two -- { one = 2; one.two = 2; two.two = 1; })
4 (two -- { one = 4; one.two = 2; two.two = 2; })
9
```

}

#### Question 2. (cont.)

(b) (4 points) VTables. When the class `one` was compiled, the compiler picked the following vtable layout for the class.

	<u>vtable layout</u>	<u>offset</u>	
one\$\$:	.quad 0	<pre># no superclass</pre>	
	.quad one\$init	# +8	
	.quad one\$one	# +8	
	.quad one\$two	# +16	
Note: We goofed	and forgot to include	one\$init. Thus, the	
following answer	r (prepared before obse	erving this bug) is	
acceptable, as i	is any answer that cor	rects the above table's	
offsets and then	n produces an accurate	corresponding table below.	

Below, show an appropriate Vtable layout for class `two`, in the same format used above for class `one`. Be sure to properly account for inherited methods in the Vtable layout.

	<u>vtable layout</u>	<u>offset</u>
two\$\$:	.quad one\$\$	<pre># superclass</pre>
	.quad two\$one	# +8
	.quad one\$two	# +16
	.quad two\$three	# +24

(c) (4 points) Object Instance Layout. Assuming we have the same memory management scheme as MiniJava, how many bytes of heap memory would each instance of a `two` take up?

32 bytes (8 for vtable pointer, 8 for one.one, 8 for one.two, 8 for two.two)

Question 3. (14 points) A bit of x86-64 coding – Perfectly Average.

The C function below calculates the signed difference between the mean and median of an integer array. For simplicity of implementation, we assume the input array is of odd length and is sorted.

```
/* Assume arr is sorted and odd-length */
double med_mean_difference(int[] arr, int length) {
    int middle = length / 2;
    int median = arr[middle];
    return median - mean(arr, length);
}
int mean(int[] arr, length) {
    // returns the closest int to the mean of all numbers in arr
}
```

Reference and ground rules for x86-64 code, (same as for the MiniJava project and other x86-64 code):

- All values, including pointers and ints, are 64 bits (8 bytes) each, as in MiniJava
- You must use the Linux/gcc assembly language, and must follow the x86-64 function call, register, and stack frame conventions:
  - Argument registers: %rdi, %rsi, %rdx, %rcx, %r8, %r9 in that order
  - Called function must save and restore %rbx, %rbp, and %r12-%r15 if these are used in the function
  - Function result returned in %rax
  - $\circ$  %rsp must be aligned on a 16-byte boundary when a call instruction is executed
  - %rbp must be used as the base pointer (frame pointer) register for this question
- The full form of a memory address is constant(%rbase,%rindex,scalefactor), which references memory address %rbase+%rindex\*scalefactor+constant. scalefactor must be 0, 2, 4, or 8.
- Please assume that the length of the array is odd, and its contents are sorted.
- This is simple C code, not a Java method, so there is no this pointer.
- The mean() function and set up a stack frame consistent with local variables declared in the med\_mean\_difference() function.
- Rather than trying to remember the sign division stuff from class, please use **divq a,b** for your answer, where **divq a,b** computes and places a divided by b in %rax with the remainder placed in %rdx.
- You do not need to mimic the code produced by your MiniJava compiler.
- Please include *brief* comments in your code to help us understand what the code is supposed to be doing (which will help us assign partial credit if it doesn't do exactly what you intended.)

Question 3. (cont.) Write your x86-64 translation of function med\_mean\_difference() below. Remember to read and follow the above ground rules carefully, including managing registers properly and using the correct argument registers for function calls, and creating a local stack frame to hold the local variables correctly while calling mean(). Your code should include a translation of all of the code in the original function. Brief comments are appreciated. Original code repeated below for convenience:

```
/* Assume arr is sorted and odd-length */
double med mean difference(int[] arr, int length) {
    int middle = length / 2;
    int median = arr[middle];
   return median - mean(arr, length);
}
int mean(int[] arr, length) {
   // returns the closest int to the mean of all numbers in arr
}
med mean difference:
  pushq %rbp
                                 # method prologue
  movq %rsp, %rbp
  subq $16, %rsp
                                # allocate space and align (mult.
of 16)
  movq $2, %rax
divq %rsi, %rax
                                # move 2 into %rax
                        # move 2 into Arax
# divide length by 2, result
in %rax
  movq 0(%rdi, %rax, 8), %rax # obtain median, result in %rax
  movq %rax, -8(%rbp)
                                # stores result on stack
                   # note: %rdi and %rsi are unchanged here
                                # calls mean function
  call mean
  subq -8(%rbp), %rax
                               # subtracts with result in %rax
                            # correct sign of subtraction
# method epilogue
  mulq $-1, %rax
  movq %rbp, %rsp
  popq %rbp
  ret
```

There are many other ways to write the function as well.

- Regardless, one must correctly set up and tear down the stack frame, safely saving values while the mean() subroutine is called.
- The division can alternately be performed using movq %rsi, %rax; shrq \$1, %rax
- The median can be stored in %rbx, or %r12-%r15 instead of placing it on the stack, in which case no stack frame needs to be allocated. Doing this can remove a lot of instructions.

**Question 4. (24 points)** Add a Feature – I am the Walrus, Coo coo ca choo! You have just launched the first version of your MiniJava compiler, and it was a huge success! Users are already wanting more features. One user, who is an avid Python programmer, is requesting an "assignment expression" construct using the walrus operator (:=).

The semantics of this operator are as follows for an expression of the form

#### <variable> := <expression>

- 1. The *<expression>* on the right hand side is evaluated.
- 2. The value of the *<expression>* is assigned to the *<variable>* on the left hand side.
- 3. The overall assignment expression returns the value of the RHS *<expression>*.

Your client is only requesting the assignment expression to offer the same level of support as standard assignment, so array assignment doesn't need to be handled.

Answer the questions below about how this new expression operator would be added to a MiniJava compiler. There is likely way more space than you will need for some of the answers. The full MiniJava grammar is attached at the end of the exam if you need to refer to it.

(a) (2 points) What new lexical tokens, if any, need to be added to the scanner and parser of our MiniJava compiler to add assignment expressions to the original MiniJava language? Just describe any necessary changes and new token(s) needed and their name(s). You don't need to give JFlex or CUP specifications or code in this part of the question, but you will need to use any token name(s) you write here in a later part of this question.

We need a new token WALRUS for the `:=` operator. Other names for the token are fine.

(Note that one should ideally create a new 2-character token. It would also be possible to create a new 1-character token for `:`. We did not deduct for this approach if it is consistently followed.)

(continued on next page)

**Question 4. (cont.) (b) (6 points)** Complete the following new AST class to define an AST node type for this new expression assignment. 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.

(Hint: recall that the AST package in MiniJava contains the following key classes: ASTNode, Exp extends ASTNode, and Statement extends ASTNode. Also remember that each AST node constructor has a Location parameter, and the supplied super(pos); statement at the beginning of the constructor below is used to properly initialize the superclass with this information.)

```
public class AssignExp extends Exp {
    // add any needed instance variables below
    public Identifier i;
    public Exp e;
```

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

```
public AssignExp (<u>Identifier i, Exp e, Location pos</u>) }
super(pos); // initialize location info in superclass
this.i = i;
this.e = e;
}
```

```
(continued on next page)
```

}

Question 4. (cont.) (c) (5 points) Complete the CUP specification below to define a production for this new assignment expression, including associated semantic action(s) needed to parse the new expression and create an appropriate AST node (as defined in part (b) above). You should use any new lexical tokens defined in your answer to part (a) as needed. Use reasonable names for any other lexical tokens that already would exist in the compiler scanner and parser if you need them. We have added additional code to the parser rule for Expression below so the CUP specification for the new expression assignment can be written as an independent grammar rule with separate parser actions.

Hint: recall that the Location of an item foo in a CUP grammar production can be referenced as fooxleft.

```
Expression ::= ...
| AssignExp:e {: RESULT = e; :}
...
;
AssignExp ::= Identifier:i WALRUS Expression:e
{: RESULT = new AssignExp(i, e, ixleft); :}
```

(d) (4 points) Describe the checks that would be needed in the checking part of the compiler to verify that a program containing this new assignment expression is legal. You do not need to give code for a visitor method or anything like that – just describe what language rules (if any) need to be checked for this new statement to verify it is used correctly.

- 1. `i` is declared and in scope
- 2. `i` is assignment compatible with the return type of `e`, i.e. it has the same type as the return type of `e` or `i` has a supertype.
- 3. Output The computed return type of the assignment expression should be the same as the return type of `e`. (alt. it could be the same as the type of `i`. In the case that the assignment involves a cast these may be different types. Either specification is reasonable)

(continued on next page)

Question 4. (cont.) (e) (7 points) Describe the x86-64 code shape for this new assignment expression, as it would be generated by a MiniJava compiler. Your answer should be similar in format to the descriptions we used in class for other language constructs. If needed, you should assume that the code generated for an expression will leave the value of that expression in %rax, as in our MiniJava project.

For example, you can use  $\langle ... \rangle$  as placeholders for code generated by a child expression. If needed, you should assume that the code generated for a child expression will leave the resulting value of that expression in %rax, as in our MiniJava project.

#### You may assume that the variable you are assigning to is a method local variable.

Use Linux/gcc x86-64 instructions and assembler syntax when needed. If you need to make any additional assumptions about code generated by the rest of the compiler you should state them.

<Generate code to evaluate e and leave the result in %rax>
movq %rax offset<sub>i</sub>(%rbp)
# This is pretty much the same as regular assignment

While we studied program analysis in order to justify the safety of compiler optimizations, analyses are also very useful for detecting and warning users about likely bugs. For the next question we will use dataflow to try to detect unclosed files.

The statements in the program include those that open files and those that close files. We would like to ensure that all files are closed before the end of execution. We want to use dataflow analysis to discover if there are any unclosed files by keeping track of which files may be open at various points in the program.

For the next two problems, assume that our language has two types: file objects and strings. The language has the following operations:

- f = open(path): opens the given file and assigns it to f (f must be file object type)
- close(f): closes the given file (f must be file object type)
- x = read(f): converts the contents of file f into a string and assigns it to x (f must be a file object type and x must be a string type)
- write(f, x): writes the string x onto at the end of file f (f must be a file object type and x must be a string type)

The following two problems refer to this dataflow graph, which uses the above operations:



**Question 5. (16 points)** Dataflow analysis – files open/closed. We would like to ensure that every file which is opened will eventually be closed. We can use a dataflow framework to analyze which variables refer to open or closed files by defining the following sets for each basic block b:

- IN(b) the set of file variables that are known to be open on entry to block b
- OUT(b) the set of file variables that are known to be open on exit from block b
- GEN(b) the set of all file variables that are opened in block b and not later closed in block b before exit from that block
- KILL(b) the set of all file variables that are closed in block b and not later opened in block b before exit from that block.

The following dataflow equations describe the relationships between these sets:

 $IN(b) = \bigcup_{x \in pred(b)} OUT(x)$  $OUT(b) = GEN(b) \cup (IN(b) - KILL(b))$ 

(a) (14 points) Complete the following table using iterative dataflow analysis to identify the open file variables in the IN and OUT sets for each block in the above flow graph. You should first fill in the GEN and KILL sets for each block (which do not depend on other blocks) and then iteratively solve for IN and OUT.

	GEN	KILL	IN	OUT	IN	OUT	IN	OUT
B0	а			а		а		
B1		е	а	а	a, b	a, b	No changes	
B2	b	а	а	b	a, b	b		
В3		e, b	a, b	а	a, b	а		

(b) (2 points) Is it possible that there is an unclosed file at the end of execution? If so, which file(s) might be open? Answer using the control flow diagram and the information about the IN and OUT sets calculated above.

Yes. The OUT set of the final block B3, contains file variable a.

**Question 6. (18 points)** Dominators and SSA. Here are the basic definitions of dominators and related concepts we have seen previously in class:

- Every control flow graph has a unique **start node** s.
- Node *x* **dominates** node *y* if every path from s to *y* must go through x.
  - A node x dominates itself.
- A node *x* strictly dominates node *y* if *x* dominates *y* and  $x \neq y$ .
- The **dominator set** of a node *x* is the set of nodes *dominated by x*.
  - $|\operatorname{Dom}(\mathbf{x})| \ge 1$
  - (note: sometimes the definition of Dom(x) is given as the set of all nodes that dominate x. For SSA it is more convenient to keep track of the set of nodes that x dominates.)
- An **immediate dominator** of a node *y*, idom(*y*), has the following properties:
  - idom(y) strictly dominates y
     (i.e., dominates y but is different from y)
  - idom(y) does not dominate any other strict dominator of y
- The **dominator tree** of a control flow graph is a tree where there is an edge from every node *x* to its immediate dominator idom(*x*).
- The **dominance frontier** of a node *x* is the set of all nodes *w* such that
  - x dominates a predecessor of w, but
  - *x* does not strictly dominate *w*



(a) (8 points) Using the same control flow graph from the previous problem, complete the following table. List for each node: the node(s) that it dominates, successor nodes to the dominated nodes, and the nodes that are in its dominance frontier (if any):

Node	Nodes dominated by this node	Successor(s) to nodes dominated by this node	Dominance Frontier of this node
В0	B0, B1, B2, B3	B1, B2, B3	
B1	B1	B3	B3
B2	B2	B1, B2, B3	B1,B2,B3
В3	B3		

Question 6. (cont.) (b) (10 points) Now redraw the flowgraph in SSA (static single-assignment) form. You need to insert all  $\Phi$ -functions that are required by the dominance frontier criteria, even if some of the variables created by those functions are not used later. Once that is done, add appropriate version numbers to all variables that are assigned in the flowgraph. You do not need to trace the steps of any particular algorithm to place the  $\Phi$ -functions as long as you add them to the flowgraph in appropriate places. Answers that have a couple of extraneous



 $\Phi$ -functions will receive appropriate partial credit, but answers that, for example, use a maximal-SSA strategy of placing  $\Phi$ -functions for all variables at the beginning of every block will not be looked on with favor. **Solution diagram** 



Question 7. (14 points) Register allocation/graph coloring.

(a) (8 points) Draw the interference graph for the temporary variables (t1-t10) in the following code. You should assume that all temporaries are dead at the conclusion of this snippet of code

```
// code for z = *((v+w) * (w*x)) * (w*x)^2 + y
a. LOAD
          t1 <- v
                       // t1 = v
b. LOAD
          t2 <- w
                       // t2 = w
c. ADD
         t3 <- t1, t2 // t3 = v+w
d. MULT
         t4 <- t2, x // t4 = w^*x
e. MULT
         t5 <- t3, t4 // t5 = (v+w) * (w*x)
f. LOAD
         t6 <- y
                   // t6 = y
g. LOAD
         t7 < -MEM[t5] // t7 = *((v+w) * (w*x))
h. SHIFT t8 <- t4, $2 // t8 = (w*x)^2
         t9 <- t7, t8 // t9 = *((v+w) * (w*x)) * (w*x)^2
i. MULT
j. ADD
         t10 <- t9, t6 // t10 = *((v+w) * (w*x)) * (w*x)^2 + y
k. STORE
           z <- t10 // store z
```



(b) (6 points) Give an assignment of groups of temporary variables to registers that uses the minimum number of registers possible based on the information in the interference graph. Use R1, R2, R3, ... for the register names.

As usual, there are many possible answers. Here is one:

R1: t1, t3, t5, t7, t9, t10 R2: t2, t4, t8 R3: t6

#### Question 8. (2 free points – all answers get the free points)

Draw a picture of something you are planning to during summer break! - or -Draw a picture of something you think one or more of your TAs will do during summer break!

A very common answer....



Have a great summer break and best wishes for the future! The CSE 401/M501 staff