ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Exercise 8

Multiple Inheritance and Linearization

November 13, 2020

Task 1 (from a previous exam)

A) In this task you will have to model a bakery, which produces and sells different types of `BakedProducts`, such as `Bread` and `SeasonalProducts`. Moreover, for different events (e.g., Saint Nicholas or Easter), the bakery offers `SeasonalBread`.

Fill in the C++ class declarations below, such that all the following subtype relations are fulfilled:

- `Bread <: BakedProduct` and `SeasonalProduct <: BakedProduct`
- `Bread` $\not<:$ `SeasonalProduct` and `SeasonalProduct` $\not<:$ `Bread`
  ($\not<:$ means *is not a subtype of*)
- `SeasonalBread <: SeasonalProduct` and `SeasonalBread <: Bread`

Make sure that your code is accepted by the C++ compiler and when executed prints:

Grittibaenz for 6th of December

Easter bread for 1st of April

```
//Bakery code:
class Product {
  public: string name;
          Product(string n) { name = n; };
};

class BakedProduct : public Product {
  public: BakedProduct(string n): Product(n) {};
};

class Bread : ____________________________________________ {
  public: Bread(string n) :

          _________________(n) {};
};

class SeasonalProduct : ____________________________________ {
   public: SeasonalProduct(string n, string e) :

          _________________(n + " for " + e) {};
};

class SeasonalBread : ______________________________________ {
  public: SeasonalBread(string a, string b):

          ____________________________________________________
```

```
                                                                    {};
};

//Client code A:
  Product* prod1 = new SeasonalBread("6th of December", "Grittibaenz");
  Product* prod2 = new SeasonalBread("1st of April", "Easter bread");
  cout << prod1->name << endl;
  cout << prod2->name << endl;
  // prints: Grittibaenz for 6th of December
  //         Easter bread for 1st of April
```

solution

```
//Bakery code:
class Product {
  public: string name;
          Product(string n) { name = n; };
};

class BakedProduct : public Product {
  public: BakedProduct(string n): Product(n) {};
};

class Bread : public virtual BakedProduct {
  public: Bread(string n) :
          BakedProduct(n) {};
};

class SeasonalProduct : public virtual BakedProduct {
   public: SeasonalProduct(string n, string e) :
           BakedProduct(n + " for " + e) {};
};

class SeasonalBread : public (virtual) Bread,
                      public (virtual) SeasonalProduct {
  public: SeasonalBread(string a, string b):
          BakedProduct(b + " for " + a),
          Bread(b),
          SeasonalProduct(b, a) {};
};
```

The classes `Bread` and `SeasonalProduct` have to inherit virtually from `BakedProduct`, otherwise the client code does not compile due to ambiguity (in case of non-virtual inheritance, `SeasonalBread` would have two copies of `Product`).

The class `SeasonalBread` can inherit both virtually and non-virtually from `Bread` and `SeasonalProduct`, the result is the same.

The order of the super constructor calls from `SeasonalBread` is arbitrary. The arguments passed to the constructors of `Bread` and `SeasonalProduct` are also arbitrary.

B) C++ supports *private inheritance*, which allows code reuse (access to methods and fields) **without** subtyping. Assume that we change the declaration of the class `BakedProduct` from **Task A** to use *private*, instead of public inheritance. All the other classes remain unchanged.

Fill in the blanks from *Client code B*, such that it compiles and when executed prints the same strings as in **Task A**. You are allowed to add methods (but *not* constructors) to any of the provided classes. For each new method that you add, please explicitly write to which class it belongs. You are *not allowed* to make any other changes.

```
//Client code B:
```

```
__________________ prod1 = new ___________________("6th of December",
                                                    "Grittibaenz");
__________________ prod2 = new ___________________("1st of April",
                                                    "Easter bread");
cout << prod1 ________________________________ << endl;

cout << prod2 ________________________________ << endl;
// prints: Grittibaenz for 6th of December
//         Easter bread for 1st of April
```

solution

```
class BakedProduct : private Product {
  public: BakedProduct(string n): Product(n) {};
          // additional method
          string getName(){ return name; }
};

//Client code B:

  BakedProduct* prod1 = new SeasonalBread("6th of December",
                                          "Grittibaenz");
  BakedProduct* prod2 = new SeasonalBread("1st of April",
                                          "Easter bread");
  cout << prod1->getName() << endl;

  cout << prod2->getName() << endl;
  // prints: Grittibaenz for 6th of December
  //         Easter bread for 1st of April
```

`prod1` and `prod2` can have any static type, but `Product`. To obtain the same strings as in **Task A**, their dynamic type has to be `SeasonalBread`.

Task 2

Consider the following declarations in 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.

solution

We can create the following types:

`C`, `D`, `T`, `U`,
`C with T` (same type as `T`, because `T extends C`),
`C with U` (same type as `U`, because `U extends C`),
`C with T with U` (same type as `C with U with T`),
`D with T`,
`D with U`,
`D with T with U` (same type as `D with U with T`).

The subtype relation is reflexive and transitive. Moreover, let X', Y' be the two base classes from which we derive X and Y by mixing in traits. Let A be the set of all traits mixed into the first class and B the set of all traits mixed into the second class. The rule is as follows:

$$X <: Y \text{ if and only if } X' <: Y' \text{ and } A \supseteq B.$$

Note: The above rule applies in our example, but it is not a general rule for subtyping in the presence of traits. Note that even if `D with T with U` and `D with U with T` are equivalent types (subtypes of each other), they can describe different behavior, causing subtle problems for behavioral subtyping.

We can also visualize the types and the subtype relations between them (the edges corresponding to reflexive and transitive subtype relations were omitted):

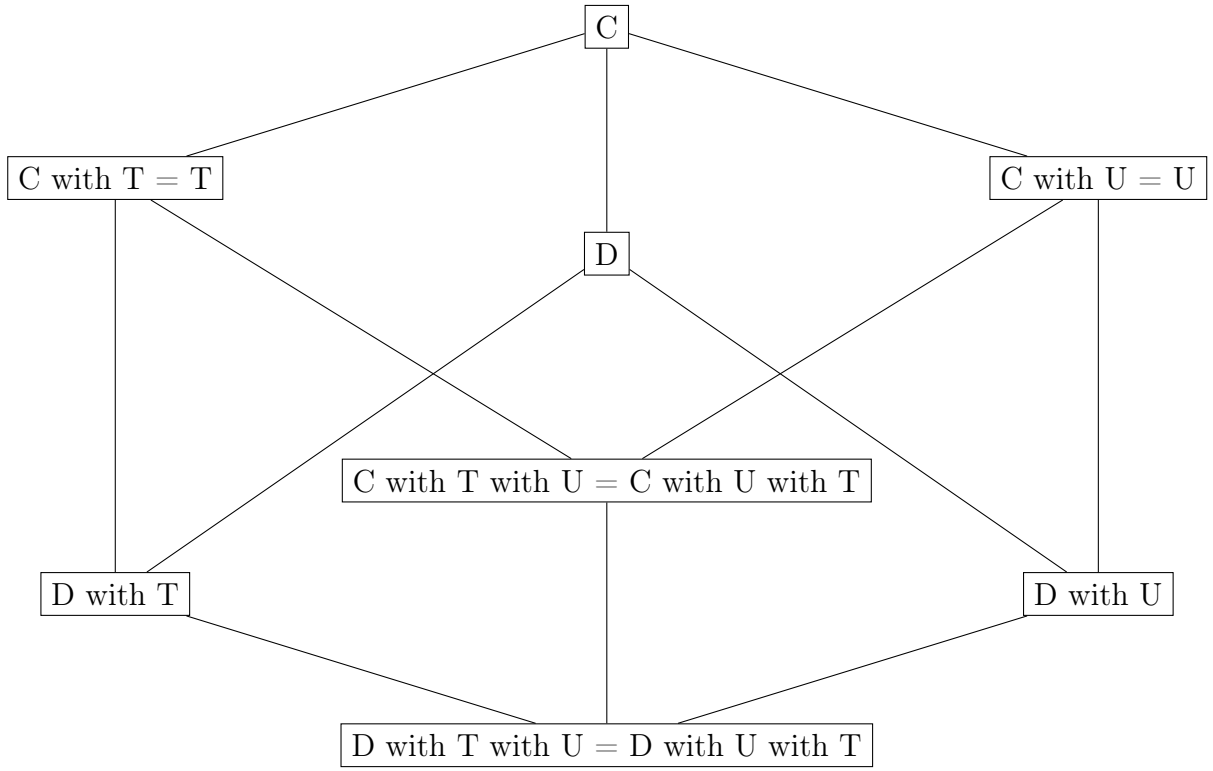

Task 3

Consider the following Scala code:

```
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?

```
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
```

solution

Object `a` behaves like a normal cell. Object `b` is also a cell, but it increases the stored value by 1. The interesting difference is between `c` and `d`. They are both cells. They have mixed in exactly the same traits. However, calling `set(i)` has a different effect on them: it stores `2i+1` in the first one and `2(i+1)` in the second one.

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

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

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

solution

Trait `Doubling` will not get mixed in twice. Scala rejects this statically.

A possible attempt to bypass the problem is to create a new trait `Doubling2` that behaves exactly like `Doubling`:

```
trait Doubling2 extends Doubling
val e = new Cell with Doubling with Doubling2
```

The code passes through, but dynamically `e` behaves as if it were a `Cell with Doubling`. Scala lets the code go through, because `Doubling2` may introduce new functionalities, but does not include `Doubling` twice in the linearization.

Our last try (the one that works), is to create a whole new trait from scratch, without reusing existing code:

```
trait Doubling3 extends Cell {
   override def set(i:Int) = { super.set(2*i) }
}
val e = new Cell with Doubling with Doubling3
```

And now `e.set` quadruples its argument, as expected.

C) Find a modularity problem in the above, or a similar, situation. Hint: a client that is given a class `C` does not necessarily know if a trait `T` has been mixed in that class.

solution

A problem is that a method that accepts `Cell with Doubling with Incrementing` as an argument could also be passed a class of the type `Cell with Incrementing with Doubling` - so what it can actually assume about its inputs is less than would be expected.

D) 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.

solution

Consider the following example:

```
class Cell {
   var x:Int = 0
   // ensures x <= i + 1
   def set(i:Int) = { x=i }
}

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

trait Incrementing2 extends Cell {
   override def set(i:Int) = { super.set(i+1) }
}
```

Both `Cell with Incrementing` and `Cell with Incrementing2` are behavioral subtypes of `Cell`. But `Cell with Incrementing with Incrementing2` is not a behavioral subtype of `Cell`, as the following example shows:

```
val c: Cell = new Cell with Incrementing with Incrementing2
c.set(4)
assert(c.x <= 5) // fails, postcondition of Cell.set not fulfilled
```

Task 4

(from a previous exam)

Consider the following Scala code:

```
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) **CORRECT:** BDC

(f) BCD

(g) None of the above

solution

The super calls are resolved based on the linear order.
L(A with D with C with B) = L(B), L(C), L(D), L(A) (*)

L(A) = A
L(D) = D, B, A
L(C) = C, B, A
L(B) = B, A
We now substitute the linearizations of A, D, C, B in (*) (in this order) and make sure the same class/trait is not included twice:
L(A with D with C with B) = L(B), L(C), L(D), A
L(A with D with C with B) = L(B), L(C), D, B, A
L(A with D with C with B) = L(B), C, D, B, A
L(A with D with C with B) = C, D, B, A

The call `x.bar()` corresponds to `C.bar()`, as `C` is the first in the linear order.
`C.super().bar()` is `D.bar()`, as `D` follows after `C` in the linear order.

Task 5 (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` $\not<:$ `PersonWithLicense`, but the methods of the class `PersonWithLicense` are available and can be overridden inside these two traits.

```
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());
```

solution

The code compiles: the traits `Lawyer` and `TaxiDriver` are mixed into a subtype of `PersonWithLicense`, as required by the self type annotation and `new PersonWithLicense with Lawyer with TaxiDriver <: Lawyer`.

The linearization of `new PersonWithLicense with Lawyer with TaxiDriver` is `TaxiDriver`, `Lawyer`, `PersonWithLicense`, `Person`. When executed, the code prints: working in court in Zurich

A.2

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

solution

The code does not compile, because `Student` is not a subclass of `Person` (the superclass of the trait `Gardener`)

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: ____

solution

```
// Client code:
val person = new PersonWithLicense with Lawyer with TaxiDriver with
   Gardener;
println(person.work());

trait TaxiDriver extends Person {
  ...
  // additional method:
  override def hasValidLicense(): Boolean = { return false; }
}
```

Note that if we try to add the method `hasValidLicense` to the trait `Gardener`, the client code does not compile, as `new PersonWithLicense with Lawyer with TaxiDriver with Gardener` inherits two methods with the same signature.