



TESTS – 2021

TRIMESTER 2

SWEN 430

COMPILER ENGINEERING

Time Allowed: TWO HOURS

CLOSED BOOK

Permitted materials: No calculators permitted.
Non-electronic Foreign language to English dictionaries are allowed.

Instructions: Answer all questions

You may answer the questions in any order. Make sure you clearly identify the question you are answering.

Question	Topic	Marks
1.	Grammars and Parsing	20
2.	Types and Type Checking	20
3.	Static Analysis	20
4.	Java Bytecode	20
5.	Machine Code	20
6.	Memory Models	20
Total		120

1. Grammars and Parsing

(20 marks)

(a) Briefly, describe the following components of a compiler.

i. (2 marks) Lexer.

A lexer is responsible for grouping characters into *tokens* that can then be fed into the parser.

ii. (2 marks) Parser.

A parser is responsible for reading a sequence of tokens and checking whether (or not) they adhere to the grammar of the language in question. A parser typically produces an abstract syntax tree representation.

iii. (2 marks) Abstract Syntax Tree.

An abstract syntax tree is a programmatic representation of source program arranged as a tree.

(b) Consider the following grammar:

$$E \rightarrow N \mid (E , E)$$

$$N \rightarrow 0 \mid 1 \mid 2 \mid 3 \mid 4 \mid 5 \mid 6 \mid 7 \mid 8 \mid 9$$

i. (4 marks) For each of the following inputs, state whether it would be accepted or not by the grammar:

0

YES

(0, 1)

YES

(00, 1)

NO

((0, 1))

NO

- ii. (4 marks) Provide suitable Java classes for an Abstract Syntax Tree representation of the grammar from page 2,

```
1  interface Expr {}
2
3  class Tuple implements Expr {
4      private final Expr lhs, rhs;
5      public Tuple(Expr l, Expr r) { lhs = l; rhs = r; }
6  }
7
8  class Number implements Expr {
9      private final int n;
10     public Number(int n) { this.n = n; }
11 }
```

- iii. (6 marks) Complete the following class `Parser` which should implement a *recursive descent parser* for the grammar given on page 2:

```

public class Parser {
    private final String input;
    private int offset = 0;

    public Parser(String input) { this.input = input; }

    public Expr parseExpr() {
        char c = input.charAt(offset);
        if(c == '(') { return parseTupple(); }
        else { return parseNumber(); }
    }
    public Expr parseTupple() {
        match('(');
        Expr l = parseExpr();
        match(',');
        Expr r = parseExpr();
        match(')');
        return new Tupple(l,r);
    }
    public Expr parseNumber() {
        char c = input.charAt(offset);
        return new Number(Integer.parseInt("" + c));
    }
    public void match(char c) {
        if(input.charAt(offset++) != c) {
            throw new RuntimeException("error");
        }
    }
}

```

Student ID:

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.

Specify the question number for work that you do want marked.

2. Types and Type Checking

(20 marks)

Consider the following simple imperative language and its corresponding typing rules.

$s ::=$ (Statements) $\quad \text{Tx} = e$ (Declarations) $\quad x = e$ (Assignments) $\quad s ; s$ (Sequence) $\quad e$ (Expressions)	$\frac{}{\vdash c : \text{int}}$ (T-INT) $\frac{x : T \in \Gamma}{\Gamma \vdash x : T}$ (T-VAR)
$e ::=$ (Expressions) $\quad c$ (Integers) $\quad x$ (Variables) $\quad *e$ (Dereference) $\quad \text{new } e$ (Allocation)	$\frac{\Gamma \vdash e : \&T}{\Gamma \vdash *e : T}$ (T-Deref) $\frac{\Gamma \vdash e : T}{\Gamma \vdash \text{new } e : \&T}$ (T-New)
$T ::=$ (Types) $\quad \&T$ (Reference type) $\quad \text{int}$ (Int type) $\quad \text{void}$ (Void type)	$\frac{\Gamma \vdash x : T \quad \Gamma \vdash e : T}{\Gamma \vdash \text{Tx} = e : \text{void}}$ (T-Decl) $\frac{\Gamma \vdash x : T \quad \Gamma \vdash e : T}{\Gamma \vdash x = e : \text{void}}$ (T-Assign) $\frac{\Gamma \vdash s_1 : T_1 \quad \Gamma \vdash s_2 : T_2}{\Gamma \vdash s_1 ; s_2 : T_2}$ (T-Seq)

- (a) (5 marks) For each of the following typing judgements identify a suitable typing environment Γ , or explain why no such typing environment exists.

 $\Gamma \vdash x = 1 : \text{void}$
 $\Gamma = \{x \mapsto \text{int}\}$
 $\Gamma \vdash \&\text{int } x = \text{new } y ; z = *y : \text{void}$
None exists because y cannot be both $\&\text{int}$ and int
 $\Gamma \vdash x = \text{new } y ; *x : \text{int}$
 $\Gamma = \{x \mapsto \&\text{int}, y \mapsto \text{int}\}$
 $\Gamma \vdash y = x ; \&\text{int } x = \text{new } 1 : \text{void}$
 $\Gamma = \{x \mapsto \&\text{int}, y \mapsto \&\text{int}\}.$
 $\Gamma \vdash x = y ; y = *x : \text{void}$
None exists because no valid type for x exists (e.g. not $\&\text{int}$, nor $\&\&\text{int}$, etc)

- (b) Suppose the language were extended with a statement “delete e” which deallocates memory as in WHILE. For example, “delete p” deallocates the memory referred to by p.

i. (4 marks) Provide a suitable typing rule for this statement.

$$\frac{\Gamma \vdash e : \&T}{\Gamma \vdash \text{delete } e : \text{void}}$$

- ii. (5 marks) Executing “&int p = new 1 ; ... ; delete p” can result in a *stuck* program. Briefly, discuss what this means using an example to illustrate.

A program is stuck if it cannot continue executing, but has not yet reduced to a value. For example, the sequence “&int p = new 1 ; delete p ; delete p” will become stuck on the last statement since the memory referred to by p was already deallocated.

- iii. (6 marks) Introducing the delete statement means the simple *progress theorem* shown in lectures no longer holds for our language. Briefly, discuss what this means.

The progress theorem states that a well-typed program is not stuck (either its a value or it can reduce). Unfortunately, the delete statement means that well-typed programs *can* get stuck. For example, the program “&int p = new 1 ; delete p ; delete p” is well typed but gets stuck. The reason this happens is that the typing environment doesn’t include information about whether heap data has been deallocated or not.

3. Static Analysis

(20 marks)

This question concerns the *uniqueness analysis* developed for WHILE which determines, at each point, whether or not a variable is *defined*. A variable is defined after it has been assigned a value, but may become *undefined* if its value is *consumed* (e.g. moved to another variable). For simplicity, assume all references `&T` are *unique references*. For example, `&int` is a reference to an `int` variable and, furthermore, must be the *only* reference to that variable. The following illustrates:

```

1  &int p = new 123;
2  &int q;
3  // p is defined, q is undefined
4  if x >= 0 {
5      q = p;
6      // p is undefined, q is defined
7  }
8  // p and q are undefined

```

- (a) (5 marks) Explain briefly, using an example, why any algorithm for uniqueness analysis must be *conservative* (i.e. imprecise) in some way.

Static analyses cannot reason with perfect precision, and must draw safe (i.e. conservative) conclusions. In uniqueness analysis, for example, the analysis may not know for sure whether a variable was moved or not. But if it thinks it could be, then it must assume it has moved. For example, consider this variation on the program above:

```

1  assert x < 0;
2  ...
3  &int p = new 123;
4  &int q;
5  if x >= 0 { q = p; }

```

In this example, we know that `q` is never moved and, hence, is defined after the last statement. But, our uniqueness analysis cannot reason about conditions in this way.

- (b) (5 marks) Using examples to illustrate, explain briefly why a depth-first traversal algorithm is *insufficient* for implementing the uniqueness analysis.

A depth-first traversal visits every node in the control-flow graph exactly once. However, this is not sufficient for tracking uniqueness information around loops. Consider the following:

```

1  &int p = new 123;
2  while x < n {
3      &int q = p;
4  }
5  return p;

```

A depth-first traversal of the CFG for this graph will, in essence, take two paths: $1 \rightarrow 2 \rightarrow 5$ and $1 \rightarrow 2 \rightarrow 3$. In both of these paths, uniqueness information appears correct. In order to catch the problem, we must propagate information coming out of 3 back around the loop so that it eventually propagates into 5.

(c) A variable `x` is *consumed* by a statement if it must be undefined *after* that statement to preserve uniqueness. The method `consume(s)` returns the set of variables consumed by evaluating statement `s`.

i. (6 marks) Sketch an implementation of `consume(s)` for statements `x = e`, `assert e`, `delete e`, expressions `x`, `x == y`, `new e` and types `int`, `&int`. You may assume `typeof(x)` returns the declared type of a variable `x`.

```
consume(assert e) = ∅
consume(delete e) = consume(e)
consume(x = e) = consume(e)

consume(x == y) = ∅
consume(new e) = consume(e)
consume(x) = ∅, if typeof(x)=int
            = {x}, if typeof(x)=&int
```

- ii. (4 marks) Using `consume(s)`, give appropriate *dataflow equations* for the uniqueness analysis.

$$\begin{aligned} \text{UNIQ}_{IN}(0) &= \text{ARGS}(0) \\ \text{UNIQ}_{IN}(v) &= \bigcap_{w \rightarrow v \in E} \text{UNIQ}_{OUT}(w) \\ \text{UNIQ}_{OUT}(v) &= (\text{UNIQ}_{IN}(v) - \text{consumed}(v)) \cup \text{DEF}_{AT}(v) \end{aligned}$$

Student ID:

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.

Specify the question number for work that you do want marked.

4. Java Bytecode

(20 marks)

(a) Consider the following method written in Java bytecode:

```

boolean f(int[], int);
  0: iconst_0
  1: istore_3
  2: iload_3
  3: aload_1
  4: arraylength
  5: if_icmpge 24
  8: aload_1
  9: iload_3
 10: iaload
 11: iload_2
 12: if_icmpne 17
 15: iconst_1
 16: ireturn
 17: iload_3
 18: iconst_1
 19: iadd
 20: istore_3
 21: goto 2
 24: iconst_0
 25: ireturn

```

i. (5 marks) In the box below, give Java source code equivalent to the bytecode above:

NOTE: Appendix A on p21 provides an overview of bytecode instructions for reference.

```

1      public boolean contains(int[] items, int item) {
2          int i = 0;
3          while(i < items.length) {
4              if(items[i] == item) {
5                  return true;
6              }
7              i = i + 1;
8          }
9          return false;
10     }

```

- (b) **(2 marks)** Branch instructions in Java bytecode use *relative addressing*. Briefly, explain what this means.

Relative addressing means that the operand in the instruction identifies the target address as an offset from the current bytecode. The alternative is *absolute addressing* where the full address is encoded in the instruction.

- (c) **(6 marks)** Using an example to illustrate both Java source and the generated bytecode, explain what is meant by the term *short circuiting*.

Short circuiting is where the right-hand side of a logical operator (e.g. &&) is not executed when the outcome of the operator is already known. For example, consider this program:

```

1    public int filter(int item) {
2        if(item < 0 || item > 16) { return 0; }
3        else { return item; }
4    }

```

This would be translated into something like this:

```

1    0: iload_1
2    1: iflt 10
3    4: iload_1
4    5: bipush 16
5    7: if_icmple 12
6   10: iconst_0

```

We see that if the first condition is true, the second condition is not even evaluated.

(d) (7 marks) Translate the following method into Java bytecode:

```
1  public void fill(int[] items, int item) {  
2      for(int i=0;i!=items.length;i=i+1) {  
3          items[i] = item;  
4      }  
5  }
```

```
1  public void fill(int[], int);  
2      iconst_0  
3      istore_3  
4      .L1  
5          iload_3  
6          aload_1  
7          arraylength  
8          if_icmpeq L2  
9          aload_1  
10         iload_3  
11         iload_2  
12         iastore  
13         iload_3  
14         iconst_1  
15         iadd  
16         istore_3  
17         goto L1  
18     .L2  
19     return
```

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.

Specify the question number for work that you do want marked.

5. Machine Code

(20 marks)

Consider the following program written in WHILE:

```
1  int max(int x, int y) {  
2    if(x < y) { return y; }  
3    else { return x; }  
4  }
```

- (a) (6 marks) In the box below, translate the above program into X86_64 machine code. You should assume: (1) parameters `x` and `y` are passed in the `%rdi` and `%rsi` registers respectively; (2) the return value is passed in the `%rax` register; (3) all other registers are *callee-saved*.

NOTE: Appendix B on page 22 provides an overview of x86_64 instructions for reference.

```
1  max:  
2      cmpl %rdi, %rsi  
3      jge .L2  
4      movl %rdi, %eax  
5      ret  
6  .L2:  
7      movl %rsi, %eax  
8      ret
```


(b) On x86_64, the `rbp` register normally holds the *frame pointer*.

i. (4 marks) Briefly, discuss what the frame pointer is used for.

The frame pointer points to the start of the method's stack frame, and is used to access local variables to the method stored on the stack frame. Space may also be used for parameters and return values if/when these are not being passed in registers.

ii. (4 marks) Briefly, discuss whether a frame pointer is needed for method `max()`.

A framepointer is not required for method `max()` because all local variables can be stored in registers and, hence, a stack frame is not required.

(c) (6 marks) Briefly, discuss why *register allocation* is important for the performance of compiled programs.

Register allocation is important because reading / writing values from registers is much faster than from main memory. Thus, allocating variables into registers improves overall performance, especially if those variables are accessed many times. For example, if there is a loop then it is desirable to have all variables used in that loop stored in registers.

6. Memory Models

(20 marks)

(a) In the following *litmus tests*, x and y are *shared* variables, whilst $r1$ and $r2$ are *local* variables. Assume all variables are initialised with 0.

i. (5 marks) Under the *Sequential Consistency* model, can executing following program ever leave both $r1=1$ and $r2=1$ at the end? Justify your answer.

Thread 1	Thread 2
$r1 = x;$	$r2 = y;$
$y = 1;$	$x = 1;$

No. Either $r1 = x$ or $r2 = y$ is always executed before the other instructions and, hence, either $r1$ or $r2$ must be zero.

ii. (5 marks) Under the *Total Store Ordering (TSO)* model, can executing the following program ever leave both $r1=0$ and $r2=0$ at the end? Justify your answer.

Thread 1	Thread 2
$y = 1;$	$x = 1;$
$r1 = x;$	$r2 = y;$

Yes. Assume the sequence $y=1 ; x=1 ; r1=x ; r2=y$. Under TSO, both writes to x and y maybe stuck in the store buffer when both $r1=x$ and $r2=y$ are executed.

(b) A *data race* can occur when two threads access the same shared variable at the same time.

i. (2 marks) Can a data race occur if both threads *read* from the shared variable?

No, at least one write is required for a data race.

ii. (2 marks) Briefly, discuss how data races can cause variables to be assigned unexpected values.

If a variable is read and written at the same time, this can result in *tearing*. Specifically, where the value read contains part of the old value and part of the new value.

(c) (6 marks) Let `c` be an instance of `Channel` (defined below) and suppose **Thread 1** repeatedly calls `c.write(1)` and **Thread 2** repeatedly calls `c.read()`.

```

1  class Channel {
2      private int value = 0;
3      private volatile boolean ready = false;
4
5      public void write(int v) {
6          value = v;
7          ready = true;
8      }
9      public int read() {
10         while(!ready) { }
11         return value;
12     } }

```

On Java 5 (or later) can **Thread 2** ever read the value `0`? Justify your answer.

No. Since `ready` is marked **volatile** this has a *synchronising* effect in Java 5 or later. In effect, this means any access to this variable introduces a memory barrier which synchronises all cached variables with main (i.e. shared) memory. For example, on X86 this would result in the store buffer for each processor being flushed.

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.

Appendix A: Java Bytecodes

aaload	Load reference element from array onto stack.	...,aref,index \Rightarrow ...,ref
aastore	Store reference element into array from stack.	...,ref,index,val \Rightarrow ...
aload <i>n</i>	Load reference from local variable <i>n</i> onto stack.	... \Rightarrow ...,ref
areturn	Return reference from method.	...,ref \Rightarrow ...
arraylength	Push array length on stack.	...,aref \Rightarrow ...,int
astore <i>n</i>	Store reference into local variable <i>n</i> from stack.	...,ref \Rightarrow ...
bipush <i>c</i>	Load integer byte constant <i>c</i> onto stack.	... \Rightarrow ...,int
dup	Duplicate top item on stack.	...,val \Rightarrow ...,val,val
iadd	Add two ints on stack.	...,int,int \Rightarrow ...,int
iaload	Load int element from array onto stack.	...,ref,index \Rightarrow ...val
iastore	Store int element into array from stack.	...,ref,index,val \Rightarrow ...
iconst_c	Load integer constant <i>c</i> onto stack.	... \Rightarrow ...,int
idiv	Divide two ints on stack.	...,int,int \Rightarrow ...,int
iload <i>n</i>	Load int from local variable <i>n</i> onto stack.	... \Rightarrow ...,int
imul	Multiply two ints on stack.	...,int,int \Rightarrow ...,int
ineg	Negate int on stack.	...,int \Rightarrow ...,int
invokeinterface	Invoke interface method.	...,oref[val,[val,...]] \Rightarrow [val]
invokespecial	Invoke special instance method (e.g. initialisation).	...,oref[val,[val,...]] \Rightarrow [val]
invokestatic	Invoke static method.	...[val,[val,...]] \Rightarrow [val]
invokevirtual	Invoke instance method.	...,oref[val,[val,...]] \Rightarrow [val]
ireturn	Return int from method.	...,int \Rightarrow ...
istore <i>n</i>	Store int into local variable <i>n</i> from stack.	...,int \Rightarrow ...
isub	Subtract two ints on stack.	...,int,int \Rightarrow ...,int
if<cond>	Branch if int comparison with zero succeeds.	...,int \Rightarrow ...
if_acmp<cond> <i>d</i>	Branch to <i>d</i> if reference comparison succeeds.	...,ref,ref \Rightarrow ...
if_icmp<cond> <i>d</i>	Branch to <i>d</i> if int comparison succeeds.	...,int,int \Rightarrow ...
ldc <i>c</i>	Load constant (e.g. integer or string) <i>c</i> on stack.	... \Rightarrow ...,int
new <i>C</i>	Create a new object of class <i>C</i> \Rightarrow ...,ref
goto <i>d</i>	Branch unconditionally to <i>d</i> \Rightarrow ...
pop	Pop top item off stack.	...,val \Rightarrow ...
return	Return from method.	... \Rightarrow ...
sipush <i>c</i>	Load integer word constant <i>c</i> onto stack.	... \Rightarrow ...,int

Appendix B: x86_64 Machine Instructions

<code>movq \$c, %rax</code>	Assign constant <code>c</code> to <code>rax</code> register
<code>movq %rax, %rdi</code>	Assign register <code>rax</code> to <code>rdi</code> register
<code>addq \$c, %rax</code>	Add constant <code>c</code> to <code>rax</code> register
<code>addq %rax, %rbx</code>	Add <code>rax</code> register to <code>rbx</code> register
<code>subq \$c, %rax</code>	Subtract constant <code>c</code> from <code>rax</code> register
<code>subq %rax, %rbx</code>	Subtract <code>rax</code> register from <code>rbx</code> register
<code>cmpq \$0, %rdx</code>	Compare constant <code>0</code> register against <code>rdx</code> register
<code>cmpq %rax, %rdx</code>	Compare <code>rax</code> register against <code>rdx</code> register
<code>movq %rax, (%rbx)</code>	Assign <code>rax</code> register to dword at address <code>rbx</code>
<code>movq (%rbx), %rax</code>	Assign <code>rax</code> register from dword at address <code>rbx</code>
<code>movq 4(%rsp), %rax</code>	Assign <code>rax</code> register from dword at address <code>rsp+4</code>
<code>movq %rdx, (%rsi, %rbx, 4)</code>	Assign <code>rdx</code> register to dword at address <code>rsi+4*rbx</code>
<code>pushq %rax</code>	Push <code>rax</code> register onto stack
<code>pushq %c</code>	Push constant <code>c</code> onto stack
<code>popq %rdi</code>	Pop qword off stack and assign to register <code>rdi</code>
<code>jz target</code>	Branch to <code>target</code> if zero flag set.
<code>jnz target</code>	Branch to <code>target</code> if zero flag not set.
<code>jl target</code>	Branch to <code>target</code> if less than (i.e. sign flag set).
<code>jle target</code>	Branch to <code>target</code> if less than or equal (i.e. sign or zero flags set).
<code>ret</code>	Return from function.

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.