ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Exercise 7
Linearization, Bytecode Verification, and Parametric Polymorphism
November 18, 2022

Tasks covered in class

Tasks 1-8 will be covered in the exercise session. The remaining tasks are material for self-study.

Task 1

Consider the following declarations in Scala:

```scala
class C
trait T extends C
trait U extends C
class D extends C
```

Find all the types that can be created with or without traits, as well as the subtype relations between them.

Task 2

Consider the following Scala code:

```scala
class Cell {
    private var x:Int = 0
    def get() = { x }
    def set(i:Int) = { x=i }
}

trait Doubling extends Cell {
    override def set(i:Int) = { super.set(2*i) }
}

trait Incrementing extends Cell {
    override def set(i:Int) = { super.set(i+1) }
}
```

A) What is the difference between the following objects?

```scala
val a = new Cell
val b = new Cell with Incrementing
val c = new Cell with Incrementing with Doubling
val d = new Cell with Doubling with Incrementing
```

B) We try to use the following code to implement a cell that stores the argument of the set method multiplied by four:

```scala
val e = new Cell with Doubling with Doubling
```

Why does it not work? What does it do? How can we make it work?

C) We propose the following solution to support traits together with behavioral subtyping: assume `C` is a class with specification `S`. Each time we create a new trait `T` that extends `C`, we must ensure that `C with T` also satisfies `S`. Show a counterexample that demonstrates that this approach does not work.

Task 3

Consider a Java class `E`, which has a method `f` with the following signature: `void f();`

The method `f` has one local variable `v` and the following body:

```
0: iconst 5
1: istore 1
2: aload 0
3: astore 1
4: iload 1
5: iconst 1
6: iadd
7: istore 1
8: return
```

The maximal stack size is equal to 1. Can the provided bytecode be verified? If so then verify it, otherwise explain which line of code causes the problem and why.

Task 4

The Java bytecode verifier is more permissive than the Java type system. Provide a program that demonstrates this.

Task 5

(from a previous exam)

Assume two Java classes `A` and `B`, where `B` is a subclass of `A`. Consider the following bytecode:

```
0: aload 1
1: astore 2
2: goto 0
```

and assume that the input to the initial node of this code is `([],[A,A,B])`, where the first list indicates the content of the stack and the second list indicates the content of the registers.

After running the bytecode type inference algorithm, what is the inferred input to the initial node?

(a) `([],[A,A,A])`

(b) `([],[A,A,B])`

(c) `([],[A,B,B])`

(d) Nothing is inferred – the type inference does not terminate

(e) Nothing is inferred – the type inference rejects the program

Task 6

Consider the following Java code:

```java
interface IFace { void m(); }

class Cl1 implements IFace {
    public void m() { System.out.println("Cl1.m"); }
}
class Cl2 implements IFace {
    public void m() { System.out.println("Cl2.m"); }
}

public class Test {
    public static void main(String[] args) {
        foo(true);
        foo(false);
    }
    public static void foo(boolean param) {
        IFace iface = null;
        if (param) { iface = new Cl1(); }
        else { iface = new Cl2(); }
        iface.m();
    }
}
```

A) What type will be calculated for the variable `iface` of the method `foo` during bytecode verification?

B) When can we decide that `iface.m()` is safe to call, during bytecode verification or during execution?

C) Would your answer from **B** be the same if `IFace` were a class instead of an interface? What if `IFace` were an abstract class?

Task 7

(from a previous midterm)

Consider the following Java program, which is rejected by the Java compiler:

```java
class Logger<T> {
    public void log(T t) {
        System.out.println(t.loggerString());
    }
}
```

A) If the code above were allowed to compile without errors, what could go wrong? To answer, write a brief code sample that uses `Logger` in a way which causes a failure at runtime.

B) How can we fix the class `Logger` so that it compiles, while preserving its functionality? You should not modify the method `log`, but otherwise can change or add any code. Your solution should include all details required to check that `Logger` is a valid Java class.

C) Assume that class `Logger` has been fixed to resolve the problem from point **A**. Let `A` and `B` be two classes such that `A` is a supertype of `B` and `Logger<A>` and `Logger<B>` are valid instantiations. Consider the following method:

```java
void foo(Logger<A> logA) {
    Logger<B> logB = logA;
    logB.log(new B());
}
```

The Java compiler rejects this code. Is the code safe? That is, if it were allowed to compile, would it run without failure?

D) Suppose we relax the Java type system rules to allow contravariant generics.

- Will the method `foo` compile now?

- What are two situations that will require dynamic checks in order to enable contravariant generics in a language, without limiting what a developer can write in a generic class?

Task 8

(from a previous exam)

A) Recall the Java interface `Comparable<T>` that was shown in the lecture:

```java
public interface Comparable<T> {
    public int compareTo(T other);
}
```

The method `compareTo` returns $\begin{cases} 1 & \text{if \texttt{this} is greater than \texttt{other}} \\ 0 & \text{if \texttt{this} is equal to \texttt{other}} \\ -1 & \text{if \texttt{this} is less than \texttt{other}} \end{cases}$

Suppose we want to turn `Comparable` into an abstract class with an additional helper method `greaterThan`, that returns true if and only if `this` is greater than `other`.

Assume the following implementation:

```java
public abstract class Comparable<T> {
    public abstract int compareTo(T other);

    public boolean greaterThan(T other) {
        return other.compareTo(this) < 0;
    }
}
```

A.1) Why does this implementation not type check?

A.2) Fix the type error by changing only the body of `greaterThan`, while preserving the intended semantics of the method.

A.3) Fix the type error by changing only the class signature and the signature of the method `compareTo`.

B) Suppose we have the following class:

```java
class A<X                    ,Y                    > {
    X a;
    Y b;
}
```

Consider a variable `v` whose type is `A<S,T>` where `S` and `T` satisfy the type bounds that you have to insert above. Your type bounds have to guarantee that for all sequences of `a` and `b` accesses on `v` (e.g., `v.b.a.b.a.a.b.b`) the following two properties hold:

- The static type of a sequence ending in `a` is `S`.

- The static type of a sequence ending in `b` is `T`.

Tasks not covered in class

Task 9

(from a previous exam)

Consider the following Scala code:

```scala
class A            {          def bar() =                  ""   }
trait B extends A { override def bar() = super.bar() + "B" }
trait C extends B { override def bar() = super.bar() + "C" }
trait D extends B { override def bar() = super.bar() + "D" }

object Main {
  def main() { foo(new A with D with C with B) }
  def foo(x: A with D) { println(x.bar()) }
}
```

What would be the output of the call `Main.main()`?

(a) BDB

(b) BBDBC

(c) BBCBD

(d) DB

(e) BDC

(f) BCD

(g) None of the above

Task 10   (from a previous exam)

Consider the following Scala code, which compiles correctly and models some jobs a `Person` may have. To work as a `Lawyer` or as a `TaxiDriver`, one needs to have a valid license. This requirement can be expressed through *self type annotations* added to the traits `Lawyer` and `TaxiDriver` (as in the given code). These annotations are checked by the compiler and allow the traits `Lawyer` and `TaxiDriver` to be mixed only into subtypes of `PersonWithLicense`. Self type annotations enable code reuse without subtyping, that is, `Lawyer` and `TaxiDriver` ⋠ `PersonWithLicense`, but the methods of the class `PersonWithLicense` are available and can be overridden inside these two traits.

```scala
class Person { def work(): String = { return "working"; }}

class Student { def work(): String = { return "studying"; }}

class PersonWithLicense extends Person {
 def hasValidLicense(): Boolean = { return false; }
}

trait Gardener extends Person {
 override def work(): String = { return super.work() + " in the garden";}
}

trait Lawyer extends Person {
 this: PersonWithLicense => // self type annotation

 override def work(): String = {
```

```
    if(this.hasValidLicense()) return super.work() + " in court";
    return "not " + super.work();
  }

  override def hasValidLicense(): Boolean = { return true; }
}

trait TaxiDriver extends Person {
  this: PersonWithLicense => // self type annotation

  override def work(): String = { return super.work() + " in Zurich"; }
}
```

A) For each of the following two code fragments (**A.1** and **A.2**), if they compile, write the output of their execution. Otherwise, briefly explain why they are rejected by the compiler.

A.1

```
val lawyer: Lawyer = new PersonWithLicense with Lawyer with TaxiDriver;
println(lawyer.work());
```

A.2

```
val student: Gardener = new Student with Gardener;
println(student.work());
```

B) Add **one** method to any of the given classes or traits **except** `PersonWithLicense` (explicitly write to which one) and fill in the instantiation from the client code below, such that it compiles and when executed prints `not working in Zurich in the garden`. You are **not allowed** to directly return this string, to use reflection, to define new classes or traits, nor to modify the given code. If this is not possible, briefly explain why.

```
// Client code:
val person = new _____
println(person.work());
```

The following method should be added to: [   ]

Task 11

Assume we have two Java classes A and B. Consider the following Java class C:

```
class C {
    void foo(A x) {
        int y = 7;
        this.bar(y, x);
    }

    B bar(int u, A v) {
        ...
    }
}
```

Assume that the method `foo` gets compiled into bytecode as follows:

```
0: iconst 7
1: istore 2
2: aload 0
3: aload 2
4: aload 1
5: invokevirtual C.bar.B(int,A)
```

Can this bytecode be verified? If so, what is the final state (after line 5)?

Task 12

(from a previous exam)

Consider an *incorrect* bytecode verifier called *BuggyVerifier*, in which due to a bug the `aload` rule assumes that the loaded element is stored at the bottom of the stack instead of at the top (see the formal description below), while all the other rules are implemented correctly.

$$\text{aload } n :$$
$$(S, R) \to (\underline{S.R(n)}, R),$$
$$\text{if } 0 \le n < MR \ \wedge \ R(n) <: Object \ \wedge \ |S| < MS$$

Assume that the initial state (stack and registers) is $([], [A, B])$, with the maximum number of registers $MR = 2$ and the maximum stack size $MS = 2$. A and B are classes such that $B <: A$.

A) Write a short bytecode program that is accepted by *BuggyVerifier*, but is **not** accepted by a correct bytecode verifier. Clearly mark the line at which the correct verifier detects an error, and briefly describe the error.

You can use in your solution all the bytecode operations seen during the lectures. As a reminder, here are some of them:

- `iconst` n: create on the stack a value n of type `int`.

- `iload` n: load on top of the stack an element of type `int` from the n-th register.

- `astore` n: remove an object from the top of the stack and store it in the n-th register.

- `goto` n: continue the execution from the operation at label n.

B) Is it possible that *BuggyVerifier* incorrectly accepts a program that overflows the stack, by pushing more than MS elements? Write yes or no, then motivate your answer.

Task 13

A) Compare dynamic type checking with the `dynamic` keyword to static type inference with `var` in C#:

- Give a correct program which can be realized with `dynamic` but not with `var`.

- Give an incorrect program which will be accepted by the compiler with `dynamic` but not with `var`.

B) C#'s most general type is `object`. Similar to `var` and `dynamic`, you can write `object x = ...` with an expression of any type on the right-hand side.

- Given a compiling program using `var`. Can we replace all `var` keywords by `object` and add explicit casts in the right places so that the program compiles and runs as before?

- Given a compiling program using `dynamic`. Can we replace all `dynamic` keywords by `object` and add explicit casts in the right places so that the program compiles and runs as before?

For both questions, either informally describe how to do the replacement, or give a counter-example where the transformation will always produce a program that does not compile or behaves differently. Note that explicit casts to `dynamic` are not allowed in the transformation.

C) Assume now a language like C#, but with covariant return types and contravariant parameter types. Given four classes A, B, C and D:

```
class A { int m (int x); }
class B { void m (dynamic x); }
class C { dynamic m (int x); }
class D { dynamic m (dynamic x); }
```

Develop a subtyping rule for the `dynamic` type annotation and informally explain the reasoning behind it. What are the potential subtypes among the four classes above?

Task 14

In this task, you have to implement (using three different approaches) a list in Java that supports the following two methods (where `i` represents an index):

```
public void add(int i, Object el)
public Object get(int i)
```

Discuss the advantages and the limitations of the three different approaches below.

A) Implement the list using only one class without generics.

B) Implement the list using one abstract class/interface and then (some) subclasses that implement it for different types.

C) Implement the list using generic types.