

Automatic Parallelisation in OO Languages with Balloons and Immutable Objects

Marco Servetto Alex Potanin

Victoria University of Wellington, New Zealand
{servetto,alex}@ecs.vuw.ac.nz

Abstract

Automatically detecting when it is safe to execute program expressions in parallel is hard. In this paper we take a step towards this ultimate goal by extending the original Almeida’s [1] balloons system with immutability and we utilise this information to figure out when the execution *order* of some expressions does not influence the result. Our type system for an imperative object-oriented language takes away the burden of parallelisation from the programmer and does not require the use of complex program analysis tools.

Categories and Subject Descriptors D.3 Programming Languages [D.3.3 Language Constructs and Features]

General Terms Languages

Keywords Full encapsulation, immutability, auto parallelisation

1. Introduction

In this paper we present Balloon Immutable Java (BI-JAVA), a language that supports both immutability and aliasing restrictions that can aid parallelisation. Traditionally, full encapsulation was considered too strict to be expressive enough for general purpose programming [8]. Further work chose to extend it with either more flexible encapsulation guarantees, such as ownership types [13], or some form of immutability, such as readonly references in Universes [11]. To be useful for parallelisation, we are concerned with *reachability* of object references and thus we consider more flexible encapsulation schemes to be not as useful for our guarantees. BI-JAVA chooses to keep a full encapsulation scheme and couple it with immutability support as the basis of our approach. We start with a tour of BI-JAVA concepts.

1.1 Immutable Objects

There is a significant recent body of work on immutable objects in Java-like languages [6, 14]. An *immutable* object can never be mutated via any reference to it. Immutability can be either deep or shallow property and most approaches distinguish between *mutable* and *immutable* objects. Additionally, references to objects can be treated as *readonly* to prevent changes via this particular reference to an object. A typical approach is to introduce types that can distinguish readonly references, references to mutable objects, and references to immutable objects. Commonly a readonly reference

type is treated as a supertype of both reference to a mutable and an immutable object [18, 19].

BI-JAVA offers transitively immutable objects, as typically found in functional programming languages. They provide capability to share and offer closed semantics. They are referentially transparent: it is irrelevant if two references point to the same immutable object or to different but equivalent immutable object graphs. Indeed a language offering immutable values can provide useful memory optimisations under the hood.

One problem with introducing immutable objects is safe initialisation of immutable (and potentially cyclic [19]) data structures. Our solution is to introduce *fresh* objects — these objects dominate their reachable object graph and only allow heap references to such objects to come from inside their reachable object graph. Moreover, references from inside a fresh reachable object graph to the outside are only allowed if they are to immutable objects. Thanks to these strong limitations, *fresh* objects can safely be converted both to immutable objects and to mutable objects.

So far, we have introduced three kinds of objects: *fresh*, *mutable*, *immutable*; and four kinds of references: *fresh*, *mutable*, *immutable*, *readonly*.

The first main contribution of this paper is a type system that can detect if a *mutable* object can be converted to *fresh*. We call this mechanism *permanent type promotion*.

1.2 Balloon Objects

Almeida [1] introduced a concept of *balloon* objects. The original balloons were class based and maintained a unique reference to the “balloon object” (thus no objects, including those “inside” the balloon, could refer to it). We propose a fourth and final kind of object in our language: *balloon* that extends Almeida’s concept by allowing references from *inside* the balloon to the “balloon object”.

The *balloon* reference is somewhat similar to the concept of a *unique* reference in a sense that it is the only reference to a balloon object from outside of the balloon¹. Thus, having a balloon reference guarantees that you are the only holder of a *balloon* reference to this balloon object. In our language balloons are object based and can be nested inside other balloons.

Just like *readonly* reference can point at both *mutable* and *immutable* objects, we introduce an *external* reference type that can point both at *mutable* and *balloon* objects. An *external* reference to an object can be passed to methods and stored in stack variables but can never be put inside a balloon. We generalise both *external* and *readonly* reference types as *external readonly* which are references to objects that neither can be put inside balloons nor mutated.

Figure 1 shows all four kinds of objects (*balloon*, *mutable*, *immutable*, and *fresh*) and seven kinds of reference types (*balloon*,

¹Our *balloon* (and also *fresh*) concept is similar to the externally unique object [5] except that we prohibit all (but immutable) objects inside a balloon from pointing to the outside of the balloon.

mutable, *immutable*, *fresh*, as well as *readonly*, *external*, and *external readonly* present in BI-JAVA.

1.3 Four Kinds of Heaps and Balloon Invariant

Almeida distinguishes two kinds of heaps: *static* (or normal) heap prohibits any references to objects inside a balloon that bypass the entry object; *dynamic* heap on the other hand allows references to objects inside balloons stored in the static heap, however no references from static heap to dynamic heap are allowed. We extend the Almeida’s heap concepts by introducing four kinds of heaps:

1. *immutable heap* contains all immutable objects;
2. *fresh heap* contains all fresh objects and their reachable object graphs, but not the objects in the immutable heap;
3. *shared heap* (similar to Almeida’s static heap) can contain mutable objects that are referred to using `m` mutable references from the stack and their reachable mutable object graphs. Notably such graphs can refer to `b` balloon objects (see Section 4.6);
4. *temporary heap* (similar to Almeida’s dynamic heap) contains all the other objects; that is objects referred to using `b` balloon, `e` external, `r` readonly or `er` external readonly references from the stack and their reachable object graphs (as long as such objects are not already captured by the above three heap definitions).

During the program execution objects inside the fresh heap can migrate to any other kind of heap; objects inside the shared heap can migrate to the temporary heap; while objects inside the temporary heap can never migrate to any other heap.

Now we can clearly re-state the balloon invariant by Almeida [1] using the object types and heap kinds from BI-JAVA:

- I_1 at most one object can refer to a balloon object using a field of balloon type. Other objects can refer to a balloon if the field is of `m` mutable type (or a supertype thereof);
- I_2 this `b` balloon reference (if it exists) is external to the reachable object graph from the balloon object itself;
- I_3 for any mutable object on the stack, its reachable mutable object graph can refer to the internal objects of at most a single balloon;
- I_4 balloons referred to from the shared heap or directly referred to from the stack refer to disjoint object graphs, and all direct sub-balloons contained in a balloon refer to disjoint object graphs as well.

Note how if some `b` balloon, `e` external, `r` readonly or `er` external readonly references refer to disjoint object graphs, any operation over these “disjoint” references will not establish any connection between these disjoint object graphs. That is, any mutable object created during such operation is: (1) injected into the reachable object graph of some `b` balloon or `e` external reference; (2) is in the reachable object graph of the result of the operation; or (3) can be safely garbage collected at the end of the operation itself.

Our second main contributions is a type system that is able to recognise expressions returning a value not injected in any `b` balloon or `e` external reference. Inside such expressions, any non `f` fresh reference can be seen as `r` readonly and thus stored inside conventional data structures. We call this mechanism *temporary type promotion*.

1.4 Field Update Operation

The limitations that are imposed by immutable (or readonly) and balloons (or external) references find its origins in the typical field update operation:

$$e_0.f = e_1$$

There are two parts to a field update: on the left, the field of e_0 is modified, and on the right, the result of e_1 is being stored.

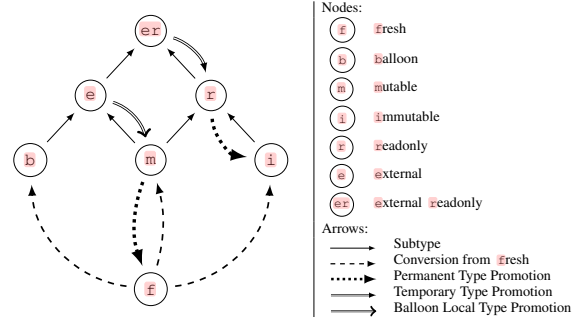


Figure 1. Object and reference kinds and possible relationships

Field update operation is prohibited if

- receiver (e_0) is `i` immutable, `r` readonly or `er` external readonly.
- assigned value (e_1) is `b` balloon, `e` external or `er` external readonly.

Moreover, if the receiver (e_0) is `b` balloon or `e` external, the assigned value has to be `f` fresh or `i` immutable.

The original proposal for Balloons uses a copy assignment. Our system clearly subsumes that proposal, since a (deep-)clone method would produce a `f` fresh result.

Our third main contribution is a type system that is able to recognise expressions that are executed “as inside” of a `b` balloon: inside such expression an `e` external reference to the inside of this balloon can be seen as `m` mutable, thus allowing free field assignment. We call this mechanism *balloon local type promotion*.

Figure 1 shows our object and reference types together with subtype relationship, conversion from fresh and the three kinds of promotion.

1.5 Why Balloons Are Useful for Parallelisation

Balloons allow modular reasoning since they offer *closed semantics*: if two `b` balloon references are different, the corresponding balloons are disjoint. Firstly, this helps the programmer by simplifying reasoning. Secondly, it is possible to use this information for automatic parallelisation. Consider the following example:

```
int h(b C x1, b C x2, i D x3, f G x4) {
    int x = x3.m(x1, x4);
    int y = new H().k(x2, x3);
    return x+y; }
```

In this simple setting, there are two expressions that have to be executed before being able to compute a result. Thanks to the property of `b` balloon (and `f` fresh and `i` immutable) references, no matter what computation will be executed by methods `D.m` and `H.k`: if the reference x_1 is not equal to the reference x_2 the two expressions can be computed in parallel, preserving the semantics of the program. That is, in many cases a simple dynamic test on pointers ($x_1 \neq x_2$) is enough to verify the correctness of this optimisation, opening the door for a completely safe speculative parallelism.

Our fourth main contribution is a compilation process able to infer simple logic expressions over pointer equalities and to dynamically enable fork-join parallelisation according to the results of these logic expressions, preserving the original semantics.

Outline. The rest of the paper is organised as follows: Section 2 and Section 3 presents the BI-JAVA language in detail, with examples showing the flexibility of our typing discipline. Section 4 provides semantics and type system rules for the BI-JAVA language without any parallelisation constructs. Section 5 adds a fork-join construct and describes a compilation process able to inject such construct in a sequential code preserving the semantics. Finally, Section 6 discusses related work and Section 7 concludes.

2. A Tour of the BI-JAVA Language

In this section we use a series of examples written in BI-JAVA to give a feel for all major aspects of our language.

2.1 Mutable and Readonly References

Mutable references are the closest kind of reference in our language to the default reference in Java. Mutable references provide capability to *share* and to *write* an object. As with usual Java, because of aliasing [8], any change to any object can potentially have an effect on the object referred to by a mutable reference.

A mutable reference can be safely cast to a readonly reference but once treated as readonly, it can never be cast back to mutable again. This guarantees that the object (and its reachable object graph) will not be modified via such readonly reference. However, this does not prevent observational exposure [4] as readonly references can be stored in any (type compatible) field that has a readonly or external readonly type and thus any changes to the object can be *observed* by the referrer. The following example shows how `d` in `client` can observe a change to a readonly field `myC` because of the alias created inside the `action` method:

```
class C {int i;...}
class D {
  er C myC=new C(8);
  int inspect()er{return this.myC.i;}
  void action(r C c)m{
    int x=c.i*2;
    this.myC=c;
    //c.i=0;//Wrong since c is readonly
  }
  void client(m C c,m D d)m{
    c.i=1;
    d.action(c);//cannot modify c
    //here c.i==1 and d.inspect()==1
    c.i=2;
    /*here d.inspect()==2*/ }
}
```

2.2 External References

A mutable reference can be safely cast to an external reference but once treated as external, it can never be cast back to mutable again. This guarantees that the object (and its reachable object graph) will not be stored anywhere on the heap using this reference. In other words, no aliases can be created to an object using an external reference to it. However, an external reference does not prevent the object from mutating. If one were to guarantee that the object graph obtaining this external reference is fully disjoint from the current object graph, then none of these mutations can be observed in the current object graph. In fact, BI-JAVA provides such a guarantee in some circumstances. The following example shows how `action` can modify `c` but cannot store a reference to it, thus not affecting the state of `d` in `client` method:

```
...
void action(e C c)m{
  int x=c.i*2;
  //this.myC=c;//Wrong since c is external
  c.i=0;
}
void client(m C c, m D d)m{
  c.i=1;
  d.action(c);
  //here c.i==0 and d.inspect()==8
  c.i=2;
  /*here d.inspect()==8*/
}
...
```

2.3 External Readonly References

As this is the common supertype of all the reference types in BI-JAVA, it only allows reading the data from the object but not storing it in fields nor modifying it as the following example shows:

```
...
void action(er C c)m{
  int x=c.i*2;
  //this.myC=c;//Wrong since c is external
  //c.i=0;//Wrong since c is readonly
}
void client(m C c, m D d)m{
  c.i=1;
  d.action(c);
  //here c.i==1 and d.inspect()==8
  c.i=2;
  /*here d.inspect()==8*/ }
...
```

2.4 Loosely Closed Expressions

A closed expression is an expression without free variables. A fundamental concept in BI-JAVA is *loosely closed expression* — which is any expression where no free variables are mutable or readonly references. There are seven kinds of references in BI-JAVA. The loosely closed expressions do not contain mutable or readonly references to begin with and our type system prevents the introduction of aliases to fresh, balloon, external, and external readonly references. We refer to all the reference kinds that are *not* mutable or readonly as *loosely closed reference types*.

2.5 Permanent Type Promotion: Mutable into Fresh

A mutable reference produced by a loosely closed expression can be promoted to a fresh one. Indeed, any expression taking as input balloon, external or external readonly references will not be able to produce a mutable reference referring to objects reachable by the original references inside its reachable object graph.

All objects inside the produced mutable reference reachable object graph are either: (1) created inside such expression; (2) immutable; or (3) inside the *fresh* heap at the start of the expression execution.

From the property of loosely closed expressions it is clear that objects existing before such expression started executing cannot refer to those newly created objects. Hence, the produced mutable reference points to a fresh object.

Consequently, since a readonly reference refers to either a mutable object or an immutable one, and since fresh can be converted into immutable, readonly reference produced by a loosely closed expression can be promoted into an immutable one.

The following example shows how the mutable result of method `read1` is promoted to fresh in `usage1`. This is possible since `read1` only requires an external reference to a `Scanner` object. Symmetrically, method `usage2` converts a readonly reference to an immutable one.

```
class BarReader{
  m Bar read1(e Scanner s)er{...}
  f Bar usage1(m Scanner s)m{
    return this.read1(s);/*promotion happens here*/ }
  r Bar read2(e Scanner s)er{...}
  i Bar usage2(m Scanner s)m{
    return this.read2(s);/*promotion happens here*/ }
}
```

Fresh References The promotion mechanism we have explained is the only way to be able to use a fresh object in a program. For simplicity of formalisation our language prevents field access and field update over fresh references.²

Fresh references can be converted to any of the other kinds once and only once. Moreover, classes cannot declare fields of fresh type. Consider the following example:

²Conversions and permanent type promotions combined allow transparent field access and easy field update over fresh references.

```

void m(f Foo f1, f Foo f2)m{
  //NB! We are statically sure that f1!=f2
  f Foo f3=f1; //f1 is not accessible anymore
  i Foo f4=f3; //ok: conversion fresh ->immutable
  i Foo f5=f3; //ok: f3 is now Immutable
  b Foo f6=f2; //conversion fresh ->balloon
}

```

Thanks to our conversion mechanism, `fresh` objects can be used to bring into existence `immutable` and `balloon` objects. For simplicity, all the constructors return `mutable` objects, and the promotion system and the conversions are used to create the other kinds of objects.

2.6 Creation and Manipulation of Objects in BI-JAVA

Constructors always create `mutable` objects. In our simplified setting we consider only conventional constructors a la Featherweight Java. As usual, every field has to be initialised with a value of the corresponding type, but we have three exceptions: (1) `balloon` type fields have to be initialised with a `fresh` value; (2) `external` type fields have to be initialised with a `mutable` value; and (3) `external readonly` type fields have to be initialised with a `readonly` value.

Field update follows the same convention (e.g. `balloon` fields are updated with fresh values etc). Method calls are allowed over all kinds of receivers if the method has the correct modifier. Field access has to be treated more carefully and is presented in Rule (F-ACCESS-T) in Section 4.

The following examples shows how fresh conversions, permanent promotions, constructors and field access and update cooperates:

```

class Bar{//Shows how immutable works
  i Foo f;
  Bar(i Foo f) { this.f=f; }
  void m(i Foo f1,i Foo f2)m {
    m Bar b1=new Bar(f1);
    f1=f2; //variable f1 Immutable, not final
    b1.f=f2; //field f Immutable, not final
    i Bar b2=new Bar(f1);
    //b2.f=f1; //Wrong, b2 Immutable
  }
}

```

Immutable References Method `Bar.m` shows that some instances of `Foo` can be `mutable` while others can be `immutable`, and a permanent type promotion inside a fresh conversion can be used to initialize a new `immutable` object `b2`.

```

class Beer{//Shows how balloon works
  b Foo f;...
  void m(m Beer b1, f Foo f1)m{
    b1.f=f1;//ok conversion fresh->balloon
    e Beer b2=b1;//implicit cast mutable->external
    b Foo f2=b1.f;//mutable b1 exposes f as balloon
    e Foo f2=b2.f;//external b2 exposes f as external
  }
}

```

Balloon References Operation `b1.f=f1` converts the fresh object referred to by `f1` into a `balloon`, and thus the object and its whole reachable object graph moves from the fresh heap to the one of the current `Beer` instance.

2.7 Temporary Type Promotion

All `external readonly` references can refer to any object except for `fresh`, the main use of such reference is to make queries over such wide variety of objects in a uniform way. In order to perform such queries over many `external readonly` references, the creation of temporary data structures holding such references could be required.

Our type system allows any loosely closed expression to consider some of its open variables as `readonly`, if the result of this expression is loosely closed. This is safe because values containing such promoted references cannot be injected into pre existing objects.

For example, in the following code a `queryMethod` requires two `readonly` parameters, but the method usage only has a `balloon` and an `external readonly` parameter. Temporary type promotion can be applied, and the result of the method is `external readonly`.

```

class BarQuery{
  r Bar queryMethod(r Bar b1,r Bar b2)er{...}
  er Bar usage(b Bar b1,er Bar b2)er{
    return this.queryMethod(b1,b2);//promotion here
  }
}

```

2.8 Balloon Local Type Promotion

All `external` values can be modified. However, simply allowing field update will break the balloon invariant. Some operations are clearly valid: e.g. removing elements from an `external` list should be allowed. On the other hand, only `fresh` or `immutable` elements could be added. Such operations cannot break the balloon invariant.

In general, in order to modify an `external` variable, we need to consider it in the scope of its own balloon. Which is, in any loosely closed expression, one `external` variable can be considered `mutable` iff all the other `balloon` variables are considered `external` and the result is loosely closed. Indeed in BI-JAVA a top level expression is executed as outside of all the balloons, but subexpression where Balloon local type promotion is applied over variable `x` are executed “as inside of” the specific balloon that contains the object referred to by `x`.

This allows the programmer to modify in a very natural way the content of such balloon, without breaking the balloon invariant. Thanks to the property of loosely closed expressions, nothing from the outside can be injected inside that balloon. The result of this expression has to be loosely closed too, in order to avoid elements from the inside of the balloon to be injected in other balloons.

This typing discipline subsumes allowing `fresh` and `immutable` objects to be used for field update over `external` receivers. Indeed `fresh` and `immutable` are loosely closed reference types. Note that this mechanism applies also to the `balloon` object itself: such object is seen as `mutable` inside the balloon reachable object graph.

The following example shows how method `mutatorMethod` requires a `mutable` `b`, however the `usage` method only provides an `external` `b`. Thanks to balloon local type promotion, the method `mutatorMethod` can be called and the result is seen as `external` instead of `mutable`.

```

class BarMutator{
  m Bar mutatorMethod(m Bar b)er{ ... }
  e Bar usage(e Bar b)er{
    return this.mutatorMethod(b);//promotion here
  }
}

```

3. Real World Example

We can now exploit BI-JAVA in a more complete real world example: a two dimensional mesh editor, where meshes are composed of triangles and can be edited using a GUI:

```

class Point{
  float x; float y;
  Point(float x, float y){ this.x=x; this.y=y; }
  boolean isNear(er Point p, float d)er{...}
}
class Triangle{
  m Point p1; m Point p2; m Point p3;
  Triangle(m Point p1,m Point p2,m Point p3){
    this.p1=p1; this.p2=p2; this.p3=p3; }
  boolean isOverlap(er Triangle t)er{...}
}

```

As you can see, methods `isNear` and `isOverlap` can be annotated as `external readonly`. Indeed any method that just takes parameters to compute a result can always be annotated in this way. If

BI-JAVA was to be extensively used, we would expect to see many such methods in real applications. Depending on the specific computation involved, the returned value could be either a primitive value, a `fresh` or an `immutable` one.

3.1 Mesh, and Two Ways to Add a Triangle

We assume vectors for mutable points (`VectorMutableP`) and triangles (`VectorMutableT`):

```
class Mesh {
  m VectorMutableP points; m VectorMutableT triangles;
  Mesh(m VectorMutableP p, m VectorMutableT t) {
    this.points=p; this.triangles=t;
  }
  m Point seekPt1(float tolerance, m Point p1) m {
    for (m Point p2: this.points)
      if (p2.isNear(p1, tolerance)) return p2;
    this.points.add(p1);
    return p1;
  }
  void addTriangle1(float tolerance,
    m Point p1, m Point p2, m Point p3) m {
    this.triangles.add(new Triangle(
      seekPt1(tolerance, p1), seekPt1(tolerance, p2),
      seekPt1(tolerance, p3)));
  }
  m Point seekPt2(float tolerance, er Point p1) m {
    for (m Point p2: this.points)
      if (p2.isNear(p1, tolerance)) return p2;
    f Point p=new Point(p1.x, p1.y);
    this.points.add(p);
    return p;
  }
  void addTriangle2(float tolerance,
    er Point p1, er Point p2, er Point p3) m {
    this.triangles.add(new Triangle(
      seekPt2(tolerance, p1), seekPt2(tolerance, p2),
      seekPt2(tolerance, p3)));
  }
}
```

As you can see, methods `seekPt1` and `addTriangle1`, `seekPt2` and `addTriangle2`, show two different ways to add a triangle to a mesh: the first does not duplicate points objects, while the second one does. This shows how the type annotation of a method makes the aliasing contract of the method explicit: whenever a parameter is taken as `balloon`, `external` or `external readonly`, no new aliases to that parameter are created.

3.2 Adding Encapsulation

The `Mesh` class by itself does not provide any guarantee about how points and triangles are shared between different `Mesh` instances. It is possible to add useful restrictions while *using* such instances as the following example demonstrates. We assume a vector of balloon meshes (`VectorBalloonM`) and a vector of readonly triangles (`VectorReadOnlyT`):

```
class Workbench {
  b VectorBalloonM meshes;
  Workbench(f VectorBalloonM m) { this.meshes=m; }
  void addMesh() e { meshes.add(new Mesh()); }

  void addTriangleGui(int mIndex, e Gui gui) e {
    e Mesh mExt=this.meshes.get(mIndex);
    m Mesh m=mExt; //balloon local type promotion
    m.addTriangle1(gui.tolerance(), gui.selectPoint(),
      gui.selectPoint(), gui.selectPoint());
  }
  void mergeMeshes(int mIndex1, int mIndex2) e {
    e Mesh m2=this.meshes.get(mIndex2);
    m Mesh m1=this.meshes.get(mIndex1); //balloon local
    for (er Triangle t: m2.triangles)
      m1.addTriangle2(0.15f, t.p1, t.p2, t.p3);
  }
}
```

```
boolean isTriangleOverlap() er {
  f VectorReadOnlyT v= new VectorReadOnlyT();
  for (er m: this.meshes)
    for (er Triangle t: m.triangles)
      v.add(t); //v converted to mutable
      //t was promoted to readonly (temp promotion)
  while (!v.isEmpty()) {
    er Triangle t1=v.pop();
    for (er Triangle t2:v)
      if (t1.isOverlap(t2)) return true;
  }
  return false;
}
```

The constructor shows how `fresh` parameters can be used to initialise `Balloon` fields. This pattern allows to avoid the inefficient *defensive cloning or copying* of objects that is required in Java in order to correctly provide encapsulation [?] (Item 24).

The method `addTriangleGui` shows how to interact with a GUI to add a triangle. For clarity we show explicitly an `external` local variable `mExt` that is converted to a `mutable` one thanks to the `balloon` local type promotion. In method `mergeMeshes` we show that `balloon` local type promotion can be applied to an expression with no need for a local variable. As you can see, since the points of the triangle are provided by the `gui`, it is possible to use such points to construct the triangle, and thus method `Mesh.addTriangle1` is called. On the other hand, in `mergeMeshes` the points come from another mesh. Hence, it is not possible to use the same point and preserve the `balloon` invariant for the meshes, so the type system force us to call the `Mesh.addTriangle2` method instead.

Finally, method `isTriangleOverlap` checks if in the whole set of triangles, from all the meshes, there are two triangles that overlap. A natural way to do this check is to put all the triangles in a collection and then consume it while checking for a couple of triangles, `t1` and `t2` such that `t1.isOverlap(t2)`. All of those triangles come from different, individually encapsulated meshes. In order to put all of them into a single collection we have to use the temporary type promotion.

3.3 Reading Meshes from Files

The following code shows how to read a `fresh` mesh using Java `Scanner` object. As you can see, a mesh is quite complex object, with aliasing involved, but our type system is able to verify that the produced objects are `fresh`.

```
class Reader {
  f Point readPoint(e Scanner s) er {
    return new Point(s.nextFloat(), s.nextFloat());
  }
  f Shape readMesh(e Scanner s) er {
    m Mesh m=new Mesh(
      new VectorMutableP(), new VectorMutableT());
    while (m.hasNextFloat())
      m.addTri1(0.15f, this.readPoint(s),
        this.readPoint(s), this.readPoint(s));
    return m;
  }
}
```

4. Formalisation

First we present a formal definition for BI-JAVA without any support for parallelisation. We show syntax, reduction and typing. Typing is composed of conventional subtyping relation (found in the Technical Report [16]), `fresh` analysis, conventional typing and promotion rules. After defining the observable behaviour of BI-JAVA, in Section 5 we show how to transparently optimise a program using parallelisation.

4.1 Syntax

In Figure 2 we show the syntax of BI-JAVA. We assume countably infinite sets of variables x , class or interface names C , method

p	$::= \overline{cd} \overline{id}$	program
cd	$::= \mathbf{class} \ C \ \mathbf{implements} \ \overline{C} \ \{ \overline{fd} \ \overline{md} \}$	class declaration
id	$::= \mathbf{interface} \ C \ \mathbf{extends} \ \overline{C} \ \{ \overline{mh}; \}$	interface declaration
fd	$::= T \ f;$	field declaration
mh	$::= T \ m \ (\overline{T} \ x) \ \mathcal{M}$	method header
md	$::= mh \ e$	method declaration
e	$::= x \mid e.m(\overline{e}) \mid \mathbf{new} \ C(\overline{e})$ $\mid e_0.f \mid e_0.f=e_1 \mid T \ x=e_0; \ e_1$	expression
T	$::= \mathcal{M} \ C$	type
\mathcal{M}	$::= \mathbf{f} \mid \mathbf{b} \mid \mathbf{m} \mid \mathbf{i} \mid \mathbf{e} \mid \mathbf{r} \mid \mathbf{er}$	type modifiers

Figure 2. Syntax

names m , and field names f . As in FJ [9] variables include the special variable **this**.

A program p is a set of class and interface declarations. A class declaration consists of a class name followed by the set of implemented interfaces, the sequence of field declarations and the set of method declarations. We assume conventional constructors as in FJ to be implicitly declared. To keep the presentation focused on our type modifiers and their application to parallelisation, we do not consider class composition operators like the *extends* operator in Java. We showed in [7, 10] how expressive composition operators (subsuming inheritance) can be added to Java-like languages.

An interface declaration consists of an interface name followed by the sets of extended interfaces and method headers. Field declarations are as in FJ. Method declarations are composed by a method header and a body. The method header is as in FJ. However since here types are richer, we complete the method header with the type modifier for the implicit parameter **this**. We write the type modifier for **this** after the parameters list, following the syntactic convention of C++ **const** modifier.

As in FJ, method bodies are simply ordinary expressions. Expressions can be variables, method or constructor calls, field access, field update and local variable declarations.

Variables can be declared in method headers or inside expressions of the form $T \ x = e_0; \ e_1$; that is equivalent to a **let** in functional languages. For simplicity, an expression is well-formed only if the same variable is not declared twice.

In method bodies we omit curly brackets and the keyword **return** inside the formalisation. In the examples we place the keyword inside the top level expression, but after any local variable declaration.

In the sequences of field declaration, parameter declaration and method or constructor calls the order is relevant. In all the other sequences the order is irrelevant, and thus, they can be considered as sets (or maps).

Types are composed of two components: a type modifier \mathcal{M} and a class or interface name C . A modifier \mathcal{M} have seven possible values: **f** (for **f**resh), **b** (for **b**alloon), **m** (for **m**utable), **i** (for **i**mmutable), **e** (for **e**xternal), **r** (for **r**eadonly) and **er** (for **e**xternal **r**eadonly).

4.2 Reduction

Figure 3 shows reduction rules for BI-JAVA (without parallelism). In order to show reduction, we have to enrich the expressions with object identifiers ι as run time expressions. As usual, the heap is modelled as a finite map from object identifiers to records annotated with a class name.

Reduction is an arrow over pairs consisting of a heap and an expression. To improve readability we will mark in grey the heap part. During reduction modifiers \mathcal{M} are completely ignored.

Rule (CTX) is standard. Rule (F-ACCESS) extracts the value of field f . We use the notation $\mu(\iota).f$ for the value of field f of object ι . Rule (CONSTRUCTOR) reduces constructor invocations.

	$e ::= \dots \mid \iota$	expressions (also run time)
	$\mu ::= \iota \mapsto C(f_1 = \iota_1, \dots, f_n = \iota_n)$	memory
	$\mathcal{E} ::= \square \mid \mathcal{E}.m(\overline{e}) \mid \iota.m(\overline{\iota}, \mathcal{E}, \overline{e})$ $\mid \mathbf{new} \ C(\overline{\iota}, \mathcal{E}, \overline{e}) \mid T \ x = \mathcal{E}; \ e$ $\mid \mathcal{E}.f \mid \mathcal{E}.f = e \mid \iota.f = \mathcal{E}$	context
	$\mu_1 \mid e_1 \rightarrow \mu_2 \mid e_2$	
(CTX)	$\frac{\mu_1 \mid e_1 \rightarrow \mu_2 \mid e_2}{\mu_1, \mu \mid \mathcal{E}[e_1] \rightarrow \mu_2, \mu \mid \mathcal{E}[e_2]}$	
(F-ACCESS)	$\frac{\mu \mid \iota_0.f \rightarrow \mu \mid \iota_1}{\text{with } \mu(\iota_0).f = \iota_1}$	
(CONSTRUCTOR)	$\frac{\mu \mid \mathbf{new} \ C(\iota_1, \dots, \iota_n) \rightarrow \mu, \iota \mapsto C(f_1 = \iota_1, \dots, f_n = \iota_n) \mid \iota}{\text{with } \iota \notin \text{dom}(\mu) \text{ and } p(C) = \mathbf{class} \ C \ \mathbf{implements} \ \{ _f_1; \dots; _f_n; \overline{md} \}}$	
(M-INVK)	$\frac{\mu \mid \iota_0.m(\iota_1, \dots, \iota_n) \rightarrow \mu \mid e[x_1 = \iota_1, \dots, x_n = \iota_n, \mathbf{this} = \iota_0]}{\text{with } \mu(\iota) = C(_) \text{ and } p(C).m = _m(_, \dots, _x_n) _e}$	
(V-DEC)	$\mu \mid T \ x = \iota; \ e \rightarrow \mu \mid e[x = \iota]$	
(F-UPDATE)	$\mu \mid \iota_0.f = \iota_1 \rightarrow \mu[\iota_0.f = \iota_1] \mid \iota_1$	

Figure 3. Reduction

Rule (M-INVK) models a conventional method call. We assume a fixed program p and we use notation $p(C).m$ to extract the method declaration. We use the notation $e[x_1 = \iota_1, \dots, x_n = \iota_n]$ for variable substitution, that is, we simultaneously replace all the occurrences of x_i in e with ι_i . Finally, Rules (V-DEC) and (F-UPDATE) are straightforward. We use notation $\mu[\iota_0.f = \iota_1]$ to override the heap location $\iota_0.f$ with the value ι_1 .

4.3 Fresh Analysis and Γ

Figure 4 shows syntax for extended types, type environments and the split relation $\Gamma = \Gamma_1 \circ \dots \circ \Gamma_n$ as explained below.

Free variables are typed using a variable environment Γ . The treatment is conventional for all variable kinds but **f**resh ones. If a **f**resh reference is not converted to another kind of reference, it can be used only one time: indeed typing rules propagate that variable to only one sub expression. If a **f**resh reference is converted to an **i**mmutable or **b**alloon one, then it can be used in many points in the expression, and thus it is treated as a normal variable of that kind. In many cases, if a **f**resh reference is converted to a **m**utable one, it is simply considered **m**utable, and it can be used at many points in the expression. However, if a **f**resh reference is converted to a **m**utable one inside a promotion expression, the expression outside that promotion can or cannot see that variable depending on the kind of promotion.

We add to the type modifiers the production $\mathbf{m}\langle \overline{TP} \rangle$ (where the metavariable \overline{TP} ranges over the three different promotion kinds) with the following meaning: if $\Gamma(x) = \mathbf{m}\langle \overline{TP} \rangle C$ then x refers to a fresh object that is going to be converted into **m**utable after entering the sequence (where order is relevant) of nested promotion \overline{TP} . Thus \mathbf{m} is just a syntactic sugar for $\mathbf{m}\langle \emptyset \rangle$.

The programmer cannot use types of the form $\mathbf{m}\langle \overline{TP} \rangle C$ with non empty \overline{TP} . They are only used internally by the type system for typing expressions with many nested promotions.

$\mathcal{M} ::= \dots \mid \overline{\mathfrak{m}}(\overline{TP})$ modifiers (also tracing fresh)
 $TP ::= \text{PTP} \mid \text{ITP} \mid \text{BLTP}$ promotion kinds
 $\Gamma ::= x:T$ variable environment

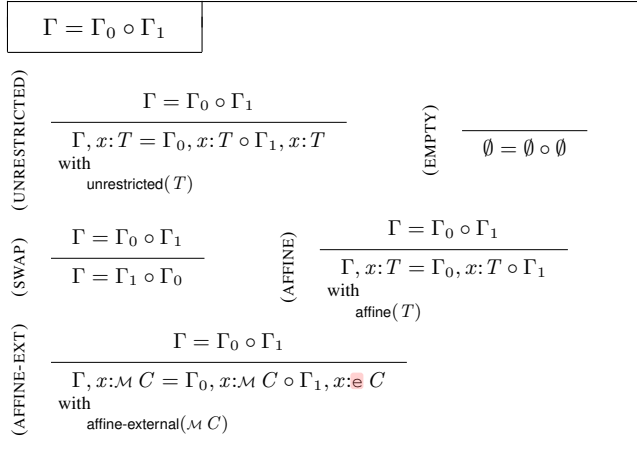


Figure 4. Variable environment split

In order to formalise the precise way \mathfrak{f} resh and $\overline{\mathfrak{m}}(\overline{TP})$ references have to be used, we employ a substructural type system [17], very similar to an affine type system. In detail, we define three different treatments for variables:

- unrestricted variables can be used any number of times in an expression; all the non \mathfrak{f} resh variables are unrestricted; \mathfrak{f} resh variable converted in $\overline{\mathfrak{m}}$ utable inside only temporary type promotions ($\overline{\mathfrak{m}}(\overline{TP})$) are also unrestricted. Formally $\text{unrestricted}(\mathcal{M} C)$ holds iff $\mathcal{M} \in \{\mathfrak{b}, \overline{\mathfrak{m}}(\overline{TP}), \mathfrak{i}, \mathfrak{e}, \mathfrak{r}, \mathfrak{er}\}$;
- affine variables can be used at most one time in an expression; \mathfrak{f} resh variables are affine, \mathfrak{f} resh variable converted in $\overline{\mathfrak{m}}$ utable inside at least one permanent type promotions are also affine. Formally $\text{affine}(\mathcal{M} C)$ holds iff $\mathcal{M} = \mathfrak{f}$ or $\mathcal{M} = \overline{\mathfrak{m}}(\overline{TP})$ and $\text{PTP} \in \overline{TP}$;
- last, we introduce the concept of affine external variables, that can be used at most one time with their type, but many time with the corresponding \mathfrak{e} external type. Formally $\text{affine-external}(\mathcal{M} C)$ holds iff $\mathcal{M} = \overline{\mathfrak{m}}(\overline{TP})$ and $\text{BLTP} \in \overline{TP}$ but $\text{PTP} \notin \overline{TP}$.

Now we can provide a generic *Context split* relation $\Gamma = \Gamma_0 \circ \Gamma_1$ as in [17], mapping a single variable environment into an unordered pair (see Rule (SWAP)) of variable environments.

Rule (UNRESTRICTED) states that an unrestricted variable can be used many times. Rule (AFFINE) states that an affine variable will be contained in only one of the two environments.

Finally, rule (AFFINE-EXT) states that an affine external variable will be contained in one of the two environments with its type, while the other one will contain that variable with the corresponding \mathfrak{e} external type.

The non-deterministic splitting relation $\Gamma = \Gamma_0 \circ \Gamma_1$ does not provide a direct guide for an efficient implementation of our type system. We believe that the implementation approach described in [17] can be easily adapted to our case.

4.4 Conventional Typing

Figure 5 shows straightforward rules for well formedness of classes and interfaces. For any well-typed program all classes and interfaces are valid. Rule (CLASS) validates a class if all methods are valid and if the interfaces \overline{C} are correctly implemented; that is, for all methods of all the implemented interfaces, a method with an analogous header is declared in the class C_0 . Note how methods are

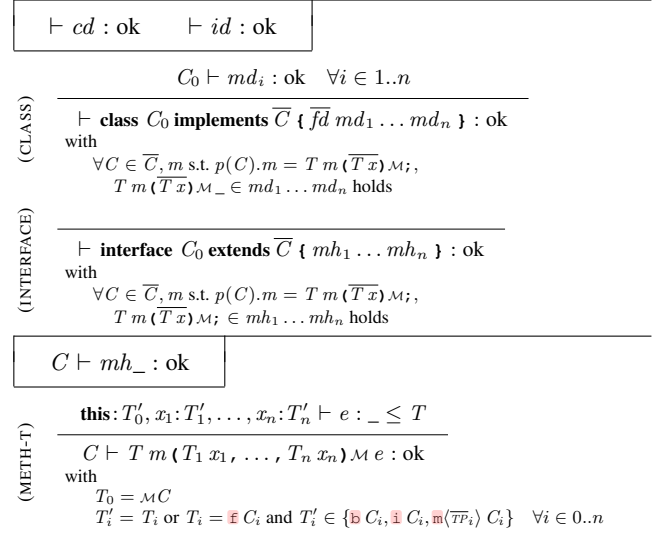


Figure 5. Typing for Classes and Methods

validated in the context of their class. Similarly, rule (INTERFACE) validates an interface if the interfaces \overline{C} are correctly implemented; that is, for all the method headers of all the implemented interfaces, an analogous method header is declared in the interface C_0 .

Rule (METH-T) derives the type for **this** from the class name C and the type modifier \mathcal{M} . Note how this rule introduces the \mathfrak{f} resh variables inside Γ : they can be non-deterministically introduced as \mathfrak{f} resh or as any of the possible conversion for a \mathfrak{f} resh variable. That is, the conversion is executed non-deterministically when the variable enters in scope. In the premise we use notation $\Gamma \vdash e : T_1 \leq T_2$ as a shortcut for $\Gamma \vdash e : T_1$ and $T_1 \leq T_2$.

Figure 6 shows conventional typing rules for expressions. Rule (VAR-T) is conventional.

Rule (F-ACCESS-T) is conventional but uses function $\text{f-access}_-(_)$ to provide the correct modifier to the resulting value. Formally:

$\text{f-access}_{\overline{\mathfrak{m}}}(\mathcal{M}) = \mathcal{M} \quad \text{f-access}_{\mathfrak{i}}(_) = \mathfrak{i} \quad \text{f-access}_{\mathfrak{i}}(\mathfrak{i}) = \mathfrak{i}$
 $\text{f-access}_{\mathfrak{b}}(\mathcal{M}) = \text{f-access}_{\mathfrak{e}}(\mathcal{M})$
 $\text{f-access}_{\mathfrak{e}}(\mathcal{M}) = \mathfrak{e}$ with $\mathcal{M} \in \{\mathfrak{b}, \mathfrak{e}, \overline{\mathfrak{m}}\}$
 $\text{f-access}_{\mathfrak{e}}(\mathcal{M}) = \mathfrak{er}$ with $\mathcal{M} \in \{\mathfrak{r}, \mathfrak{er}\}$
 $\text{f-access}_{\mathfrak{r}}(\mathcal{M}) = \mathfrak{r}$ with $\mathcal{M} \in \{\mathfrak{m}, \mathfrak{r}\}$
 $\text{f-access}_{\mathfrak{r}}(\mathcal{M}) = \mathfrak{er}$ with $\mathcal{M} \in \{\mathfrak{b}, \mathfrak{e}, \mathfrak{er}\}$
 $\text{f-access}_{\mathfrak{er}}(\mathcal{M}) = \mathfrak{er}$ with $\mathcal{M} \in \{\mathfrak{b}, \mathfrak{m}, \mathfrak{e}, \mathfrak{r}, \mathfrak{er}\}$

Rule (M-INVK-T) is conventional. Note how here and in the following we use relation $\Gamma = \Gamma_0 \circ \dots \circ \Gamma_n$.

Rules (F-UPDATE-T) and (F-UPDATE-TB) manage field update. Rule (F-UPDATE-T) covers the case of non \mathfrak{b} alloon fields. In this case to force correct field update it is enough to state that the assigned value is not \mathfrak{b} alloon, \mathfrak{e} external or \mathfrak{e} external \mathfrak{r} eadonly. Indeed, references of these kinds cannot be assigned. Rule (F-UPDATE-TB) covers the case of \mathfrak{b} alloon fields. It requires the assigned value to be \mathfrak{f} resh. In any case the modified value has to be $\overline{\mathfrak{m}}$ utable. Promotion can be used to modify \mathfrak{e} external values.

Rule (NEW-T) is conventional, but, just like the field update, has to ensure that all the \mathfrak{b} alloon fields are initialised with \mathfrak{f} resh values, and that non assignable values are not assigned.

Rule (V-DEC) is conventional but, as for (METH-T), takes care of the fresh conversion.

4.5 Promotions Rules

First, we define a function $\text{l-close}(_)$ that returns a loosely closed type for any reference kind as follows:

$\Gamma \vdash e : T$	
(VAR-T)	$\frac{\Gamma \vdash x : \Gamma(x)}{\Gamma \vdash e_0 : \mathcal{M}_0 C_0}$
(F-ACCESS-T)	$\frac{\Gamma \vdash e_0.f : \text{f-access}_{\mathcal{M}_0(\mathcal{M})} C}{\Gamma \vdash e_0.f : \text{f-access}_{\mathcal{M}_0(\mathcal{M})} C}$ <p style="text-align: center; margin-top: -10px;">with $p(C_0) = \text{class } C_0 \text{ implements } \{ _ \mathcal{M} C f ; _ \}$</p>
(M-INVK-T)	$\frac{\Gamma_0 \vdash e_0 : _ C_0 \leq \mathcal{M} C_0}{\Gamma_0 \circ \dots \circ \Gamma_n \vdash e_0.m(e_1, \dots, e_n) : T}$ <p style="text-align: center; margin-top: -10px;">with $p(C_0).m = T m(T_1 x_1, \dots, T_n x_n) \mathcal{M} _$</p>
(F-UPDATE-T)	$\frac{\Gamma_0 \vdash e_0 : \bar{\mathfrak{m}} C_0}{\Gamma_0 \circ \Gamma_1 \vdash e_0.f = e_1 : \mathcal{M} C}$ <p style="text-align: center; margin-top: -10px;">with $p(C_0) = \text{class } C_0 \text{ implements } \{ _ \mathcal{M}_1 C_1 f ; _ \}$ $\mathcal{M} \notin \{ \bar{\mathfrak{b}}, \bar{\mathfrak{e}}, \bar{\mathfrak{e}\mathfrak{r}} \}$</p>
(F-UPDATE-TB)	$\frac{\Gamma_0 \vdash e_0 : \bar{\mathfrak{m}} C_0}{\Gamma_0 \circ \Gamma_1 \vdash e_0.f = e_1 : \mathcal{M} C}$ <p style="text-align: center; margin-top: -10px;">with $p(C_0) = \text{class } C_0 \text{ implements } \{ _ \bar{\mathfrak{b}} C_1 f ; _ \}$</p>
(NEW-T)	$\frac{\Gamma_i \vdash e_i : \mathcal{M}_i C_i \leq \mathcal{M}'_i C'_i \quad \forall i = 1..n \text{ where } \mathcal{M}'_i \neq \bar{\mathfrak{b}}}{\Gamma_1 \circ \dots \circ \Gamma_n \vdash \text{new } C(e_1 \dots e_n) : \bar{\mathfrak{m}} C}$ <p style="text-align: center; margin-top: -10px;">with $p(C) = \text{class } C \text{ implements } \{ \mathcal{M}'_1 C'_1 _ ; \dots ; \mathcal{M}'_n C'_n _ ; \bar{m}d \}$ $\mathcal{M}_i \notin \{ \bar{\mathfrak{b}}, \bar{\mathfrak{e}}, \bar{\mathfrak{e}\mathfrak{r}} \} \quad \forall i = 1..n$</p>
(V-DEC-T)	$\frac{\Gamma_0 \vdash e_0 : \mathcal{M} C_0 \leq T_0}{\Gamma_1, x:T \vdash e_1 : T_1}$ <p style="text-align: center; margin-top: -10px;">with $\Gamma_0 \circ \Gamma_1 \vdash T_0 x = e_0 ; e_1 : T_1$ either $T = T_0$ or $T_0 = \bar{\mathfrak{f}} C$ and $T \in \{ \bar{\mathfrak{b}} C, \bar{\mathfrak{a}} C, \bar{\mathfrak{m}}(\bar{\mathfrak{T}\mathfrak{P}}) C \}$</p>

Figure 6. Type system expressions

$$\begin{aligned} \text{l-close}(\bar{\mathfrak{m}} C) &= \bar{\mathfrak{e}} C \\ \text{l-close}(\bar{\mathfrak{r}} C) &= \bar{\mathfrak{e}\mathfrak{r}} C \\ \text{l-close}(T) &= T \text{ otherwise} \end{aligned}$$

Rules (PERMANENT1-TP) and (PERMANENT2-TP) define permanent type promotions: as you can see $\bar{\mathfrak{m}}$ utable results can be promoted into $\bar{\mathfrak{f}}$ resh while $\bar{\mathfrak{r}}$ eadonly results can be promoted into $\bar{\mathfrak{a}}$ mmutable. We use function $\text{l-close-p}(\Gamma)$, that behave like $\text{l-close}(\Gamma)$ but is extended in order to manage the types of the form $\bar{\mathfrak{m}}(\bar{\mathfrak{P}\mathfrak{T}\mathfrak{P}}, \bar{\mathfrak{T}\mathfrak{P}}) C$.

Formally: $\text{l-close-p}(\bar{\mathfrak{m}}(\bar{\mathfrak{P}\mathfrak{T}\mathfrak{P}}, \bar{\mathfrak{T}\mathfrak{P}}) C) = \bar{\mathfrak{m}}(\bar{\mathfrak{T}\mathfrak{P}}) C$ and $\text{l-close-p}(T) = \text{l-close}(T)$ otherwise.

Rule (TEMPORARY-TP) models temporary type promotion. As you can see, many variable of (a supertype of) an $\bar{\mathfrak{e}}$ xternal $\bar{\mathfrak{r}}$ eadonly type can be promoted to $\bar{\mathfrak{r}}$ eadonly type. We use function $\text{l-close-t}(\Gamma)$, that behave like $\text{l-close}(\Gamma)$ but is extended in order to manage the types of the form $\bar{\mathfrak{m}}(\bar{\mathfrak{T}\mathfrak{P}}, \bar{\mathfrak{T}\mathfrak{P}}) C$.

Formally: $\text{l-close-t}(\bar{\mathfrak{m}}(\bar{\mathfrak{T}\mathfrak{P}}, \bar{\mathfrak{T}\mathfrak{P}}) C) = \bar{\mathfrak{m}}(\bar{\mathfrak{T}\mathfrak{P}}) C$ and $\text{l-close-t}(T) = \text{l-close}(T)$ otherwise.

Rule (BALLOON-L-TP) models balloon local type promotion. It promotes a variable x_0 to $\bar{\mathfrak{m}}$ utable if it is (a supertype of) an $\bar{\mathfrak{e}}$ xternal reference. We use function $\text{l-close-b}(\Gamma)$, that behave like $\text{l-close}(\Gamma)$ but is extended in order to manage the types of the form $\bar{\mathfrak{m}}(\bar{\mathfrak{B}\mathfrak{L}\mathfrak{T}\mathfrak{P}}, \bar{\mathfrak{T}\mathfrak{P}}) C$, and to ensure that $\bar{\mathfrak{b}}$ alloon references are converted

$\Gamma \vdash e : T$	
(PERMANENT1-TP)	$\frac{\text{l-close-p}(\Gamma) \vdash e : \bar{\mathfrak{m}} C}{\Gamma \vdash e : \bar{\mathfrak{f}} C}$
(PERMANENT2-TP)	$\frac{\text{l-close-p}(\Gamma) \vdash e : \bar{\mathfrak{r}} C}{\Gamma \vdash e : \bar{\mathfrak{a}} C}$
(TEMPORARY-TP)	$\frac{\text{l-close-t}(\Gamma), x_1:\bar{\mathfrak{r}} C_1, \dots, x_n:\bar{\mathfrak{r}} C_n \vdash e : \mathcal{M} C}{\Gamma, x_1:\mathcal{M}_1 C_1, \dots, x_n:\mathcal{M}_n C_n \vdash e : \text{l-close}(\mathcal{M}) C}$ <p style="text-align: center; margin-top: -10px;">with $\mathcal{M}_i C_i \leq \bar{\mathfrak{e}\mathfrak{r}} C \quad \forall i = 1..n$</p>
(BALLOON-L-TP)	$\frac{\text{l-close-b}(\Gamma), x_0:\bar{\mathfrak{m}} C_0 \vdash e : \mathcal{M} C}{\Gamma, x_0:\mathcal{M}_0 C_0 \vdash e : \text{l-close}(\mathcal{M}) C}$ <p style="text-align: center; margin-top: -10px;">with $\mathcal{M}_0 C_0 \leq \bar{\mathfrak{e}} C_0$</p>
(SUB-EXP-T)	$\frac{\Gamma \vdash e_0 : T_0}{\Gamma, x_0:T'_0 \vdash \mathcal{E}[x_0] : T}$ <p style="text-align: center; margin-top: -10px;">with either $T'_0 = T_0$ or $T'_0 = \bar{\mathfrak{f}} C$ and $T \in \{ \bar{\mathfrak{b}} C_0, \bar{\mathfrak{a}} C_0, \bar{\mathfrak{m}}(\bar{\mathfrak{T}\mathfrak{P}}) C_0 \}$</p>

Figure 7. Type system special

in $\bar{\mathfrak{e}}$ xternal ones.

Formally: $\text{l-close-b}(\bar{\mathfrak{m}}(\bar{\mathfrak{B}\mathfrak{L}\mathfrak{T}\mathfrak{P}}, \bar{\mathfrak{T}\mathfrak{P}}) C) = \bar{\mathfrak{m}}(\bar{\mathfrak{T}\mathfrak{P}}) C$, $\text{l-close-b}(\bar{\mathfrak{b}} C) = \bar{\mathfrak{e}} C$ and $\text{l-close-b}(T) = \text{l-close}(T)$ otherwise.

Note how in all the three variation of $\text{l-close}(_)$ the other kinds of modifiers of the form $\bar{\mathfrak{m}}(\bar{\mathfrak{T}\mathfrak{P}})$ are undefined. This ensure that the rules are applicable only if such modifiers are not involved.

Finally, rule (SUB-EXP-T) allows to see any sub expression as a local variable, removing from the programmer the need of declare a local variable only to please the type system.

4.6 Soundness

In order to define Soundness, we need to define concept of stack, well formedness of the memory in a stack, and how to type a run time expression. Figure 8 defines a stack (Σ) as a map from object identifiers to types. The stack associated with a run time expression maps all the identifiers appearing inside such expression to a type.

Rule (LOC-T) states that if an expression e with some free variable is well typed, then the same expression with some variable replaced with object identifier is well typed too, as long as the stack Σ is rich enough to provide the type information for all the object identifiers introduced inside e .

Rule (LOC-T) models typing of object identifiers and allows us to design the rest of the type system for the language without object identifiers, so that we can avoid tedious duplication of the definition of Γ split operator. Moreover, promotions are simpler in the language without object identifiers. In a conventional language rule (LOC-T) would make the proof non trivial, however in BI-JAVA the same issues are already caused by (SUB-EXP-T). A careful reader could notice how rule (LOC-T) is just a specialisation of a set of applications of (SUB-EXP-T) over subexpression of the form ι .

We introduce typing rule (MEM-T) stating when a memory is well formed under a stack. It uses judgement $\mu \vdash \iota$ to check all object identifiers to be well formed with respect to the memory. Formally $\mu \vdash \iota$ holds iff $\mu(\iota) = C(f_1 = \iota_1 \dots f_n = \iota_n), \forall i \in 1..n$
 $\mu(\iota_i) = C_i(_), p(C).f_i = \mathcal{M} C'_i$ and $\mathcal{M} C_i \leq \mathcal{M} C'_i$
As you can see this judgement does not check that location ι actually refers to an object compatible with the type modifiers \mathcal{M} .

In order to define validity of immutable ($\text{valid-}\mu(\iota_0)$), fresh ($\text{valid-}\mathfrak{F}_\mu^\Sigma(\iota_0)$) and balloon ($\text{valid-}\mathfrak{B}_\mu^\Sigma(\iota_0)$) invariant we introduce the following *link* notation:

$\Sigma ::= \overline{\iota:T}$ stack

$\Gamma; \Sigma \vdash e : T$	
(LOC-T) $\frac{\Gamma, x_1:T_1 \dots x_n:T_n \vdash e : T}{\Gamma; \iota_1:T_1 \dots \iota_n:T_n \vdash e[x_1 = \iota_1 \dots x_n = \iota_n] : T}$	
$\Sigma \vdash \mu$	
(MEM-T) $\frac{\Sigma \vdash \mu \text{ with } \begin{array}{l} \Sigma = \iota_1:T_1 \dots \iota_n:T_n \\ \mu \vdash \iota \quad \forall \iota \in \text{dom}(\mu) \\ \text{valid-}I_\mu(\iota) \quad \forall \iota \in \{\iota_i \mid i \in 1..n, T_i = \underline{\mathfrak{I}}\} \\ \text{valid-}F_\mu^\Sigma(\iota) \quad \forall \iota \in \{\iota_i \mid i \in 1..n, \text{affine}(T_i)\} \\ \text{valid-}B_\mu(\iota) \quad \forall \iota \in \text{BC}_\mu^\Sigma \\ \text{rog-}I_\mu(\iota_1) \cap \text{rog-}I_\mu(\iota_2) = \emptyset \quad \forall \iota_1, \iota_2 \in \text{RB}_\mu(\Sigma), \iota_1 \neq \iota_2 \\ \text{rog-}I_\mu(\iota_1) \cap \text{SH}_\mu^\Sigma \neq \emptyset \text{ implies } \text{rog-}I_\mu(\iota_2) \cap \text{SH}_\mu^\Sigma = \emptyset \quad \forall \iota_1, \iota_2 \in \text{BC}_\mu^\Sigma, \iota_1 \neq \iota_2 \end{array}}{\Sigma \vdash \mu}$	
(MEM-MEM-OK) $\frac{\Sigma_0 \rightsquigarrow \Sigma_1 \vdash \mu_0 \rightsquigarrow \mu_1}{\Sigma, \Sigma_0 \rightsquigarrow \Sigma, \Sigma_1 \vdash \mu_0 \rightsquigarrow \mu_1}$ with $\begin{array}{l} \text{dom}(\Sigma_0) \cap \text{dom}(\Sigma_1) = \emptyset \\ \text{IH}_{\mu_0} \subseteq \text{IH}_{\mu_1} \\ \mu_0(\iota) = \mu_1(\iota) \quad \forall \iota \in \text{IH}_{\mu_0} \\ \text{rog-}I_{\mu_0}(\iota) \cap \text{rog-}I_{\mu_0}(\iota') = \emptyset \\ \text{implies } \text{rog-}I_{\mu_1}(\iota) \cap (\text{rog-}I_{\mu_1}(\iota') \cup \text{rog-}I_{\mu_0}(\iota')) = \emptyset \quad \forall \iota, \iota' \in \text{BC}_{\mu_0}^\Sigma \end{array}$	

Figure 8. Typing memory and locations

- $\iota_1.f \xrightarrow{\overline{M}} \iota_2$ iff $\mu(\iota_1) = C(_f = \iota_2)$ and $p(C).f = M_$ that reads as “there is a link named f of kind $M \in \overline{M}$ between ι_1 and ι_2 ” For example $\iota_1.f \xrightarrow{\underline{\mathfrak{b}}} \iota_2$ means that ι_1 have a field f of kind $\underline{\mathfrak{b}}$ balloon referring to ι_2 .

We generalise the link notion in the following way:

- If we omit the field name, we refer of a generic reference without keeping in account the field name, as in $\iota_1 \xrightarrow{\underline{\mathfrak{i}}, \underline{\mathfrak{r}}} \iota_2$ that reads as “there is a link of kind $\underline{\mathfrak{i}}$ or $\underline{\mathfrak{r}}$ between ι_1 and ι_2 ”
- If we omit \overline{M} then we refer to a generic reference without keeping in account the kind of the reference, as in $\iota_1 \xrightarrow{_} \iota_2$ iff $\mu(\iota_1) = C(_f = \iota_2)$ that reads as “there is a link between ι_1 and ι_2 ”.
- We use $*$ to indicate reflexive and transitive closure, as usual. For example $\iota_1 \xrightarrow{*\underline{\mathfrak{m}}, \underline{\mathfrak{r}}} \iota_2$ reads as “there is a path composed by links of kind $\underline{\mathfrak{m}}$ or $\underline{\mathfrak{r}}$ between ι_1 and ι_2 ”

Now we can introduce the $\underline{\mathfrak{i}}$ immutable invariant:

$\text{valid-}I_\mu(\iota_0)$ iff $\forall f, \iota_1 \in \text{rog}_\mu(\iota_0), \iota_2 \in \text{dom}(\mu) \setminus \text{rog}_\mu(\iota_0), (\iota_2.f \xrightarrow{_} \iota_1 \text{ implies } \iota_2.f \xrightarrow{\underline{\mathfrak{i}}, \underline{\mathfrak{r}}, \underline{\mathfrak{e}}, \underline{\mathfrak{e}}^*} \iota_1)$ where $\text{rog}_\mu(\iota)$ denotes the reachable object graph of ι .

As you can see, an object is immutable if all the reference to it in the heap are a supertype of $\underline{\mathfrak{i}}$ immutable. The type of circular references inside the immutable object graph are not relevant ($\iota_2 \in \text{dom}(\mu) \setminus \text{rog}_\mu(\iota_0)$).

The immutable heap IH_μ is simply composed by all the valid immutable objects: $\text{IH}_\mu = \{\iota \in \text{dom}(\mu) \mid \text{valid-}I_\mu(\iota)\}$

Now we can introduce the $\underline{\mathfrak{f}}$ fresh invariant:

$\text{valid-}F_\mu^\Sigma(\iota_0)$ iff $\forall f, \iota_1 \in \text{rog-}I_\mu(\iota_0) (\iota_1.f \xrightarrow{_} \iota_0 \text{ implies } \iota_1.f \xrightarrow{\underline{\mathfrak{m}}, \underline{\mathfrak{e}}, \underline{\mathfrak{r}}, \underline{\mathfrak{e}}^*} \iota_0)$ and $\forall \iota_2 \in \text{rog-}I_\mu(\iota_0), \iota_3 \in \text{dom}(\mu) (\iota_3 \xrightarrow{_} \iota_2 \text{ implies } \iota_3 \in \text{rog-}I_\mu(\iota_0) \text{ or } \text{gc}_\mu^\Sigma(\iota_3))$

with $\text{rog-}I_\mu(\iota_0) = \{\iota \in \text{rog}_\mu(\iota_0) \mid \text{not valid-}I_\mu(\iota)\}$ and predicate $\text{gc}_\mu^\Sigma(\iota)$ holds if ι is ready to be garbage collected.

As you can see, an object is fresh if all the reachable object graph (using $\text{rog-}I_\mu(\iota_0)$ to avoid keeping in consideration immutable objects) is only internally aliased ($\iota_3 \in \text{rog-}I_\mu(\iota_0)$), and all the (internal) links to the fresh object itself are of $\underline{\mathfrak{m}}$ mutable kind or subtype thereof.

The fresh heap FH_μ^Σ is composed by all the fresh objects and their reachable object graphs

$\text{FH}_\mu^\Sigma = \{\iota_0 \in \text{rog-}I_\mu(\iota_1) \mid \iota_1 \in \text{dom}(\mu), \text{valid-}F_\mu^\Sigma(\iota_1)\}$

An object ι_0 is a valid balloon, that is $\text{valid-}B_\mu^\Sigma(\iota_0)$ iff

- at most one $\underline{\mathfrak{b}}$ balloon reference exists to the balloon object, and is outside the balloon itself: $\forall f_1, f_2, \iota_1, \iota_2 \in \text{dom}(\mu) \setminus \text{IH}_\mu (\iota_1.f_1 \xrightarrow{\underline{\mathfrak{b}}} \iota_0 \text{ and } \iota_2.f_2 \xrightarrow{\underline{\mathfrak{b}}} \iota_0 \text{ implies } \iota_1 = \iota_2, f_1 = f_2 \text{ and } \iota_1 \notin \text{rog}_\mu(\iota_0))$
- all the internal aliasing of the balloon object itself are of a $\underline{\mathfrak{m}}$ mutable kind or supertype thereof: $\forall f, \iota_2 \in \text{rog}_\mu(\iota_0) (\iota_2.f \xrightarrow{_} \iota_0 \text{ implies } \iota_2.f \xrightarrow{\underline{\mathfrak{m}}, \underline{\mathfrak{e}}, \underline{\mathfrak{r}}, \underline{\mathfrak{e}}^*} \iota_0)$
- finally, every balloon has a tree shape: that is, all the direct subballoons are disjoint $\forall \iota_1, \iota_2 \in \text{RB}_\mu(\iota_0)$ such that $\iota_1 \neq \iota_2, \text{rog-}I_\mu(\iota_1) \cap \text{rog-}I_\mu(\iota_2) = \emptyset$ holds.

Where $\text{RB}_\mu(\iota_0) = \{\iota \in \text{dom}(\mu) \mid \iota_0 \xrightarrow{*\underline{\mathfrak{m}}} \iota, \iota_1 \xrightarrow{\underline{\mathfrak{b}}} \iota\}$

and BC_μ^Σ , balloon candidates are objects identifiers for which a balloon reference exists in the memory or on the stack:

$\text{BC}_\mu^\Sigma = \{\iota_0 \mid \iota_1 \in \text{dom}(\mu) \setminus \text{IH}_\mu, \iota_1 \xrightarrow{\underline{\mathfrak{b}}} \iota_0\} \cup \{\iota \mid \Sigma(\iota) = \underline{\mathfrak{b}}_ \}$

Finally, in (MEM-T) the last two side conditions check that all the reachable balloons in the stack are disjoint and that all the balloon data stored in the shared heap belongs to a single balloon.

Where $\iota \in \text{RB}_\mu(\Sigma)$ iff $\Sigma(\iota) = \underline{\mathfrak{b}}_ \text{ or } \iota \in \text{RB}_\mu(\iota_0)$ with $\Sigma(\iota_0) = \underline{\mathfrak{m}}_ ,$ and with the following definition of shared heap:

$\text{SH}_\mu^\Sigma = \{\iota \in \text{dom}(\mu) \mid \Sigma(\iota_0) = \underline{\mathfrak{m}}_ , \iota_0 \xrightarrow{*\underline{\mathfrak{m}}} \iota\}$.

Finally, rule (MEM-MEM-OK) verifies that the invariants are preserved during reduction: that is, a memory μ_1 can be the result of an operation over μ_0 only if the two memories are well formed in their respective stacks, and:

- all the immutable objects are preserved in the same state (side conditions 2 and 3),
- if two balloons are disjoint in the common part of the stack in μ_0 then the first balloon in μ_1 is disjoint with the second balloon, either if considered in μ_0 or μ_1 .

Theorem 1 (soundness).

If p is well typed, $\emptyset \vdash e : T$ and $\emptyset \mid e \xrightarrow{*} \mu_0 \mid e_0$ then either

- $e_0 = \iota, \iota : T_0 \vdash \mu_0$ and $T_0 \leq T$, or
- $\mu_0 \mid e_0 \rightarrow \mu_1 \mid e_1, \Sigma_0 \rightsquigarrow \Sigma_1 \vdash \mu_0 \rightsquigarrow \mu_1, \emptyset; \Sigma_0 \vdash e_0 : T_0 \leq T$ and $\emptyset; \Sigma_1 \vdash e_1 : T_1 \leq T_0$

As usual $_ \xrightarrow{*} _$ is the reflexive and transitive closure over $_ \rightarrow _$. It is possible to divide the proof into progress and subject reduction:

Theorem 2 (progress).

If p is well typed, $\Sigma \vdash \mu$ and $\emptyset \vdash \mu \mid e : T$ then either $e = \iota$ or $\mu \mid e \rightarrow _ \mid _$.

Proof. Since reduction ignores type modifiers the proofs is equivalent to that of FJ.

Theorem 3 (subject reduction).

If p is well typed, $\emptyset; \Sigma_0 \vdash e_0 : T_0, \Sigma_0 \vdash \mu_0$ and $\mu_0 \mid e_0 \rightarrow \mu_1 \mid e_1$ then $\Sigma_0 \rightsquigarrow \Sigma_1 \vdash \mu_0 \rightsquigarrow \mu_1$ and $\emptyset; \Sigma_1 \vdash e_1 : T_1 \leq T_0$

Proof. The proof is in the companion technical report [16].

e	$::= \dots \mid \{\bar{J}\} \text{ in } e$	expressions (also fork-join)
J	$::= T x = e; [\overline{dep}]$	Job
dep	$::= x \text{ if } L$	dependencies
L	$::= \text{true} \mid \text{false}$	logic constants
	$\mid L_1 \vee L_2 \mid \pi_1 \equiv \pi_2$	logic operations
π	$::= x \mid \iota$	
\mathcal{E}	$::= \square \mid \mathcal{E}.m(\bar{e}) \mid \iota.m(\bar{\iota}, \mathcal{E}, \bar{e})$	
	$\mid \text{new } C(\bar{\iota}, \mathcal{E}, \bar{e}) \mid T x = \mathcal{E}; e$	
	$\mid \mathcal{E}.f \mid \mathcal{E}.f = e \mid \iota.f = \mathcal{E}$	
	$\mid \{\bar{J}_1 T x = \mathcal{E}; [] \bar{J}_2\} \text{ in } e$	context




Figure 9. Adding the Fork Join Expression

5. Adding Parallelisation

The language BI-JAVA as defined before offers no support for parallelisation. Formally, the reduction arrow is deterministic, and thus is trivially confluent.³

We now provide a way to automatically add parallelisation as optimisation. Informally, it means that no observable behaviour has to change, to use the terminology of [3] we want to obtain a parallel performance model without altering the well known sequential semantic model.

We extend the language with fork-join expressions. Such expressions are added during compilation, replacing sequences of normal local variable declaration expressions, and thus are completely transparent from the point of view of the programmer.

Formally, we extend the language with fork-join expressions, allowing non deterministic reduction. However a well typed program have a confluent reduction. Then, we show a program transformation preserving the semantic, but replacing local variable declaration expressions with the new fork-join expressions.

In Figure 9 we show syntax for fork-join and logical expressions, context, reduction and typing rules for fork-join expressions.

5.1 New Syntax

Fork join expressions are composed by a sequence of variable initialisations and a conclusive expression. Note that the order of variable initialisations is relevant since it induces the (perceived) order of execution for the sub-expressions. Every variable initialisation is annotated with a logic expression encoding the dependency of such expression. Such logic expressions are composed by the constants **true** and **false**, the \vee (logical *or*) operation and the identity check over variables, or object identifiers. Any logic expression is related

³Only object identifiers are nondeterministically chosen.

to a specific variable, declared *inside* the same fork-join block, the syntax $x \text{ if } L$ means that a dependency exists with the expression that initialize x if L holds, that is, is equivalent to **true**. For example,

```
{ int x=1+2; []
  int y=1+2; []
  int z=x*y; [x if true, y if true]
} in z+1;
```

Means that x and y depends from nothing and can be computed in parallel, while before computing z we have *surely* to wait for the computation of x and y , since z depend from them.

Assuming a class C with methods $\bar{m} C.m(), \bar{m} C.k(), \bar{m} C.h()$ and **int** $h(\bar{m} C)\bar{m}$, we can show a more complex example:

```
{ m C x=l1.m(); []
  m C y=l2.k(); [x if l1 ≡ l2]
} in x.h(y);
```

The expression associated with x can be computed, while in order to compute the one associated with y we have to check if $l_1 \equiv l_2$.

- if the two object identifiers are different, we can start executing that expression in parallel,
- otherwise, we have to wait for the completion of the x expression.

5.2 New Reduction Rules

Any expression with an empty dependency declaration can be reduced, see last case of the context declaration.

Rule (OUT-E) is straightforward, while rule (COMMIT) replace occurrences of a local variable with its computed value. Dependency relation that are now equivalent to **false** are removed. Formally:

- substitution $J[x = \iota]$:

$$(T x = e; [\overline{dep}])[x = \iota] = T x = e[x = \iota]; [\overline{dep}][x = \iota]$$
- substitution $[\overline{dep}][x = \iota]$:

$$[] [x = \iota] = []$$

with $[\overline{dep}][x = \iota] = [\overline{dep}']$ and $x \neq x'$

$$[\overline{dep}, x \text{ if } L][x = \iota] = [\overline{dep}']$$

$$[\overline{dep}, x' \text{ if } L][x = \iota] = [\overline{dep}'] \text{ if } L[x = \iota] \text{ equivalent to false}$$

$$[\overline{dep}, x' \text{ if } L][x = \iota] = [\overline{dep}', x' \text{ if } L[x = \iota]] \text{ otherwise}$$
- Substitution over logic expressions $L[x = \iota]$ is straightforward: distributed on \vee operation, identity over constants and conventionally defined over variables.

In future work, it is clearly possible to perform many kinds of optimisations over dependency and logic expressions, either during the generation at compile time or during variable substitution.

We show now a possible reduction sequence for the following example; we assume the same class C as before, we omit the memory and show only some relevant steps:

```
1 { C a=new C().m(); [] C b=l1.m(); []
  C c=a.k(); [a if true, b if l1 ≡ a] } in b.h(c);
```

expressions associated with a and b can be computed in parallel,

```
2 { C a=l2; [] C b=e1; []
  C c=a.k(); [a if true, b if l1 ≡ a] } in b.h(c);
```

here we suppose the computation over $l_1.m()$ to be quite long, so in the shown point a is already associated with the value ι_2 , while b is of the form e_1 .

If $\iota_2 \neq \iota_1$ then the reduction will proceed in the following way:

```
3 { C b=e2; [] C c=l2.k(); [] } in b.h(c);
```

The execution of c can start now, while b is transparently continuing its execution, and it is now in e_2 .

4 { C b=e₃; [] C c=e; [] } in b.h(c);

In this stage both b and c are associated with long standing, autonomous, parallel computations.

5 { C b=l₃; [] C c=l₄; [] } in b.h(c);

Now they have both concluded the associated computation, and the fork-join can be reduced to $\iota_4.h(\iota_3)$.

5.3 New Typing Rule

The flow of parallel execution is regulated by the logic expressions inside the fork-join expressions, and the typing rule (FORK-JOIN-T) constraint the shape of such logic expression. Rule (FORK-JOIN-T) is quite complicated, since it performs different tasks:

- introduces local variables $x_1:T'_1, \dots, x_n:T'_n$ taking care of fresh conversions;
- find $\Gamma_0 \dots \Gamma_n$ using the non deterministic splitting relation;
- compute $\Gamma'_0 \dots \Gamma'_n$ restricting the domain of $\Gamma_0 \dots \Gamma_n$ on the free variable of the corresponding expression. For $i \in 1..n$ we assume expression e_i to not have free variables $x_1 \dots x_i$; finally, dependency can be computed. Since dependencies $\overline{dep}_1 \dots \overline{dep}_n$ are completely calculated inside this rule, they can be automatically inferred during our program transformation. Formally:

x_i if **true** \in compute-dep($x_1 \dots x_n, \Gamma_1 \dots \Gamma_n, \Gamma_0$)
 iff $x_i \in \text{dom}(\Gamma_0)$ and $i \in 1..n$
 x_i if $\text{dep}(x:T, x_0:T_0) \in$ compute-dep($x_1 \dots x_n, \Gamma_1 \dots \Gamma_n, \Gamma_0$)
 iff $\Gamma_i(x) = T, \Gamma_0(x_0) = T_0, i \in 1..n$
 and $\text{dep}(x:T, x_0:T_0)$ not equivalent to **false**.

The definition of $\text{dep}(x_1:T_1, x_2:T_2)$ is the subject of the next section.

5.4 Primitive Dependency $\text{dep}(\pi^r_1, \pi^r_2)$

For convenience we will use meta variables of the form π^r to indicate pairs $\pi:T$.

Defining $\text{dep}(\pi^r_1, \pi^r_2)$ is not straightforward, we have many possible design choices. A clearly correct solution is to define it as always **true**. This makes the assumptions of Theorem 4 and Theorem 5 trivially verified, since the program have to be executed in sequential fashion anyway. Of course our expressive type system allows us to do much more than that:

First of all, we know that if one of the two π^r is **i**mmutable or **f**resh, then there is no dependency at all:

$\text{dep}(\pi^r_1, \pi^r_2) = \text{false}$
 if π^r_1 or π^r_2 have **f**resh or **i**mmutable type

If both the π^r are **r**eadonly (or **e**xternal **r**eadonly), than no dependency is caused by this specific interaction.

$\text{dep}(\pi^r_1, \pi^r_2) = \text{false}$
 if π^r_1 and π^r_2 have **r**eadonly or **e**xternal **r**eadonly type

Moreover, if the two π^r both have a **b**alloon type, we can simply check the pointer equivalence:

$\text{dep}(\pi^r_1, \pi^r_2) = (\pi^r_1 \equiv \pi^r_2)$
 if π^r_1 and π^r_2 have **b**alloon type

where two equal variables will always refer to the same object:

$\pi:\mathcal{M}_1 C_1 \equiv \pi:\mathcal{M}_2 C_2 = \text{true}$

and two variables with incompatible types will never be equal:

$\pi_1:T_1 \equiv \pi_2:T_2 = \text{false}$ if $T_1 \not\leq T_2$ and $T_2 \not\leq T_1$

otherwise we generate the dynamic test over pointer equivalence:

$\pi_1:\mathcal{M}_1 C_1 \equiv \pi_2:\mathcal{M}_2 C_2 = (\pi_1 \equiv \pi_2)$

For space reasons we stop here and declare that in all the other cases $\text{dep}(\pi^r_1, \pi^r_2)$ is **true**.

However it is possible to extend the $\text{dep}(\pi^r_1, \pi^r_2)$ function in a variety of ways, greatly increasing the set of cases where an

efficient test can allow safe parallel execution. One extension can be found in the Technical Report [16].

5.5 Parallelisation Guarantee

The main result of our work is that in any BI-JAVA program, any sequence of variable declaration of the form

$$T_1 x_1 = e_1; \dots T_n x_n = e_n; e_0$$

can be optimized in the form

$$\{ T_1 x_1 = e_1; [\overline{dep}_1] \dots T_n x_n = e_n; [\overline{dep}_n] \} \text{ in } e_0$$

(where $\overline{dep}_1 \dots \overline{dep}_n$ are inferred by (FORK-JOIN-T)) and the semantic is not influenced in any way.

In the following theorems we express the preservation of the semantic in term of equivalence of the result modulo alpha conversions over object identifiers (symbol \simeq).

Theorem 4 (Confluence).

A well typed program have a confluent reduction. Formally:

$\forall e$ if $\emptyset \vdash e : T, \emptyset \mid e \xrightarrow{*} \mu_0 \mid \iota_0$ and $\emptyset \mid e \xrightarrow{*} \mu_1 \mid \iota_1$
 we have that $\mu_0 \mid \iota_0 \simeq \mu_1 \mid \iota_1$.

Theorem 5 (Safe Program Transformation).

For all well typed programs p_0 containing a method declaration md_0 containing a sequence of local variable declaration $T_1 x_1 = e_1; \dots T_n x_n = e_n$, for any well typed program p_1 identical to p_0 but where such sequence of local variable declaration is replaced with a corresponding fork-join operation and for any expression e well typed under p_0 ,

if such expression reduces to a pair $\mu_0 \mid \iota_0$ under p_0 ,

then such expression reduces to a pair $\mu_1 \mid \iota_1$ equivalent to $\mu_0 \mid \iota_0$ modulo object identifiers alpha renaming.

Formally:

$\forall p_0 = \text{class } C \text{ implements } \overline{C} \{ \overline{fd} \overline{md} \overline{md}_0 \} \overline{cd} \overline{id}$, where
 $md_0 = T m (T_1 x_1 \dots T_n x_n) M$

$\mathcal{E}[T_1 x_1 = e_1; \dots T_n x_n = e_n; e_0]$ and p_0 well typed,

$\forall p_1 = \text{class } C \text{ implements } \overline{C} \{ \overline{fd} \overline{md} \overline{md}_1 \} \overline{cd} \overline{id}$, where
 $md_1 = T m (T_1 x_1 \dots T_n x_n) M$

$\mathcal{E}[\{ T_1 x_1 = e_1; [\overline{dep}_1] \dots T_n x_n = e_n; [\overline{dep}_n] \} \text{ in } e_0]$

and p_1 well typed,

$\forall e$ where $\emptyset \vdash e : T$ under p_0 and $\emptyset \mid e \xrightarrow{*} \mu_0 \mid \iota_0$ under p_0

we have that $\emptyset \mid e \xrightarrow{*} \mu_1 \mid \iota_1$ under p_1 and $\mu_0 \mid \iota_0 \simeq \mu_1 \mid \iota_1$.

Thanks to these theorems, it is possible to start from a program with no fork-join expressions (and thus no parallelism) and to introduce automatic parallelisation in place of sequence of local variable declarations. Of course converting all the local variable declarations in fork-join expression could not be an optimal solution, investigating the optimal set of points to apply our transformation is future work. We are currently working on finalising the proof of these theorems and more discussion can be found in the companion technical report [16].

5.6 Application Example

There exists many well known operations transforming a mesh in some way, for example to provide a simplification. Thanks to our system, the following code can be transparently parallelised:

```
class Simplify{
  ...//some fields
  int simplify(Mesh m){
    ...//returns number of removed triangles
  }
}
class Workbench{ ...//as before, plus
  int simplifyMeshes(int mIndex1, int mIndex2){
    Balloon ml=this.meshes.get(mIndex1);
```

```

Balloon m2=this.meshes.get(mIndex2);
int x1=new Simpl().simplify(m1);
int x2=new Simpl().simplify(m2);
return x1+x2; } }

```

m_1 and m_2 would be (rapidly) computed sequentially, then x_1 will start its execution. At the same time the run time check $m_1 \equiv m_2$ will be performed. If $m_1 \neq m_2$ x_2 will instantaneously start its execution, otherwise if $m_1 \equiv m_2$ x_2 will start its execution only after x_1 is completed. This can happen if, for example, the user provides the same value for $mIndex_1$ and $mIndex_2$.

Note how the programmer could naturally write this code without even ever wondering about the parallelisation concept itself. The code will be well typed thanks to a sequence of promotions of the expression `new Simpl().simplify(m1): new Simpl()` is promoted to `fresh` using (PERMANENT1-TP); m_1 is promoted to `mutable` using (BALLOON-L-TP) thanks to (SUB-EXPR-T) where `new Simpl()` is extracted as a variable with type $\bar{m}_{\langle BLTP \rangle}$.

6. Related Work

The most closely related work to our approach is Deterministic Parallel Java (DPJ) [2, 3]. While our motivation is present in other recent works [15] we are among the few people who chose to build on top a *full encapsulation* mechanism such as balloons as opposed to a less restrictive and more flexible ones such as *ownership types* [13].

DPJ [2, 3] requires a programmer to think about the parallelisation explicitly and to manually insert the parallelisation constructs at the appropriate points in the program. DPJ uses an extended type checker with regions and effects and requires explicit manipulation of memory regions.

BI-JAVA takes the approach whereupon a straightforward reference comparison kind of check can be inserted automatically at the right points inside the program allowing safe parallel execution of expressions without the need for speculative parallelism (such as transactional memory etc). These safety checks capture the cases when speculative parallelism will succeed and thus avoid any need for potential roll backs. In other words, the cost of rolling back is transferred to the preventative checks that BI-JAVA inserts on behalf of the programmer. Future work might include an empirical study comparing the cost of these two approaches.

Finally, recent work by Naden et al. [12] extends Plaid with a similar set of access permissions to ours so that they can support borrowing of unique, shared, and immutable objects. One of their motivations is an ability to detect noninterference of concurrency, which we believe BI-JAVA achieves.

7. Conclusion

In this paper we have demonstrated how balloons and immutable objects can be utilised to guarantee when expressions can be executed in parallel. Our system employs a combination of a static and efficient dynamic check (e.g. a simple pointer equality). This is a step towards removing the burden of “thinking in parallel” from day to day programmers who want to take advantage of modern multicore architectures.

In the future we plan to extend our formalism and proofs to support a richer subset of a Java-like language. In particular, our companion technical report [16] contains a section that extends the $\text{dep}(\pi^r_1, \pi^r_2)$ which we would like to investigate further. We also plan to complete a prototype implementation of our approach and perform a user study evaluating the ease of use of our approach in day to day programming.

References

- [1] P. S. Almeida. Balloon types: Controlling sharing of state in data types. In *ECOOP*. Springer-Verlag, June 1997.
- [2] J. Bloch. *Effective Java: Programming Language Guide*. Addison-Wesley, 2001.
- [3] R. Bocchino. *An Effect System and Language for Deterministic-by-Default Parallel Programming*. PhD thesis, University of Illinois at Urbana-Champaign, 2010.
- [4] R. L. Bocchino, Jr., V. S. Adve, D. Dig, S. V. Adve, S. Heumann, R. Komuravelli, J. Overbey, P. Simmons, H. Sung, and M. Vakilian. A type and effect system for deterministic parallel java. In *OOPSLA*, pages 97–116. ACM Press, 2009.
- [5] J. Boyland. Why we should not add readonly to Java (yet). In *FTJJP*, Glasgow, Scotland, July 2005.
- [6] D. Clarke and T. Wrigstad. External Uniqueness is Unique Enough. In *ECOOP*, volume 2473 of *LNCS*, pages 176–200, Darmstadt, Germany, July 2003. Springer-Verlag.
- [7] D. Clarke, J. Noble, and T. Wrigstad, editors. *Aliasing in Object-oriented Programming*, volume In Print of *Lecture Notes in Computer Science*. Springer, 2012.
- [8] A. Corradi, M. Servetto, and E. Zucca. DeepFJig - Modular composition of nested classes. In C. Wimmer and C. W. Probst, editors, *PPPJ’11 - Principles and Practice of Programming in Java*, ACM International Proceedings Series, pages 101–110. ACM Press, 2011.
- [9] J. Hogg, D. Lea, A. Wills, D. de Champeaux, and R. Holt. The Geneva convention of the treatment of object aliasing. *OOPS Messenger*, 3(2): 11–16, April 1992.
- [10] A. Igarashi, B. C. Pierce, and P. Wadler. Featherweight Java: a minimal core calculus for Java and GJ. *TOPLAS*, 23(3):396–450, May 2001.
- [11] G. Lagorio, M. Servetto, and E. Zucca. Featherweight Jigsaw - a minimal core calculus for modular composition of classes. In *ECOOP’09 - Object-Oriented Programming*, volume 5653. Springer, 2009.
- [12] P. Müller and A. Poetzsh-Heffter. *Programming Languages and Fundamentals of Programming*. Technical report, Fernuniversität Hagen, 2001. Poetzsh-Heffter, A. and Meyer, J. (editors).
- [13] K. Naden, R. Bocchino, J. Aldrich, and K. Bierhoff. A type system for borrowing permissions. In *POPL*, pages 557–570. ACM Press, 2012.
- [14] J. Noble, J. Vitek, and J. Potter. Flexible Alias Protection. In *ECOOP*, volume 1445 of *LNCS*, pages 158–185. Springer-Verlag, July 1998.
- [15] A. Potanin, J. Östlund, Y. Zibin, and M. D. Ernst. Immutability. In Clarke et al. [6].
- [16] C. Reichenbach, Y. Smaragdakis, and N. Immerman. PQL: A purely-declarative Java extension for parallel programming. In *ECOOP*, volume 7313, pages 53–78. Springer-Verlag, 2012.
- [17] M. Servetto and A. Potanin. Balloon immutable java. Technical Report 12-18, ECS, VUW, 2012. <http://ecs.victoria.ac.nz/Main/TechnicalReportSeries>.
- [18] D. Walker. *Advanced Topics in Types and Programming Languages (Edited by B. Pierce)*, chapter Substructural Type Systems. MIT Press, 2005.
- [19] Y. Zibin, A. Potanin, S. Artzi, A. Kiezun, and M. D. Ernst. Object and reference immutability using Java generics. In *Foundations of Software Engineering*, 2007.
- [20] Y. Zibin, A. Potanin, P. Li, M. Ali, and M. D. Ernst. Ownership and immutability in generic java. In *OOPSLA*, pages 598–617. ACM Press, 2010.

A. Subtyping Relation

In our setting subtype $T_1 \leq T_2$ is defined as follows:

- $T \leq T$;
- $\mathcal{M} C_1 \leq \mathcal{M} C_2$ if $p(C_1) = \text{class } C_1 \text{ implements } C, _ \{ _ \}$ and $\mathcal{M} C \leq \mathcal{M} C_2$;

- $\mathcal{M} C_1 \leq \mathcal{M} C_2$ if $p(C_1) = \text{interface } C_1 \text{ extends } C, _ \{ _ \}$ and $\mathcal{M} C \leq \mathcal{M} C_2$;
- $\mathcal{M} C \leq \mathbf{e} C$ iff $\mathcal{M} = \mathbf{b}$ or $\mathcal{M} = \mathbf{m}$;
- $\mathcal{M} C \leq \mathbf{r} C$ iff $\mathcal{M} = \mathbf{i}$ or $\mathcal{M} = \mathbf{m}$;
- $\mathcal{M} C \leq \mathbf{er} C$ iff $\mathcal{M} = \mathbf{e}$ or $\mathcal{M} = \mathbf{r}$;

B. Extending $\text{dep}(\pi^{r_1}, \pi^{r_2})$

We show one possible extension allowing to safely parallelize also some mutable objects.

We extend the syntax to use π to indicate a variable or an object identifier possibly followed by a set of field accesses.

$\pi ::= x.\bar{f} \mid \iota.\bar{f}$ extended comparison

It often happens to use objects as records containing some highly related data, and to use such object as parameter for complex computation. In those cases the object graph have a very simple structure, so simple that the shape of their object graph is statically predictable up to the point where immutable or balloon fields are reached. We can inductively define a *predictable* class as a class where all the fields are of immutable, balloon or predictable type. Note that such definition rules out, for example, a linked list where the next field is mutable. We call mutable-predictable types composed by a mutable modifier and the name of a predictable class.

We can define the following two function over mutable predictable variables:

Function $\text{m-pred}(\pi:\mathbf{m} C)$ denoting all the paths that brings to another mutable entity inside the reachable object graph of π .

$\pi:\mathbf{m} C \in \text{m-pred}(\pi:\mathbf{m} C)$
 $\pi.\bar{f}:\mathbf{m} C_2 \in \text{m-pred}(\pi:\mathbf{m} C_0)$
 iff $\pi.\bar{f}:\mathbf{m} C_1 \in \text{m-pred}(\pi:\mathbf{m} C_0)$
 and $p(C_1) = \text{class } C_1 \text{ implements } _ \{ _ \mathbf{m} C_2 f; _ \}$

Function $\text{b-pred}(\pi:\mathbf{m} C)$ denoting all the path that brings to a balloon entity inside the reachable object graph of π .

$\pi.\bar{f}:\mathbf{m} C_2 \in \text{b-pred}(\pi:\mathbf{m} C_0)$
 iff $\pi.\bar{f}:\mathbf{m} C_1 \in \text{m-pred}(\pi:\mathbf{m} C_0)$
 and $p(C_1) = \text{class } C_1 \text{ implements } _ \{ _ \mathbf{b} C_2 f; _ \}$

Path to immutable references are not relevant, and by our definition of predictable class other kinds of objects are simply not reachable.

Now is possible to define when $\text{dep}(\pi^{r_1}, \pi^{r_2})$ is true over variables with mutable-predictable type.

$\text{dep}(\pi^{r_1}, \pi^{r_2}) = (\pi^{r_1} \in \text{b-pred}(\pi^{r_2}))$
 if π^{r_1} have balloon type and π^{r_2} have mutable-predictable type
 $\text{dep}(\pi^{r_1}, \pi^{r_2}) =$
 $\text{m-pred}(\pi^{r_1}) \cap \text{m-pred}(\pi^{r_2})$
 if π^{r_1} and π^{r_2} have mutable-predictable type

where⁴

non-disjointness: $\pi^{r_1} \dots \pi^{r_n} \cap \pi^r = \pi^{r_1} \in \pi^r \vee \dots \vee \pi^{r_n} \in \pi^r$
 inclusion: $\pi^r \in \pi^{r_1} \dots \pi^{r_n} = \pi^r \equiv \pi^{r_1} \vee \dots \vee \pi^r \equiv \pi^{r_n}$.

We now assume the dependency to hold for all the other cases:

$\text{dep}(\pi^{r_1}, \pi^{r_2}) = \text{true}$ otherwise

However, there are clearly many other cases in witch is possible, and convenient, to study a precise logic expression allowing safe parallelization. For example it would be possible to improve our study over external variables. This would require to enrich our logical expression with dynamic tests over the actually type of a variable.

⁴ No need to check $(\text{b-pred}(\pi^{r_1}) \cap \text{b-pred}(\pi^{r_2}))$ in the second case.

C. Discussion

C.1 Using Balloon invariant for encapsulation

In many cases an instance has full control of its balloon fields and their reachable object graph⁵. For example

```
class Foo{
  Balloon Bar bar;
  /*other fields*/
  void meth(/*some parameters*/ m{...} )
```

Independently of the other fields, if the parameters do not contains balloon, external or external readonly elements the method execution will have full control of the balloon field.

For example a pattern like the following will be safe in any arbitrary complex (and potentially circular) object graph:

```
class Foo{
  private Balloon Bar bar; //no public getter method
  private boolean opInProgress=false; //no public setter
  void meth(Mutable Boh pl)Mutable{
    if (this.opInProgress)
      throw new Error("unexpected indirect recursion");
    try { this.opInProgress=true; //this.bar only set in Foo
          //this.bar only accessed in Foo
          doStuff(this.bar, pl); }
    finally { opInProgress=false; }
    ... }
}
```

We use variable opInProgress to prevent undesired indirect multiple accesses to the bar data.⁶ If the method have balloon parameters then a simple dynamic check (straightforward identity test $x!=y$) can ensure full control. However if external or external readonly parameter are present, and they are of a type that can potentially reach a Bar object in the reachable object graph, ensuring full control is not trivial. If an object points to a complex object graph that have non trivial consistency relations, and modification over this object graph requires to temporary break such relations, the pattern above allows is a way to encode such modification operation with the guarantee that no other client will be able to see such uncoherent state.

C.2 Memory management optimisations

immutable objectss can have an ad hoc memory management. Whenever a fresh object is converted into immutable, the hash code is computed and the object is stored in a system Map. If an equivalent object is already present in the map, the new version can be instantaneously garbage collected. System generated hashcode and equals, supporting circularity, should be automatically introduced. Moreover, keys of a HashSet/HashMap should be only Immutable, such maps will be faster and with a simpler semantic.

Moreover, from F_3 any expression that does not return a balloon, external or external readonly reference can be safely garbage collect all the balloon, external and external readonly created references at the end of its execution. Moreover, every l-closed expression that returns an immutable value (or a primitive type) can safely garbage collect all the (non immutable) created objects.

This opens the door for a simple language extension with a safe try-with-resource statement, when the resource object is guaranteed to be collected at the end of the statement.

⁵ Reachable objects are not taken into consideration by many approaches, so our guarantee is stronger, more useful and helps produce code that is maintainable w.r.t. the change of internal representation

⁶ Other approaches replace this pattern with temporary borrowing facilities or destructive reads (i.e. the field returns null when accessed). We believe this pattern provides better error control, even if more verbose.

D. Proofs

Proof of Theorem 3. By induction over the reduction rules. All cases but rule (CTX) does not require the use of the inductive iphotesis. Since we have a call by value semantic all cases of direct reduction work over very simple expressions where there is little space for the application of promotions. In the case of (V-DEC) we expect promotions over the whole expression $T x = \iota; e$ to be always applicable directly to the conclusive expression e .

$$r ::= \iota.m(\bar{\iota}) \mid \mathbf{new} C(\bar{\iota}) \iota_0.f \mid \iota_0.f = \iota_1 \mid T x = \iota; e \quad \text{redex}$$

Case (CTX)

Choosing

(C1) $\iota \in \text{dom}(\Sigma_0)$ iff ι is a subterm of r_0

(C2) $\text{dom}(\mu_0) = \text{rog}_{\mu_0, \mu}(\text{dom}(\Sigma_0))$

Given

(G1) $\mu_0, \mu \mid \mathcal{E}[r_0] \rightarrow \mu_1, \mu \mid \mathcal{E}[e_1]$

(G2) $\Sigma_0, \Sigma \vdash \mu_0, \mu$

(G3) $\emptyset; \Sigma_0, \Sigma \vdash \mathcal{E}[r_0] : T'_0$

Show

(R1) $\Sigma_0, \Sigma \rightsquigarrow \Sigma_1, \Sigma \vdash \mu_0, \mu \rightsquigarrow \mu_1, \mu$

(R2) $\emptyset; \Sigma_1, \Sigma \vdash \mathcal{E}[e_1] : T'_1 \leq T'_0$

we obtain:

(I1) $\mu_0 \mid r_0 \rightarrow \mu_1 \mid e_1$ by (G1),

(I2) $\Sigma_0^{\overline{TP}} \vdash \mu_0$ by (G2), (G3) and Lemma 7.i,

(I3) $\emptyset; \Sigma_0^{\overline{TP}} \vdash r_0 : T_0$ by (G2), (G3) and Lemma 7.ii,

(IH1) $\Sigma_0^{\overline{TP}} \rightsquigarrow \Sigma_1^{\overline{TP}} \vdash \mu_0 \rightsquigarrow \mu_1$ and

(IH2) $\emptyset; \Sigma_1^{\overline{TP}} \vdash e_1 : T_1 \leq T_0$ by (I1), (I2), (I3) and inductive iphotesis.

Now we can obtain our results: By (IH1), (G2) and Lemma 9.ii we obtain (R1) By (I3), (IH2), (G3), and Lemma 9.i we obtain (R2)

Lemma 6 (promotionPreservation).

$$\Sigma_0 \vdash \mu_0$$

implies

$$\Sigma_0^{\overline{TP}} \vdash \mu_0$$

Proof. By induction over sequence of promotions:

Case empty set of promotion

Trivial

Case Promotion sequence ends with a permanent type promotion

all the **mutable** or **readonly** $\iota:T$ in Σ_0 are **external** or **external readonly**, in $\Sigma_0^{\overline{TP}}$; respectively. This means that $\text{RB}_{\mu}(\Sigma_0^{\overline{TP}}) \subseteq \text{RB}_{\mu}(\Sigma_0)$. Thus rule (MEM-T) is still applicable, since all side condition stay the same, but last is checked on less cases.

Case Promotion sequence ends with a balloon local type promotion

it can be shown correct in two steps:

- All the **mutable** or **balloon** variable in the stack are seen as **external**.

At this point rule (MEM-T) is still applicable, since all side condition stay the same, but last is checked on zero cases.

- Then we promote a single **external** in the stack to **mutable**. This repopulate the set $\text{RB}_{\mu}(\Sigma_0^{\overline{TP}})$ shared heap with (potentially) all the element inside the balloon containing such promoted object identifier.

Rule (MEM-T) is still applicable: the last two side conditions of (MEM-T) still holds:

- for the last clause of $\text{valid-B}_{\mu}(\iota)$ we know that all the sub balloon are disjoint, and
- only elements inside a single balloon are added to the shared heap.

Case Promotion sequence ends with a temporary type promotion
As in the permanente type promotion case, we simply reduce the number of cases the last side condition of (MEM-T) is checked.

Lemma 7 (promotion application).

(G1) $\text{dom}(\mu_0) = \text{rog}_{\mu_0, \mu}(\text{dom}(\Sigma_0))$

(G2) $\Sigma_0, \Sigma \vdash \mu_0, \mu$

(G3) $\emptyset; \Sigma_0, \Sigma \vdash \mathcal{E}[r_0] : T'_0$

implies

i $\Sigma_0^{\overline{TP}} \vdash \mu_0$

ii $\emptyset; \Sigma_0^{\overline{TP}} \vdash r_0 : T_0$

Proof.

i By (G2) and the property of the choosed division between Σ_0, Σ and μ_0, μ we know that $\Sigma_0 \vdash \mu_0$ We conclude by Lemma 6.

ii We choose \overline{TP} to coincide with the set of promotion applied over subexpression r_0 . By rule (LOC-T) $\emptyset; \Sigma_0^{\overline{TP}} \vdash r_0 : T_0$ iff $\Gamma_0^{\overline{TP}} \vdash r_0[\iota_1 = x_1 \dots \iota_n = x_n] : T_0$ with $\text{dom}(\Sigma) = \iota_1 \dots \iota_n$, $\text{dom}(\Gamma) = x_1 \dots x_n$, $\Gamma_0(x_i) = \Sigma(\iota_i)$

Thus, for our chosen \overline{TP} we can conclude by straightforard application of the typing rules over $\Gamma_0, \Gamma \vdash \mathcal{E}[r_0][\iota_1 = x_1 \dots \iota_n = x_n] : T'_0$

Lemma 8 (substitution zero).

(G1) $\mu_0 \mid r_0 \rightarrow \mu_1 \mid e_1$

(G2) $\emptyset; \Sigma_0 \vdash r_0 : T_0$

(G3) $\emptyset; \Sigma_0, \Sigma \vdash \mathcal{E}[r_0] : T'_0$

(G4) $\emptyset; \Sigma_1 \vdash e_1 : T_1 \leq T_0$

(G5) $\Sigma_0 \rightsquigarrow \Sigma_1 \vdash \mu_0 \rightsquigarrow \mu_1$

(G6) $\Sigma_0, \Sigma \vdash \mu_0, \mu$

exists a Σ_1 such that

i $\emptyset; \Sigma_1, \Sigma \vdash \mathcal{E}[e_1] : T'_1 \leq T'_0$

ii $\Sigma_0, \Sigma \rightsquigarrow \Sigma_1, \Sigma \vdash \mu_0, \mu \rightsquigarrow \mu_1, \mu$

Proof. Induction over the typing rules for term (G3). For cases (F-ACCESS-T), (M-INVK-T) and (V-DEC-T) $\mu_0 = \mu_1$ and thus (ii) trivially holds. Cases for rules (F-ACCESS-T) and (M-INVK-T) are simple also with respect to (ii).

Case (F-UPDATE-T)

(i) holds by (F-UPDATE-T) second premise. The receiver is by construction inside the shared heap. in case the receiver is also inside a balloon, (ii) requires a case analysis over the possible kinds of the assigned value:

• **i** (ii) trivially holds if the assigned value is in the immutable heap.

• **r** (ii) holds, since readonly links does not contribute for the set of reachable ballons $\text{RB}_{\mu}(\Sigma)$.

\bar{m} (MEM-T) last side condition states that all the element in the shared heap belongs to at most one balloon. Since the receiver is also on the shared heap (ii) holds.

Case (F-UPDATE-TB)

(i) holds by (F-UPDATE-TB) second premise. (ii) trivially holds since the assigned value is inside the fresh heap in μ_0 .

Case (NEW-T)

The application of this rule ignect a new element on the stack Σ_1 , of kind \bar{m} utable. (i) holds by inductive hypotesis and sub-type transitivity. Since this new element is inside the shared heap but outside any baloon, also (ii) holds.

Case (PERMANENT1-TP)

In this case we know that r_0 can be typed as mutable in $\text{l-close-p}(\Gamma_0) \vdash r_0[l_1 = x_1 \dots l_n = x_n] : \bar{m} C_0$ By cases depend- ing on the kind of redex:

$\iota.m(\bar{l})$ (i) holds since also the expression e_1 , body of the method m can be typed with (PERMANENT1-TP) in this environment. (ii) holds since $\mu_0 = \mu_1$.

new $C(\bar{l})$ The created object clearlyly respect both the fresh invariants and the immutable invariants. So it can be consid- ered inside the fresh heap. Indeed all refeence in \bar{l} are either in the fresh heap or in the immutable heap.

$\iota_0.f$ The produced reference is inside the fresh heap, indeed ι_0 must be of a fresh kind $\bar{m}(\text{PTE})$ in Σ_0 , that is $\Sigma_0(\iota_0) = \bar{m}(\text{PTE})_-$.

In the detail, (i) holds since $\iota_{n+1}[l_1 = x_1 \dots l_{n+1} = x_{n+1}]$ can be typed with (VAR-T) as a fresh reference. (ii) holds since the newly created object, ι_{n+1} , can be considered inside the fresh heap.

$\iota_0.f = \iota_1$ clearlyly $\Sigma_0(\iota_1) = \bar{m}(\text{PTE})_-$, thus ι_1 is already inside the fresh heap in Σ_0 .

In the detail, (i) holds since $\iota_1[l_1 = x_1 \dots l_n = x_n]$ can be typed with (VAR-T) as a fresh reference. (ii) holds since (G5).

$T x = \iota; e$ (i) holds since also the expression e can be typed with (PERMANENT1-TP) in this environment. (ii) holds since $\mu_0 = \mu_1$.

Case (PERMANENT2-TP)

In this case we know that r_0 can be typed as readonly in $\text{l-close-p}(\Gamma_0) \vdash r_0[l_1 = x_1 \dots l_n = x_n] : \bar{r} C_0$ By cases depend- ing on the kind of redex, we omit details about (i) and (ii) since they are similar to the former point.

$\iota.m(\bar{l})$ also the expression e_1 , body of the method m can be typed with (PERMANENT1-TP) or (PERMANENT2-TP) in this environment.

new $C(\bar{l})$ The created object clearlyly respect both the fresh invariants and the immutable invariants. So it can be consid- ered inside the immutable heap.Indeed all refeence in \bar{l} are either in the fresh heap or in the immutable heap.

$\iota_0.f$ The produced reference is inside the immutable heap, indeed ι_0 must be of a fresh kind $\bar{m}(\text{PTE})$ or of an immutable kind in Σ_0 , that is $\Sigma_0(\iota_0) \in \{\bar{m}(\text{PTE})_-, \bar{i}_-\}$.

$\iota_0.f = \iota_1$ clearlyly $\Sigma_0(\iota_1) \in \{\bar{m}(\text{PTE})_-, \bar{i}_-\}$, thus ι_1 is already inside the fresh or immutable heap in Σ_0 .

$T x = \iota; e$ also the expression e can be typed with (PERMANENT1- TP) or (PERMANENT2-TP) in this environment.

Case (BALLOON-L-TP)

In this case we know that r_0 can be typed in $\text{l-close-b}(\Gamma), x_0 : \bar{m} C_0 \vdash r_0[l_1 = x_1 \dots l_n = x_n] : \mathcal{M} C$

and to proof (i) we have to show that

$\Gamma, x_0 : \mathcal{M}_0 C_0 \vdash e_1[l_1 = x_1 \dots l_n = x_n] : \text{l-close}(\mathcal{M}) C$

That is, an \bar{e} xternal reference to ι_0 is seen as \bar{m} utable. By cases depending on the kind of redex:

$\iota.m(\bar{l})$ also the expression e_1 , body of the method m can be typed with (BALLOON-L-TP) in this environment. (ii) holds since $\mu_0 = \mu_1$.

new $C(\bar{l})$ (i) holds since $\iota_{n+1}[l_1 = x_1 \dots l_{n+1} = x_{n+1}]$ can be typed with (VAR-T) as a mutable reference. (ii) holds since the created object mutable reachable object graph clearlyly refers only to elements inside the balloon that contains ι_0 . So it can be considered inside the shared heap.

$\iota_0.f$ The produced reference is also inside the reachable object graph of ι_0 .

$\iota_0.f = \iota_1$ Following the same process used for case (F- UPDATE-T) we can show that both ι_0 and ι_1 are in the shared heap, thus we are not violating the balloon invari- ants, thanks to (MEM-T) last side condition.

$T x = \iota; e$ also the expression e can be typed with (BALLOON- L-TP) in this environment.

Case (TEMPORARY-TP)

-

Case (SUB-EXP-T)

-

Lemma 9 (substitution).

(G1) $\mu_0 \mid r_0 \rightarrow \mu_1 \mid e_1$

(G2) $\emptyset; \Sigma_0^{\overline{TP}} \vdash r_0 : T_0$

(G3) $\emptyset; \Sigma_0, \Sigma \vdash \mathcal{E}[r_0] : T'_0$ applying \overline{TP} to r_0

(G4) $\emptyset; \Sigma_1^{\overline{TP}} \vdash e_1 : T_1 \leq T_0$

(G5) $\Sigma_0^{\overline{TP}} \rightsquigarrow \Sigma_1^{\overline{TP}} \vdash \mu_0 \rightsquigarrow \mu_1$

(G6) $\Sigma_0, \Sigma \vdash \mu_0, \mu$

exists a Σ_1 such that

i) $\emptyset; \Sigma_1, \Sigma \vdash \mathcal{E}[e_1] : T'_1 \leq T'_0$

ii) $\Sigma_0, \Sigma \rightsquigarrow \Sigma_1, \Sigma \vdash \mu_0, \mu \rightsquigarrow \mu_1, \mu$

Proof. By induction over sequence of promotions:

Case empty set of promotion

Proved in Lemma 8

Case Promotion sequence ends with a permanent type promotion

By cases over the different kinds of redex. In all cases but con- structor invocation is enough to consider $\Sigma_0 = \Sigma_1$

The constructor invocation is thus the more intresting case.

Is required to conside the newly created object as fresh. This is possible since we are inside a perment type promotion: indeed outside of the promotion all the parameter of the constructed object will be either \bar{i} mmutable or $\bar{m}(\text{PTE})$.

Case Promotion sequence ends with a balloon local type promotion

-

Case Promotion sequence ends with a temporary type promotion

-