

CSE P 501 19au Exam 11/21/19

Name _____ UW ID # _____

There are 8 questions worth a total of 125 points. Please budget your time so you get to all of the questions. Keep your answers brief and to the point.

The exam is closed books, closed notes, closed electronics. Please turn off all cell phones, personal electronics, alarm watches, and pagers, and return your tray tables and seat backs to their full upright, locked positions. Sound recording and the taking of photographs is prohibited.

If you have a question during the exam, please raise your hand and someone will come to help you.

There is an extra blank page at the end of the exam you can use if your answer(s) do not fit in the space provided. Please indicate on the original page(s) if your answer(s) is(are) continued on that last page.

Following the blank page for additional answers is a single page that contains a copy of the MiniJava grammar and a summary of x86-64 information that may be useful in answering some of the questions. Feel free to remove that page to use while you are working.

Please wait to turn the page until everyone is told to begin.

Score _____

1 _____ / 10

2 _____ / 12

3 _____ / 24

4 _____ / 12

5 _____ / 23

6 _____ / 14

7 _____ / 14

8 _____ / 16

CSE P 501 19au Exam 11/21/19

Question 1. (10 points, 1 each) Compiler phases. There are many possible errors that can occur in a program. For each of the following possible MiniJava errors, indicate when it would be detected, either at compile time or during execution, and, if the error can be detected by the compiler, indicate the earliest point in the compiler (scanner, parser, typecheck/semantics) where the error can definitely be detected. Assume that the compiler is a conventional one that generates native code for a single target machine (say, x86-64), and assume that the source language is the MiniJava for our project (if it helps, a copy of the MiniJava grammar is attached at the end of the exam for reference). Use the following abbreviations for the stages:

scan – scanner

parse – parser

sem – semantics/type check

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

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

_____ Standard MiniJava does not include a `>>` (right shift) operator

_____ In the assignment statement `a=b;`, expression `b` has type `int` and `a` has type `boolean` and these types are incompatible for assignment.

_____ The program contains an infinite loop and will not terminate if it is executed.

_____ Standard MiniJava does not include a `++` “increment” operator.

_____ If `x` has type `T` and the program contains a method call `x.f(17)`, there is no method `f` that has one integer parameter in class `T` or any of its superclasses.

_____ In the method call `x.f(17)`, variable `x` is `null` (does not reference an object)

_____ Class `C` contains two definitions for method `f` that have different numbers of parameters (recall that MiniJava does not support method overloading).

_____ The array reference `a[5]` is incorrect because `a` has too few elements.

_____ `System.out.println(x<y)` is not supported in MiniJava.

_____ `System.out.println()` is not supported in MiniJava.

CSE P 501 19au Exam 11/21/19

Question 2. (12 points) Regular expressions. All reasonable file systems have directories and files, where directories can contain other directories and files, and there is a naming convention for referencing files.

We have a Unix/Linux-like operating system where file names are specified as follows (read carefully – this might be a *tiny* bit different from what you might assume). A name can start with an optional leading slash ('/'). Then there is a set of zero or more directory names separated by slashes. Directory names consist of one or more upper- or lower-case letters. Following the directory names, if any, is the required file name. File names include an identifier with one or more upper- or lower-case letters, followed by an optional one- to three-letter file extension. If the file extension is present, it is separated from the file name by a period ('.').

Examples of valid names: `/long/directory/path/with/a/file.txt;`
`some/subDirectory/lib.a;` `/x.yz;` `simpleFileNameWithNoExtension;`
`simpleFileNameWithExtension.cpp`

Examples of invalid names: `a/b//c` (adjacent '/' characters), `a/b/` (no file name after final '/'), `foo.docx` (more than three letters following '.'), `file1` (digit in file name) `file.` (no letters following '.')

(a) (6 points) Give a regular expression that generates strings representing file names as described above. (Hint: you may want to work on parts (a) and (b) at the same time.)

Ground rules (the fine print): You may only use the basic regular expression operations of concatenation, choice (`|`), and repetition (`*`) plus the derived operators `?` and `+`, and simple character classes like `[abc0-9]` and `[^a-z]`. You may use abbreviations like `vowels = [aeiou]`. You may not use more complex operators found in various software tools that handle extended regular expressions and you should not use `\` or other escape characters. If you need to differentiate between terminal characters and regular expression operators, underline the terminal characters to distinguish them or do something equally simple and easy to read.

(continued on next page)

CSE P 501 19au Exam 11/21/19

Question 2. (cont.) (b) (6 points) Draw a DFA that accepts file names as defined in the question and that are generated by your answer to part (a).

CSE P 501 19au Exam 11/21/19

Question 3. (24 points) The you're-probably-not-surprised-to-see-it LR parsing question. In C, C++, Java, and other languages, ++ can be used as both a prefix and postfix increment operator. Consider the following grammar, which uses the terminal `inc` instead of ++ to avoid any possible confusion about whether it is a single terminal symbol (which it is), rather than a pair of separate + symbols (which it isn't).

0. $exp' ::= exp \$$ (\$ is end-of-file)
1. $exp ::= inc\ exp$
2. $exp ::= exp\ inc$
3. $exp ::= x$

(a) (14 points) Draw the LR(0) state machine for this grammar. You do not need to include the parse table with shift/reduce and goto actions, although you can write all or part of that later if you find it helpful to answer the rest of this question on the next page.

(continued on next page)

CSE P 501 19au Exam 11/21/19

Question 3. (cont.) Grammar repeated for reference

- 0. $exp' ::= exp \$$ (\$ is end-of-file)
- 1. $exp ::= inc\ exp$
- 2. $exp ::= exp\ inc$
- 3. $exp ::= x$

(b) (4 points) Compute *nullable* and the FIRST and FOLLOW sets for the single nonterminal *exp* in the above grammar:

Symbol	nullable	FIRST	FOLLOW
<i>exp</i>			

(c) (3 points) Is this grammar LR(0)? Why or why not?

(d) (3 points) Is this grammar SLR? Why or why not?

CSE P 501 19au Exam 11/21/19

Question 4. (12 points) (LL parsing/grammars) Consider the following grammar:

$$S ::= b \mid M$$
$$M ::= a \mid M c \mid c M \mid b b$$

Is this a LL(1) grammar suitable for top-down predictive parsing? If yes, give a specific technical justification for your answer. If not, give a quick explanation of why it is not, and then give a grammar that generates the same language and is LL(1) if that is possible. If no LL(1) grammar can generate the same language produced by the original grammar, give an explanation of why this is not possible.

Hint: You are not required to compute FIRST and FOLLOW information for the nonterminals in the grammar, but you might find some of that information useful in explaining your answer.

CSE P 501 19au Exam 11/21/19

Question 5. (23 points) Compiler hacking: MiniJava++. We would like to add a new ++ pre-increment operator to MiniJava. (A copy of the MiniJava grammar is included at the end of the test for reference if needed.) The new ++ operator is a single token and is a *prefix* operator. We will add the following new rule to the MiniJava grammar:

Expression ::= “++” Expression

The new ++ operator has the same semantics as it does in C/C++/Java and similar languages. The Expression is evaluated, and it must be a l-value, i.e., an assignable location designating a variable, array element, or something similar. The ++ operator increments the contents of that location, and then the returned (computed) value of ‘++Expression’ is the newly incremented value. Since we are extending MiniJava, the operand of ++ must designate a location containing an integer, since those are the only values in the language that can be incremented.

For example, suppose we have the following two statements in a MiniJava program:

```
x = 7;  
System.out.println(++x);
```

The `System.out.println` statement should print 8, and the final value of `x` is 8.

Note: for this problem, we are only adding a prefix ++ operator, i.e., for expressions like ++x. We are *not* also including a postfix x++ operator.

(a) (3 points) What new lexical tokens and/or keywords would need to be added to the scanner and parser of our MiniJava compiler to add this new ++ operator to the original MiniJava grammar? Just list the new tokens, if any; you don’t need to give JFlex or CUP specifications for them.

(continued on next page)

CSE P 501 19au Exam 11/21/19

Question 5. (cont.) (c) (5 points) Complete the CUP specification below to add an appropriate grammar rule and semantic actions for the new ++ operator, including creating a new Increment node (as defined in part (b) above). Your CUP grammar rule should use any new tokens or keywords defined in part (a) as appropriate. We have included the production for an Identifier expression (as found in the project starter code) as a reminder; your job is to add an equivalent case for the new increment (++) operator.

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

```
Expression ::= ...
  | IDENTIFIER:name
  { : RESULT = new IdentifierExp(name, namexleft); : }
  | /* add the CUP grammar rule and code for the ++
     operator in the box below */
```

(d) (4 points) Describe the checks that would be needed in the semantics/type-checking part of the compiler to verify that an expression using the ++ preincrement operator 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.

(continued on next page)

CSE P 501 19au Exam 11/21/19

Question 5. (cont.) (e) (6 points) Give the code shape for this new ++ preincrement operator, i.e., what code should be generated in the assembly language program to properly evaluate an expression with the new ++ operator.

You should show any instructions, labels, or other assembly language-level details needed, and also show where generated code to evaluate the subexpression that is the operand of the increment ++ operator would appear. Be sure you are clear about which parts of the code are computing or processing addresses of storage locations vs. the contents of those locations (i.e., lvalues vs. rvalues).

Your code does not need to be precisely correct x86-64 assembly code (i.e., you are not expected to have memorized instruction details), but it should be close enough so that your intent is clear and it basically equivalent to real x86-64 code.

Hints: you *definitely* won't need all the space on this page for your answer. But be careful that your code matches the semantics of the ++ operator precisely.

CSE P 501 19au Exam 11/21/19

Question 6. (14 points) x86-64. Translate the following C function into x86-64 assembly language. You must use the standard x86-64 C language conventions for function calls, register usage, and so forth. Note that this is just a C function, not a Java or MiniJava method, so there are no objects, no `this` pointer, and no vtables involved. You must translate the function exactly as given – you may not omit the recursive function call or other code, or rewrite the calculation into a loop or some other form. Assume that `ints` are 64-bit numbers, as in the MiniJava project.

There is a page at the end of the exam with rules for writing x86-64 code for this question along with some basic x86-64 assembly language reference information. Refer to it while answering this question, and feel free to remove that page from the exam for convenience.

```
// return the sum x+(x+1)+...+(y-1)+y
int sum(int x, int y) {
    if (x == y)
        return x;
    else
        return x + sum(x+1, y);
}
```

CSE P 501 19au Exam 11/21/19

Question 7. (14 points) Value numbering. For the following block, (i) apply local value numbering (including version numbers for variables) to the statements in the block, then (ii) rewrite the code to eliminate redundant expressions using the value numbering information from part (i).

(i) Show the results of value numbering by writing value numbers and variable version numbers on the code below. Do not change the code (not yet – that’s the next part). Recall that the notation v_i^n means that version i of variable v has value number n .

$$a = x$$

$$b = x + y$$

$$c = a + b$$

$$a = a + y$$

$$d = b$$

$$b = d + a$$

(ii) Now, rewrite the above code to show the results of eliminating redundant expressions using the value numbering information computed in part (i) above. Write your answer below.

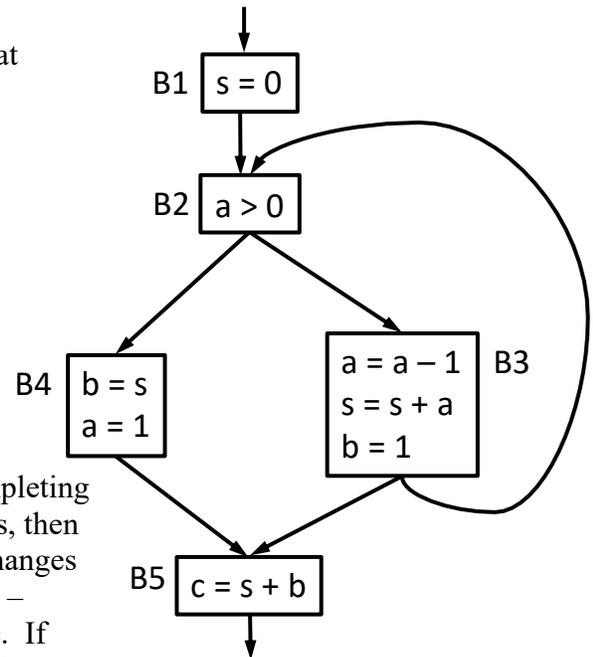
CSE P 501 19au Exam 11/21/19

Question 8. (16 points) Dataflow Analysis – Liveness. In class we looked at the liveness dataflow problem. Recall that a variable v is *live* at program point p if there is any path from p to a use of v along which v is not defined. For a basic block b we define these sets:

- $use[b]$ = variables used in b before any definition of that variable occurring in b .
- $def[b]$ = variables defined in b and not killed later in b .
- $in[b]$ = variables live on entry to b .
- $out[b]$ = variables live on exit from b .

These sets are related by these equations:

- $in[b] = use[b] \cup (out[b] - def[b])$
- $out[b] = \bigcup_{s \in succ[b]} in[s]$



Calculate the live variables for each block in this CFG by completing the table below. You should first calculate the use and def sets, then iteratively calculate in and out sets until there are no further changes in the solution. Note that the blocks are listed in reverse order – information for the last block (5) is in the first row of the table. If you need additional space for your answer, continue on the following blank page, and write a note at the end of this page to indicate that you are doing that.

Block	use	def	out	in	out	in	out	in
5								
4								
3								
2								
1								

CSE P 501 19au Exam 11/21/19

Additional space for answers if needed. Please be sure to label your answers and indicate on the original question that your answers are continued here.

CSE P 501 19au Exam 11/21/19

Extra scratch space or additional space for answers if needed.

CSE P 501 19au Exam 11/21/19

Reference information for use during the exam. Feel free to remove this page from the exam for convenience while answering questions.

The reverse side of this page contains the grammar for MiniJava.

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

- 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`
 - `%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 exam, even though this is not strictly required by the x86-64 specification.
- Pointers, Booleans, and ints are 64 bits (8 bytes) each, as in MiniJava.
- Your x86-64 code must implement all of the statements in the original method. You may *not* rewrite the method into a different form that produces equivalent results (i.e., replacing a function call by the function body or rewriting the code to compute the result in a different way from the original). Other than that, you can use any reasonable x86-64 code that follows the standard function call and register conventions – you do not need to mimic the code produced by a 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.)