A type checked prototype-based model with linearity

Andi Bejleri, Jonathan Aldrich, and Kevin Bierhoff

December 2004
CMU-ISRI-04-142

Institute for Software Research International
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

This technical report is an expanded version of a paper submitted for publication, and is also a preliminary draft of the first author's thesis.

This work was supported in part by the High Dependability Computing Program from NASA Ames cooperative agreement NCC-2-1298, NSF grant CCR-0204047, and the Army Research Office grant number DAAD19-02-1-0389 entitled "Perpetually Available and Secure Information Systems."

Keywords: change code dynamically, prototype-based languages, Self, static type checker, type safety

Abstract

Dynamic inheritance, originating in the SELF programming language, is the ability of an object to change the code that it inherits at run time. This ability is useful for modeling objects that behave in different ways at different points in the object's lifecycle. Unstructured dynamic inheritance, however, allows arbitrary changes to the interface of the object, and thus is incompatible with statically typechecked languages such as C++, C# and Java.

This paper provides a more structured facility for dynamic inheritance, where a type system tracks the changes in an object's interface that occur as the inheritance hierarchy is changed. We define a formal model of a language and type system with dynamic inheritance, and prove that the type system is sound in that it prevents run-time type errors. The type system tracks the linearity of objects and methods in order to ensure that objects whose interfaces change are not aliased.

```
|socket < −() | "a new empty object"
socket AddSlots: (|bind = (< code > ...
   socket AddSlots: (|port < −Nil|) "adding a new data slot"
   socket AddSlots: (|listen = (< code > ...
      socket AddSlots: (|accept = (< code > ...
         socket AddSlots: (|read = (< code >)|)
         socket AddSlots: (|write = (| : data...| < code >)|)
         socket AddSlots: (|close = (< code >)|)
)|))|))|)
```

Figure 1: TCP socket example illustrating the expressiveness of Self and the challenges the type system has to deal with

1. Introduction

Objects, by their nature, often have different behavior in different stages of their lifecycle. However, in classical object-oriented languages the type of an object and the messages it understands are fixed at compile time and cannot be changed at run time. SELF [15] is a prototype-based object-oriented language that allows programmers to dynamically change the inheritance hierarchy and the set of methods that each object understands. Thus SELF objects can have different behavior at different moments. This model is appealing for implementing a large variety of software systems. Several interesting properties of the language are:

- There are no classes in SELF. Instead, a *prototype* mechanism is used for object creation. A new object is created by *cloning* (copying) another object that serves as its prototype.
- There is no distinction between state and behavior. The methods and fields of an object are unified into slots. A slot with a function can be used to model a method and a slot with data can be used to model a field. Consequently, there is no distinction between accessing a field and sending a message. Every object is simply represented by a list of slots.
- There are one or more *delegation slots* that refer to objects from which behavior is inherited. This mechanism allows objects to share behavior or state. The object can change the objects to which it delegates at any point in program execution. Section 2 describes a number of situations where the expressiveness of dynamic inheritance is beneficial.
- SELF also supports adding methods dynamically by adding a new slot. A method body (or field) can be changed by changing the value of the slot.

Unfortunately, the additional expressiveness of SELF comes at a cost: A programmer might experience "message not understood" errors at runtime. Figure 1 demonstrates how a Berkeley TCP socket might be implemented in SELF. By adding e.g. the *port* field only after *bind* was called we can ensure that clients cannot read uninitialized values from the object. The same applies to calling methods in the wrong order: since the body of *bind* adds the *listen* method, clients cannot call *listen* until after *bind* has been called.

However, changing delegation and adding methods at run time to objects can be the source of bugs if they are not controlled. The SELF compiler has no way of statically detecting whether the *port* is defined on a socket object at a particular point in the program. Instead, an access to a non-existing slot will cause a "message not understood" error at run time. It is easier to identify the cause of this error than it would be if the method call succeeded but corrupted the socket's data structures (as might happen without the user of dynamic method addition), but nevertheless it would be nice to avoid both errors using static checking. As we will see later, aliasing in particular makes it very hard for the programmer to manually make sure that such errors cannot occur. As a result, the potential benefits of dynamic inheritance and method change at run time are underutilized.

The contribution of this paper is a type system that statically ensures that all accesses to object slots will succeed at runtime, even in the presence of dynamic inheritance and method changes. We formally define a new language, EGO, which is similar to SELF but restricts SELF's flexibility somewhat. In particular we control dynamic changes to aliased objects. We designed EGO in such a way that a static type checker can guarantee that a well typed program will lack of "message not understood" errors at run time. The type safety proof for EGO directly implies this property.

The type system of EGO blends the features of several previous type systems in order to achieve soundness. For each object it keeps track of all methods a client can invoke. The type system distinguishes between linear (non-aliased) and non-linear (aliased) objects [8]. It statically ensures that a linear variable or function is used exactly once, while allowing aliasing and multiple uses of non-linear objects and functions. To our knowledge, our system is the first to integrate first-class linear functions in an object-oriented environment.

The use of linearity in typing objects solves crucial aliasing and typing issues. Dynamic changes to the type of the object (e.g. by adding a method) are only permitted to linear objects. A new object has a linear type when it is created and the type system guarantees its linearity during the program unless the client explicitly makes it an aliased object (on which fewer changes are allowed).

Our system can be considered a foundation for research into more flexible typestate systems for objects [6,7]. As a foundational system, it may not be as succinct or easy to use as a source-level language, but instead is designed to further

understanding into the core mechanisms of typestate and to explore more flexible implementation strategies for typestate, such as dynamic changes to the methods and superclass of an object. Incorporating this additional flexibility into easy-to-use source-level languages is an interesting area of future work.

The remainder of this paper is organized as follows. Section 2 gives an intuitive presentation of EGO illustrated with a number of examples. Section 3 introduces the core language, its dynamic semantics, static semantics, and a brief presentation of the type safety proof. Section 4 summarizes related work, and the last section concludes.

2. Overview of EGO

This section gives an informal introduction to our language. After giving a brief intuition of its constructs, we show how to encode some common object-oriented programming idioms. We then discuss how EGO tackles the important problem of aliasing. That forms the basis for a detailed description of how methods are defined. Finally we demonstrate EGO's expressive power with a number of examples. Throughout the section we highlight the challenges static typechecking has to deal with.

2.1 Language Intuition

A program in EGO is an expression. An expression can be anything from a simple value to a complex object manipulation. Some kinds of expressions can contain other nested expressions. We use the notion of lambda abstractions to define a function and bind a variable in its body expression. Moreover, we introduce a number of primitives for object manipulation that are inspired by the work on SELF [15].

- *clone* allows duplicating an object.
- *delegate* sets the super field of an object, thus determining from whom the object inherits.
- *addMethod* is used for adding a method to an object (or changing the implementation of an existing method).
- *change_linearity* is a technical primitive used for dealing with aliasing, as we shall see later.
- Finally, in typical object-oriented style, $e.m$ invokes a method on an object.

The first four primitives yield the object created or manipulated to be used in the surrounding expression. The last one is used for method calls and thus yields the body of that method.

In the following sections we will develop a number of examples that show these primitives in action.

2.2 Elementary Programming Idioms

EGO is designed as a core language for expressing dynamic inheritance and method addition. We can define a number of derived forms for well-known and convenient idioms that will help us write more concise programs. That will also help us in presenting more advanced examples in the remainder of the section.

This section focuses on the notions of a *let* construct and instance fields for objects. We will also discuss how to create new objects and how to use them like traits in SELF (thus in a class-like manner).

The *let* construct is well-known from languages like ML [13]. It typically binds a variable name in a subsequent expression (e.g. $let\ x = 5\ in\ x{+}1$). We can simulate this behavior with a simple lambda expression as reflected in the following definition. It also allows us to define sequences of operations.

$$let\ x : \tau = e_1\ in\ e_2 \stackrel{\text{def}}{=} (\lambda x{:}\tau.e_2)\ e_1$$

$$e_1; e_2 \stackrel{\text{def}}{=} let\ _ = e_1\ in\ e_2$$

Instance fields as defined in object-oriented languages is the canonical way of holding values in an object over which that object abstracts. Methods are used to manipulate the fields of the object. Being a core language, EGO does not incorporate any form of information hiding (thus everything is public). It also does not support fields as a primitive. Instead we can encode fields as parameterless methods. Defining a field would look like the following.

$$e_1.f = v \stackrel{\text{def}}{=} e_1.addMethod(f, let\ x = e_2\ in\ \lambda self{:}\tau.x)$$

This will also work for reassigning a field value. In this case, *addMethod* will just redefine the method body. Note that e_1 has to be an object and we use a let binding to evaluate e_2 to a value before the method body is created. Access of a field then becomes invoking a parameterless method (with $e.f$, where e is an object and f the name of a field).

In fact we can use the above derived form to add or change an arbitrary method on an object: If f is itself a lambda expression then it simply defines a method body that relies on additional arguments. (We will discuss method definitions in detail below.)

How do we get an object in the first place? EGO is a prototype-based language that allows us to *clone* existing objects. A program can assume the variable *Object* bound to the first object in the system. Thus creating a new object, adding two methods, and invoking the first one can be realized as follows.

```
// trait for s
let b = change_linearity(clone(Object).addMethod(service, λx:Nat.x + 1)) in
// now define s itself
let s = change_linearity(clone(Object).delegate(b)) in
// and finally the clients
  let c1 = clone(Object).addMethod(r, s) in
  let c2 = clone(Object).addMethod(r, s) in
  ⋮
    // invalid: let _ = s.delegate(a) in
    c2.r.service(5)
```

Figure 2: A server object s referenced by multiple clients

$$clone(Object).addMethod(m_1, e_1).addMethod(m_2, e_2).m_1$$

Expressions for a method body have to evaluate to a lambda abstraction for *self*. When a method is executed, the receiver object will be applied to this outermost lambda. Objects are recursive types, thus methods can refer to their receiver and its (other) methods by accessing the bound variable *self*.

We often want to use an object in a class-like manner, meaning that the object contains instance methods to be used by other objects. Such an object is called a *trait* in the SELF literature [15]. We can use the *let* construct in combination with delegation to realize traits as shown below.

$$\begin{array}{l} let\ Trait = clone(Object).addMethod(succ, \lambda self{:}\tau.self.f + 1)\ in \\ \quad clone(Object).delegate(Trait).addMethod(f, \lambda self{:}\tau.5).succ \end{array}$$

The result of this expression would be 6. Obviously, an arbitrary number of objects can be defined that inherit their behavior from the *Trait* object above by delegation and define their own f field. Another option is to simply clone the trait object, which would result in simply duplicating all of the methods of *Trait* rather than sharing them through delegation. We will present an example of this more *prototype-oriented* approach in a later section.

2.3 Dealing with Aliasing

So far we have ignored a major complication of our system: Aliasing. An aliased object is (possibly) referred to by multiple names (references) in a program as opposed to *linear* objects that have only one name. Aliased objects are also called "non-linear", and linear ones are sometimes called "non-aliased".

In an object-oriented setting, aliasing is almost inevitable because of the state held in instance fields. A very common notion is that a server object s is used by multiple clients c_i that all hold a reference to s in their fields $c_i.r$. s is then heavily aliased as in the following definition.

```
// trait for s
let b = clone(Object).addMethod(service, λx:Nat.x + 1) in
// now define s itself
let s = clone(Object).delegate(b) in
// and finally the clients
  let c1 = clone(Object).addMethod(r, λself:τc..s) in
  let c2 = clone(Object).addMethod(r, λself:τc..s) in
    ...
```

If we now change the configuration of s e.g. by changing its delegate from b to a with $s.delegate(a)$, obviously all clients are affected. In particular, it is hard to tell whether s will still work the way its current clients expect it to.

For this reason, we forbid a change of delegation for aliased objects as well as adding or changing methods for such objects if it changes the method's signature. We allow methods to be modified for aliased objects as long as the new method has the same signature as the old method. This allows us to model field updates, for example.

Moreover, we forbid delegation to a linear object (because that would be just like a second explicit reference to that object). Instead, we introduce the *change_linearity* primitive mentioned earlier to explicitly convert a linear into an aliased object that can then be a delegee. Note that there is no way of turning an aliased object back into a linear one. Figure 2 correctly encodes the situation described above in EGO.

Intuitively, these restrictions have to do with the typing of objects. Changing a method signature or the delegation changes the type of the object. That means that the aliases to that expression somehow would have to invisibly change their types as well, which would be difficult or impossible for a static type system to track in the general case. Conversely, changing a linear object affects only the type of the expression at hand, which is what a static type checker tracks anyway.

```
let lin = clone (Object)
  let oclone (Object) .addMethod (l, ¡λself:¡τ.(self, lin))  in
  // lin is no longer available
    let (o2, lin2) = o.l in
      // instead we can now use lin2
      // o2 replaces o, but does not contain l any more
```

Figure 3: A linear method consuming a variable on the stack and its linear receiver

On the typing level we introduce a *linearity flag* for objects and lambdas that we write as "¡". *change_linearity* explicitly removes this flag for an object, thus making it possibly aliased. Bodies of linear lambda abstractions have access to the linear variables defined in the context of that abstraction (i.e. on the stack). The type system guarantees that such linear variables are used only once. (We say they are "consumed" on usage.) Figure 3 gives an example assuming pairs written as (x, y). Non-linear lambda-abstractions, on the other hand, can only access the non-linear variables in the context. Non-linear variables can be used multiple times.

We call all methods linear that have a linear lambda. Linear methods are consumed upon invocation, i.e. they are effectively removed from the receiver object. This guarantees the linearity of the context variables: If we could call the linear method l from figure 3 twice, then we would gain two aliases to *lin* "through the back door". As recursive calls to the same linear method would have the same harmful effect, we have to remove a linear method from its object *before* that method's body is evaluated. Thus the method l in figure 3 is not only no longer available after l was evaluated, but l cannot invoke itself on *self* again either.

We forbid cloning of objects with linear methods for the same reason: That would result into pairs of linear methods accessing the very same variable. However, the object can be linear (because it is completely duplicated), and the resulting clone is linear in any case. Thus all objects are linear in the beginning of their lifetime and can be converted into a non-linear object explicitly using *change_linearity* (but not back into a linear object).

An alternative to the solution of consuming linear methods upon invocation would be to consume the receiver as a whole. We consider this a bad choice: Only one method could be ever executed on a linear object.

Independently from the linearity of a lambda itself, its argument can be linear or non-linear. A linear lambda argument requires a linear object. The object applied to such a lambda is no longer available at the invocation site after that application (again, we say it is "consumed"). However, the lambda abstraction can return its argument to the caller as the method $o.l$ in figure 3 illustrates. o is no longer available after the last line, but it is passed back into $o2$.

2.4 Method Definition

In order to capture the dynamic manipulations of objects statically, EGO types objects with a recursive record type [1,9] containing an explicit list of all the methods the object defines together with a field for its delegate. A linear object containing an integer field as well as a linear method that takes an integer argument and yields an integer would be typed as follows. The object delegates to an empty object like *Object*.

$$t.¡ < field : t \to int, linMeth : t \multimap int \multimap int, super :<>>$$

We use $\multimap$ for typing linear lambda abstractions and $\to$ for non-linear ones. Every method body definition must be an explicit lambda abstraction for *self*, the receiver object. The type of *self* essentially lists *all* methods expected to be defined for the receiver – not just those needed in the body. Additional arguments can be captured with nested lambdas.

The requirement that *self* must be typed with a recursive record type is essentially not different from typing an object with a class name in e.g. Java: Since the methods in a Java class cannot be manipulated the class name can be used as a (shorter) synonym for a record type containing all methods defined for that class.

In fact, our system is much more flexible in that different methods can declare different receiver object types. The programmer can make explicit what he expects will change over the lifetime of the object. He can hereby enforce possible sequences of method invocations on the object, i.e. the object's *protocol*. The receiver object type for a method then reflects the *typestate* the object has to be in [4] for invocations of that method. Figure 4 gives an example of method definitions using typestate. Note that we give *typedefs* for several record types in the beginning to improve readability. They are not part of the core EGO language.

We illustrate the business logic of a Web-based phonebook. Such applications are characterized by *two-phased actions*: First, the user indicates the type of action he wants the system to perform (e.g. create a new entry with *prepareNew*). The phonebook application will then present a form to enter the new contact information. The user can now complete the action by sending an *ok* message (or cancel, which we omit).

Our phonebook therefore has a *default* and an *action* state. We see that objects in the *default* state have three methods, while those in *action* have four. The methods applicable to the respective states can be easily identified by the types of their *self* variables. The triggers to switch from one state to the other are the business methods and *ok*, respectively.

The type system ensures that a method can only be called on a receiver that matches its expected receiver type exactly,

```
typedef entry = t. < name : t→string, number : t→string, super :<>>
typedef default = t.¡ < prepareNew : t→action, makeEditable : t→entry→action,
  confirmDelete : t→entry→action, super :<>>
typedef action = t.¡ < prepareNew : default→t, makeEditable : default→entry→t,
  confirmDelete : default→entry→t, ok : default→∘default, super :<>>

let Entry = clone(Object).
addMethod(name, λself:entry.“”).
addMethod(phone, λself:entry.“”) in

let WebPhonebook = clone(Object).
addMethod(prepareNew, λself : default.
  let curEntry = clone(Entry) in
  self.addMethod(ok, ¡λself : default./* save new entry */)).
addMethod(makeEditable, λself : default.
  λcurEntry : entry
  self.addMethod(ok, ¡λself : default./* save edited entry */)).
addMethod(confirmDelete, λself : default.
  λcurEntry : entry.
  self.addMethod(ok, ¡λself : default./* delete selected entry */))
```

Figure 4: Web phonebook business logic

after the method itself has been removed in the case of linear methods.[1] Thus, in the *default* state, the business methods can be called because they expect a receiver of type *default*, but the *ok* method cannot be called because it is not even part of the *default* state.

In the action state, the three business methods are still part of the type, but they cannot be called because these non-linear methods expect an object in the *default* state and the receiver is in the *action* state, which has the additional method *ok*. On the other hand, the *ok* method is linear, and it can be called in the *action* state because once you take the *ok* method out of the *action* type, you get the *default* type which is what the *ok* method expects.

Note that *ok* behaves differently depending on the action that is to be performed. Therefore each business method defines its own *ok* method.

The exact type matching in our system is essential in the case of linear objects to make sure that changes to the object are legal with respect its complete current type. In particular, when changing delegation the type system has to determine the exact new record type of the object, which can only be done on the basis of the exact old type of the object and its new delegate. Otherwise, the object could define a method with a name also used in the new delegate but with a different return type. If that method were not listed in the object type (which could happen if we allowed subtyping for linear objects) then the system would expect the wrong return type (the one defined in the delegate object) from a later call to that method.

The restriction of exact type tracking could be relaxed for aliased objects. Here, subtyping could be introduced to accept objects with more methods than expected, because the object type cannot change in a way that would introduce the problem mentioned above. Even though subtyping is well defined for record types, we elide this extension from our formal core system to keep it as simple as possible.

2.5 Expressive Power

The examples we have seen so far were mostly intended to illustrate syntax and semantics of EGO. This section will present higher-level examples in order to demonstrate the expressiveness of the language. In fact, one was already given in the previous section (figure 4) to illustrate the application of EGO to typestates. We will see typestates again in the examples that follow. The final one will implement the TCP socket from the introduction in EGO.

The examples rely on dynamic inheritance and adding new methods to objects over time. They are therefore not directly expressible in languages with static inheritance like Java. They are expressible in SELF, but SELF would not be able to statically guarantee that the program evaluation will succeed at runtime. Our system does guarantee successful evaluation of the presented examples by virtue of the type safety proof presented later.

Throughout the examples we rely on the intuition of the reader to assume the semantics of certain objects to which we merely refer by name. It would exceed the limitations of this paper to define a sufficiently large library explicitly on which interesting high-level examples can rely.

Consider the EGO program in figure 5. It models the workflow in a company between a manager, her secretary, and her designated worker. We first implement the secretary who can do some work. We also define a prototype worker who, no surprise, can also do some work. We define a concrete secretary as opposed to a prototype worker for purely pedagogical reasons. Both could be prototypes. Also note that we do not use the *trait* idiom known from SELF to generate a worker

[1] Removing the method from the type is necessary to ensure that linear methods cannot recursively call themselves. Recursive calls would break the invariant that no linear method is called more than once.

$$
\begin{array}{l}
typedef\ init = t.¡<sec : worker \to u. < doWork : secret \to unit >, \\
\quad setWorker : (t{-}setWorker)\ \multimap \\
\qquad u. < doWork : worker{\to}unit, workerSick : worker{\to}secret, super{:} <> > \multimap worker, \\
\quad super{:} <> > \\
typedef\ worker = t.¡<sec : t \to u. < doWork : secret \to unit >, \\
\quad myworker : secret \to u. < doWork : t \to unit, workerSick : t \to secret >, \\
\quad super{:}\ < doWork : t \to unit, workerSick : t \to secret >> \\
typedef\ secret = t.¡<sec : worker \to u. < doWork : secret \to unit >, \\
\quad myworker : t \to u. < doWork : worker \to unit, workerSick : worker \to t >, \\
\quad workerRecover : (t{-}workerRecover) \multimap worker,\ super :< doWork : t{\to}unit >>
\end{array}
$$

$$
\begin{array}{l}
let\ Secretary = change_linearity(clone(Object). \\
addMethod(doWork, \lambda self{:}secret.\lambda task{:}\tau. \ldots)) \ in
\end{array}
$$

$$
\begin{array}{l}
let\ WorkerPrototype = clone(Object). \\
addMethod(doWork, \lambda self{:}worker.\lambda task{:}\tau. \ldots) \\
addMethod(workerSick, \lambda self{:}worker\ .\ self.delegate(self.sec). \\
\quad addMethod(workerRecover, ¡\lambda self{:}secret\ .\ self.delegate(self.myworker))) \ in
\end{array}
$$

$$
\begin{array}{l}
let\ Manager = clone(Object). \\
addMethod(sec, \lambda self : secret.Secretary). \\
addMethod(setWorker, ¡\lambda self : init. \\
\quad ¡\lambda\ newworker{:}u\ .\ < doWork : worker \to unit, workerSick : worker \to secret >\ . \\
\quad self.addMethod(myworker, \\
\qquad \lambda self{:}secret\ .\ (self, newworker)). \\
\quad delegate(newworker)). \\
setWorker(change_linearity(clone(WorkerPrototype)))\ \ldots
\end{array}
$$

Figure 5: Using delegation to implement workflows

$$
\begin{array}{l}
let\ PowerSupply = clone(Object). \\
addMethod(generatePower, \lambda self{:}t.\ < generatePower : t \to power > \ldots.).
\end{array}
$$

$$
\begin{array}{l}
let\ On = change_linearity(clone(Object). \\
addMethod(getPower, \lambda self{:}t.¡<supply, on, off, super :< getPower >>\ . \\
\quad self.supply.generatePower)) \ in
\end{array}
$$

$$
let\ Off = change_linearity(clone(Object)) \ in
$$

$$
\begin{array}{l}
let\ PowerSwitch = clone(Object). \\
addMethod(on, \lambda self{:}t.¡<supply, on, off, super :<>> .self.delegate(On). \\
addMethod(off, \lambda self{:}t.¡<supply, on, off, super :< getPower >> .self.delegate(Off) \ in
\end{array}
$$

$$
\begin{array}{l}
let\ ps = clone(PowerSwitch).addMethod(supply, \\
\quad \lambda self{:}t.¡<supply, on, off, super :< getPower >> .PowerSupply) \ in \\
ps.on.getPower.getPower.off.on.getPower.off
\end{array}
$$

Figure 6: A kernel power network using composition and delegation

"class". Instead we define the worker prototype as an object to be cloned to create instances. We feel that this more closely resembles the real world where different workers are different autonomous individuals.

Finally we implement the manager who has fields for her secretary and her worker. By default, the manager forwards all the work she has to do to her worker. We do this simply by delegation. (That forces the complicated typing of *self* in the two *doWork* implementations.) We stress that this exactly models the situation in a real company, where work is delegated from one to the other person. Slightly confusing might be the implication that our manager does not even "see" the work items she delegates to her subordinate. But maybe this is not too unrealistic, either.

Now imagine the worker gets sick. We would invoke the *workerSick* method on our manager. That causes the manager to dump her work onto her secretary from now on. The secretary cannot get sick, so that's a safe guess. But also, the manager expects her worker to recover eventually. Thus she defines an additional method *workerRecover* to anticipate this event. Note that this changes the manager's signature. She is now in a different state, the "worker sick" state. *workerRecover* is defined to

$$
\begin{array}{l}
typedef\ open = t.¡ < port : t \to (t, int), read : t \to (t, \tau), \\
\quad write : t \to \tau \multimap t, close : t \multimap unit, super :<>> \\
\\
let\ Socket = clone(Object). \\
addMethod(bind, ¡\lambda self : t.¡ <> .self.(\ldots). \\
\quad addMethod(port, ¡\lambda self : open.(self, prt)). \\
\quad addMethod(listen, ¡\lambda self : t.¡ < port : open \to (t, int) > .self.(\ldots). \\
\qquad addMethod(accept, ¡\lambda self : t.¡ < port : open \to (t, int) > .self.(\ldots). \\
\qquad\quad addMethod(read, \lambda self : open.(self, result)). \\
\qquad\quad addMethod(write, \lambda self : open.¡\lambda data{:}\tau.\ldots; self). \\
\qquad\quad addMethod(close, ¡\lambda self{:}open.\ldots; unit))))
\end{array}
$$

Figure 7: A TCP socket object in Ego

be linear and thus will be consumed on invocation. The method will also redelegate to the now recovered worker, effectively transferring the manager back to her original state. As a final remark concerning states we point out that the manager is in a sort of initialization state before *setWorker* is called. Only then can she do (or rather, delegate) work.

Next we implement a kernel power network in figure 6. It consists of a power supply, an on-off-switch and a client that requests power. In this example we use delegation to model the different states of the power switch (on and off). Obviously, only the *On* object has a *getPower* method that forwards the power request to the supply configured for that object. Thus our client first has to connect the switch to the supply by adding the *supply* field. Then it can switch on, get power for a while, switch off, and on again to get more power.

The *On* and *Off* objects that implement the two controller states can be aliased by an arbitrary number of switches that all delegate to one of these two objects. The power supply is unique to each switch (both being physical devices) and therefore represented as an instance field to the switch. Without that field defined, the switch is not functional as the signatures for the *on* and *off* methods do not match. It can redirect to a different source later, though.

The power network example uses an implementation strategy that is quite the opposite to the workflow example above. In the power network, we use delegation to express states (on and off) and explicit forwarding (similar to composition in object-oriented programming) to transfer the power from the supply to the consumer. In the workflow example, on the other hand, we added and removed methods to change the state of the manager object. We used delegation to (implicitly) forward calls from one object to the other.

Finally, we look into the TCP socket example from the introduction section again. Figure 7 gives an implementation in Ego. We do not use delegation at all but rather manipulate the object with each method call. The implementation relies on linear methods to enforce that *bind*, *listen*, and *accept* are called exactly once. Each of these generates the following method; therefore a client must follow the prescribed call sequence.

We show as an example how *bind* also generates a field that contains the port on which the socket is going to listen in order to demonstrate that a real socket implementation is a full-blown data structure. Derived fields, as the *port* here, can be added to the object when they are available in Ego, effectively preventing reads from not yet defined fields.

The call to *accept* will generate *read*, *write*, and *close* methods. The first two can now be called an arbitrary number of times. They require a linear *self* and return it unchanged upon completion of the call. *close* also requires a linear *self* but does not give it back, effectively making the object inaccessible. Lending [2] or borrowing [4] for the methods returning *self* would make this explicit return unnecessary. We elide this possible extension to Ego for simplicity.

2.6 Summary

In the preceding sections we gave an informal introduction to Ego. We have seen in detail how programs can be implemented in the language. We discussed its handling of aliasing as well as the notion of typestates which it naturally supports through its method definitions. Finally we could express a number of relevant examples in Ego. We saw that delegation and dynamic method changes are somewhat substitutable, effectively allowing different programming styles.

The examples were complex enough to imagine that an ad-hoc Self programmer can introduce bugs that result in runtime errors. That motivates the need for static typechecking for such programs in order to make sure that all object manipulations and method invocations will succeed. Throughout this section we described the restrictions Ego imposes on the programmer to control Self's "power of simplicity". We have seen that they are loose enough to implement interesting programs in Ego, and although the current type system is somewhat complex we believe this can be simplified considerably in a practical system. It is the main result of this paper that these restrictions are also strong enough to ensure Ego's type safety. This will be formalized in the next section.

3. Formal Model

We now introduce the core Ego language to formalize the intuitions given above. This section contains the full dynamic semantics, the full static semantics, and a summarized type safety proof of Ego. The full type safety proof is available in [3].

3.1 Syntax

Figure 8 presents the syntax of our model. We do not include base types, control flow structures, exceptions, and subtyping into the model as they are well-known from the literature. We omit multiple inheritance and polymorphism as these are orthogonal to the typing issues at hand. Note that an overbar is used to represent a sequence.

Each program is an expression. An expression is a variable **(x)**, a value **(v)**, a clone of an object **(clone)**, a method invocation **(.m)**, an object delegation change **(delegate)**, the addition of a new method to an object or the change of a method body **(addMethod)**, a function application **(f a)** and a change of the type linearity of an object **(change_linearity)**. A method **(M)** is defined as a pair: the name of the method **(m)** and an expression that reduces to a method body. A method body definition is a lambda expression with a linearity type for the function. We will enforce that the outermost lambda types the receiver object. Our store **(S)** is a set of pairs: the location of the object and the object descriptor **(Odescr)**. An object descriptor is a pair: the location of the super object, and a sequence of methods defined for this object. There are four kinds of types: for variables **(t)**, for non-linear functions **(→)**, for linear functions **(⊸)** and finally for objects **(t.R)**. The object type is a recursive type where t is bound to R. The record type **(R)** is a list of the types of the methods **(B)** defined for the object and the type of the super object **(super)**. The type of a linear object is presented by ¡. We use **[¡]** to represent that the object might be linear or non-linear. Optional syntax is enclosed in [ ].

Instance variables are represented by parameterless methods. Locations L, l are not part of the source code. We assume to have a first object **(Object)** defined when we want to evaluate a program.

(programs)	p	::=	e
(expressions)	e	::=	$x \mid v \mid e.m \mid e_1.delegate(e_2) \mid clone(e)$
			$\mid e.addMethod(M) \mid apply(e_1, e_2)$
			$\mid change_linearity(e)$
(method sig)	M	::=	(m, e)
(values)	v	::=	$L \mid [¡]\lambda x : \tau.e_0$
(heap)	S	::=	$(Object, < Object, () >) \mid (L, Odescr), S$
(object desc)	$Odescr$	::=	$< l, (M_1, ..., M_n) >$
(types)	τ	::=	$t \mid \tau' \rightarrow \tau'' \mid t.R \mid \tau' \multimap \tau''$
(records)	R	::=	$[¡]<> \mid [¡]< B, super : \tau >$
	B	::=	$<> \mid < m : \tau, B >$
(heap location)	L, l		
(variable)	x		
(type variable)	t		
(method name)	m		

Figure 8: Syntax of the language, store, types. Angle bracket (<>) inside other angle brackets are deleted in our examples.

3.2 Dynamic Semantics

The dynamic semantics we defined for EGO is a standard small step operational semantics. The store **(S)** is a function from locations **(L)** to object descriptors **(Odescr)**. Figure 9 summarizes the rules for evaluating expressions. We describe each rule in turn.

$(R - Appl)$ shows how a method is applied to its arguments. We write $[v/x]e_0$ for the result of replacing x by v in expressions e_0.

$(R - LInvk)$ invokes a linear method on an object. The method is owned by the receiver and is linear. As the type system does not allow another call to that linear method we remove it from the store. The location **L** is passed as an argument to the method because **self** is not a free variable in the lambda expression. The type system does not support this in order to not have aliasing issues. The result of the reduction is a method apply with **L** as an argument and a store without the method **m** in it.

$(R - NInvk)$ invokes a non-linear method. The result of the reduction is the same as the one above except that the store is unchanged now: The type system allows the client to invoke a non-linear method more than once .

$(R - Clone)$ creates a new object from an existing one. The list of methods and the address of the super object are copied from the cloned object to the newly created location.

$(R - Deleg)$ changes the reference to the super object of the receiver object. The result of the reduction is the modified location of the receiver.

$(R - AddM)$ adds a new method to the receiver object. The result returned is the modified location of the receiver.

$(R - ChanMBd)$ changes the body of method **(m)** of the receiver.

$(R - ChanLin)$ does not effect the memory. It changes the linearity of an object from linear to non-linear. The result of the reduction is the location passed as argument for the expression.

$$\frac{}{([\text{¡}]\lambda x:\tau'.e_0)v, S \longrightarrow [v/x]e_0, S}\ \text{R}-\text{Appl}$$

$$\frac{S[L] = (l, ((m_1, v_1), ..., (m, v), ...)) \quad v \text{ is linear} \quad S' = S[L \rightarrow (l, ((m_1, v_1), ...))]}{L.m, S \longrightarrow vL), S'}\ \text{R}-\text{LInvk}$$

$$\frac{mbody(S[L], m) = v \quad v \text{ is nonlinear}}{L.m, S \longrightarrow vL, S}\ \text{R}-\text{NInvk}$$

$$\frac{S[L](l, \overline{M}) \quad L' \notin domain(S) \quad S' = S[L' \rightarrow (l, \overline{M})]}{clone(L), S \longrightarrow L', S'}\ \text{R}-\text{Clone}$$

$$\frac{S[L_1] = (l_1, \overline{M'}) \quad S[L_2] = (l_2, \overline{M''}) \quad S' = S[L_1 \rightarrow (L_2, \overline{M'})]}{L_1.delegate(L_2), S \longrightarrow L_1, S'}\ \text{R}-\text{Deleg}$$

$$\frac{S[L] = (l, \overline{M}) \quad m \notin dom(\overline{M}) \quad S' = S[L \rightarrow (l, (\overline{M}, (m, v)))]}{L.addMethod((m, v)), S \longrightarrow L, S'}\ \text{R}-\text{AddM}$$

$$\frac{\begin{array}{c} S[L] = (l, ((m_1, v_1), ..., (m, v), ...(m_n, v_n))) \\ m \in dom(\overline{M}) \\ S' = S[L \rightarrow ((m_1, v_1), ..., (m, v'), ..., (m_n, v_n))] \end{array}}{L.addMethod(m, v'), S \longrightarrow L, S'}\ \text{R}-\text{ChanMBd}$$

$$\frac{}{change_linearity(L), S \longrightarrow L, S}\ \text{R}-\text{ChanLin}$$

Figure 9: Evaluation rules for expressions

3.3 Static Semantics

Figure 11 presents the typing rules for expressions. Every typing rule has the standard form, $\Sigma; A \vdash e : \tau \Longrightarrow list_e$ that contains a store type **(Σ)**, an assumption list **(A or A')**, an expression that is typed **(e)**, the type of the expression **(τ)**and the list of linear objects **(list$_e$)** that are used to type the expression.

We use a type store Σ to store the types of our objects:

$$\Sigma ::= Object : t. <> \ \mid \Sigma; l : \tau$$

The assumption list A (or A') contains the types of the bound variables in the expression **e** that is typechecked. An assumption list, A, is defined as:

$$A ::= \cdot \mid A, x : \tau$$

We use · to present the empty assumption list. An assumption list is non-linear if each assumption x_i:τ_i in it has a non-linear type τ_i. Note that linear variables will be removed from the assumption list upon usage.

The type expression $t.[\text{¡}]< m_1 : \tau_1, ..., m_k : \tau_k, super : t'.[\text{¡}]< m'_1 : \tau'_1, ..., m'_j : \tau'_j >>$ is a type t with the property that when we invoke a method m_i for $1 \leq i \leq k$ or a method m'_i for $1 \leq i \leq j$ to any element x of this type, like $x.m_i$, the result has type τ_i or τ'_i with t substituted for $t.R$.

Let us describe each rule and give a brief justification with examples for selected cases. Note that the word location is used somewhat ambiguous because sometimes it refers to the label of a location and sometimes it is used to refer to an object.

$$\frac{\overline{M}[m] = v}{mbody((l, \overline{M}), m) = v}$$

$$\frac{m \notin dom(\overline{M})}{mbody((l, \overline{M}), m) = mbody(S[l], m)}$$

Figure 10: Rules for lookup of methods body.

However, what is meant is always obvious from the context.

($T - Loc$) A location is well-typed if it is defined in Σ. The list returned is empty if the location is non-linear or $\mathbf{L}$ if the location is linear.

($T - Method$) A non-linear method is well-typed if its body is well-typed. The restriction on the assumption list to be non linear is because we want each variable needed to type the expression to be non-linear so we can safely call the method more than once. There is no restriction on the arguments of the methods because if they are linear there is no way of duplicating them. The returned list is empty as the objects used here are all non-linear.

($T - LMethod$) A linear method, too, is well-typed if its body is well-typed. However in this rule, there is no restriction on the assumption list. Linear variables can safely be used in a linear method because it will be called only once during the program. The list returned is the one returned from the type rule applied to the method body.

($T - Var$) The type of a variable is the one that it has in the assumption list. The returned list is empty as there are no object used to type it. The assumption list has only the record to type the variable. The type system does not require to forget information in the assumption list during the typing of an expression.

($T - Kill$) This rule is used in the case we have to delete a record from the assumption list in order to typecheck an expression. We need it in typing cases like ¡λ¡$x : X \lambda y : Y.x$ where y can be non-linear. As long as it is not used, ($T - Kill$) can remove it from the context to type this linear method. The list returned is the same as the expression that is typed with the new assumption list.

($T - Copy$) This rule makes another copy of a non linear variable in the assumption list. We use it in cases like $\lambda x : Nat.x + x$ or $\lambda x : X.(x.addMethod(m, \lambda y : Y.x))$ where x is non-linear. The type system has to explicitly duplicate x in order to use it multiple times.

($T - Clone$) A *clone* expression is well-typed if $\mathbf{e}$ (the prototype object) is well-typed and the super object of $\mathbf{e}$ has a non-linear type (which is true automatically by virtue of ($T - Deleg$). The methods defined for the cloned object must all be non-linear in order not to copy references to linear objects through the back door. The new object created has a linear type.

($T - Invk$) A $.m$ expression is well-typed in the non-linear case if $\mathbf{e}$ (the receiver) and $\mathbf{m}$ are well-typed, both non-linear and the argument type of $\mathbf{m}$ is the same as the type of the object. The **mtype** function (see Figure 12) returns the type of the method that is invoked. The type system does not allow **self** as a free variable for aliasing issues. Instead it requires an explicit lambda for **self**. The following example will help clarifying the problem

$let\ o\ =\ clone(Object).$
$addMethod(m, \lambda_ : unit.self).$
$addMethod(m', \lambda_ : unit.self.m)\ in$
$(o.m).m'$ /*method m is invoked twice*/

($T - LInvk$) The difference of this rule from the one above is that $\mathbf{e}$ (the receiver) is linear and $\mathbf{m}$ can be either linear or non-linear. The new type of $\mathbf{e}$ does not allow the client to call $\mathbf{m}$ again if $\mathbf{m}$ is linear. We do that by deleting the record for $\mathbf{m}$ from the record type of the object. This prevents aliasing of linear objects in the assumption list (the stack). The following example illustrates the idea

$let\ obj\ =\ clone(Object).addMethod(m,$¡$\lambda$$.unit.self)\ in$
$let\ obj2\ =\ obj.m()\ in$
$let\ obj3\ =\ obj2.m()$ /*we have two references to obj*/

($T - AddM$) The type-system adds new methods only to linear objects because aliases to an object would not be aware of the new method. The assumption list used to type the expression is split to type the two different expressions, e_1 and e_2, in order to track the linearity of the objects. The list returned is the concatenation of the lists returned from the typing rules of $\mathbf{e}$ and $\mathbf{m}$.

($T - LChanMBd$) This rule checks if the object is linear and then checks if the new method body is well-typed. We can change the type of the method when the receiver is linear just like we can add new methods.

($T - ChanMBd$) This rule checks if the object is non-linear and that the new method body has the same type as the existing one. We do that for the same typing problem we can have in the $T - AddM$ or $T - Deleg$.

($T - Deleg$) This rule permits only to change delegation for a linear objects for the same reason the type-system only permits new methods for linear objects. The rule assures also to delegate to a non-linear object because if the type

system allows the client to delegate to linear objects then we effectively have more than one reference to it.

($T - ChanLin$) This rule changes the linearity of an object from linear to non-linear. The super object is non-linear anyway.

($T - Appl$) This rule checks if the first expression $\mathbf{e}_1$ has a function type and that the second expression $\mathbf{e}_2$ has the same type of the argument of e_1.

Figure 13 contains the rules for type-checking the store **S**. $list_i$ is a list of all linear objects used to type the location $S[L_i]$. The store **S** is well-typed if every location in the store is well-typed. A location **L** is well-typed if the super object which it inherit is well-typed and each method's body defined for that location is well-typed. $list_{e_i}$ is a list of linear objects used during the typing of the expression **e**.

3.4 Type Safety

In this section we describe the approach taken for proving type safety for our system.

As a program executes, the number of locations in the store can expand as **clone** operations are performed, and the types of locations can change as a result of method addition or delegation changes. We formalize the way the store type can change as a store extension operation $\Sigma' \geq_l \Sigma$. This judgment means that Σ' differs from Σ because of l in one of two cases :

1. Σ' may have an additional l in its domain

$$\frac{dom(\Sigma') = dom(\Sigma) \cup \{l\} \quad \forall l' \in dom(\Sigma).\Sigma(l') = \Sigma'(l')}{\Sigma' \geq_l \Sigma}$$

2. l was linear in Σ but is non-linear in Σ'

$$\frac{dom(\Sigma') = dom(\Sigma) \quad \forall l' \in \{dom(\Sigma) - l\}.\Sigma(l') = \Sigma'(l')}{\Sigma' \geq_l \Sigma}$$

The first case of Σ' and Σ is introduced by the **clone** typing rule and the second case is introduced by the **addMethod**, **delegate** or **change_linearity** typing rules. This lemma is used for the proof of T-AddM, T-Deleg and T-Appl.

1. (Preservation)If $\Sigma; \cdot \vdash e : \tau \Longrightarrow list_e$ and $\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$ and there are no duplicates in $list_e, list_S$ and $e, S \rightarrow e', S'$ then for some $\Sigma' \geq_l \Sigma$ we have $\Sigma'; \cdot \vdash e' : \tau \Longrightarrow list_e$ and $\Sigma'; \cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$ and there are no duplicates in $list_{e'}, list_{S'}$.

2. (Progress) If $\Sigma; \cdot \vdash e : \tau \Longrightarrow list_e$ and $\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$ then either

   (i) $e, S \rightarrow e', S'$ for some S' and e', or

   (ii) e is a value v

Preservation.. Preservation ensures that the type of an expression is preserved during its evaluation. For the proof of preservation, we need two properties about the substitution operation as it occurs in the case of function application and two lemmas.

Theorem 1 (Properties of Typing)
(i) (Weakening) If $\Sigma; A, A' \vdash e' : \tau' \Longrightarrow list_{e'}$ *and* τ *is nonlinear then* $\Sigma; A, x : \tau, A' \vdash e' : \tau' \Longrightarrow list_{e'}$.

(ii) (Substitution)If $\Sigma; A, x : \tau, A' \vdash e' : \tau' \Longrightarrow list_{e'}$ *and* $\Sigma; \cdot \vdash e : \tau \Longrightarrow list_e$ *then* $\Sigma; A, A' \vdash \{e/x\}e' : \tau' \Longrightarrow list_e, list_{e'}$

Proof: Property (i) follows directly by the T-Kill rule.

Property (ii) follows by a rule induction on the given derivation of $\Sigma; A, x : \tau, A' \vdash e' : \tau' \Longrightarrow list_{e'}$. Since typing and substitution are both compositional over the structure of the term, the only interesting cases are where e' is x or $[\mathfrak{i}]\lambda x.e_0$.

Case:(Rule T-Var). with $e' = x$. Then $\tau' = \tau, list_{e'} = \{\}$ and $\{e/x\}e' = \{e/x\}x = e$. But our assumption is $\Sigma; \cdot \vdash e : \tau \Longrightarrow list_e$ so we can conclude this by weakening property.

Case: (Rule T-[Non]Linear Method). with $e' = [\mathfrak{i}]\lambda y : \tau''.e_0$.

1. $x \neq y$; $\{e/x\}([\mathfrak{i}]\lambda y : \tau''.e_0) = [\mathfrak{i}]\lambda z : \tau''.(e_0\{z/y\}\{e/x\})$ where z $\notin$ FV(e') & FV(e) $\neq$ BV(e'). So we have to show $\Sigma; A, A' \vdash [\mathfrak{i}]\lambda z : \tau''.(e_0\{z/y\}\{e/x\}) : \tau'' \rightarrow \tau' \Longrightarrow list_e, list_{e_0}$, where $list_{e'} = list_{e_0}$ by definition. We apply for two times the inductive hypothesis for the new expression, $(e_0\{z/y\}\{e/x\})$.

2. If $x = y$ then $\{e/x\}e' = e'$. Then $\Sigma; A, A' \vdash e' : \tau' \Longrightarrow list_{e'}$.

$$\frac{\Sigma(L) = t.R}{\Sigma; \cdot \vdash L : t.R \Longrightarrow [L]}\ \mathrm{T-Loc}$$

$$\frac{\Sigma; A, x : \tau' \vdash e_0 : \tau'' \Longrightarrow \{\}}{\Sigma; A \vdash (\lambda x : \tau'.e_0) : \tau' \to \tau'' \Longrightarrow \{\}}\ \mathrm{T-Method\ (x \notin A, nonlinear\ A)}$$

$$\frac{\Sigma; A, x : \tau' \vdash e_0 : \tau'' \Longrightarrow list_{e_0}}{\Sigma; A \vdash (¡\lambda x : \tau'.e_0) : \tau' \multimap \tau'' \Longrightarrow list_{e_0}}\ \mathrm{T-LMethod\ (x \notin A)}$$

$$\frac{}{\Sigma; x : \tau \vdash x : \tau \Longrightarrow \{\}}\ \mathrm{T-Var}$$

$$\frac{\Sigma; A \vdash u : U \Longrightarrow list_u}{\Sigma; A, x : X \vdash u : U \Longrightarrow list_u}\ \mathrm{T-Kill}$$

$$\frac{\Sigma; A, x : X, x : X \vdash u : U \Longrightarrow list_u}{\Sigma; A, x : X \vdash u : U \Longrightarrow list_u}\ \mathrm{T-Copy\ (nonlinear\ X)}$$

$$\frac{\begin{array}{c}\Sigma; A \vdash e : t.[¡] < B, super : t'. < B', super : t''.R'' > \Longrightarrow list_e \\ \forall m \in B.\Sigma; A \vdash M(m) : \tau' \to \tau'' \Longrightarrow \{\}\end{array}}{\Sigma; A \vdash clone(e) : t.¡< B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_e}\ \mathrm{T-Clone}$$

$$\frac{\begin{array}{c}\Sigma; A \vdash e : t.¡< (..., m : \tau'[\to / \multimap]\tau''), super : t'.R > \Longrightarrow list_e \\ \tau' = t.¡< (..., [m{:}\tau' \to \tau''/\ ]\ ), super : t'.R >\end{array}}{\Sigma; A \vdash e.m : \tau'' \Longrightarrow list_e}\ \mathrm{T-LInvk}$$

$$\frac{\begin{array}{c}\Sigma; A \vdash e : t. < B, super : t'.R > \Longrightarrow list_e \\ mtype(m, t. < B, super : t'.R >) = \tau' \to \tau'' \\ \tau' = t. < B, super : t'.R >\end{array}}{\Sigma; A \vdash e.m : \tau'' \Longrightarrow list_e}\ \mathrm{T-Invk}$$

$$\frac{\begin{array}{c}\Sigma; A' \vdash e_2 : \tau \Longrightarrow list_{e_2} \\ m \notin B \\ \Sigma; A \vdash e_1 : t.¡< B, super : t'.R > \Longrightarrow list_{e_1}\end{array}}{\Sigma; A, A' \vdash e_1.addMethod((m, e_2)) : t.¡<< B, m : \tau >, super : t'.R > \Longrightarrow list_{e_1}, list_{e_2}}\ \mathrm{T-AddM}$$

$$\frac{\begin{array}{c}\Sigma; A \vdash e_1 : t.¡<< ..., m : \tau', ... >, super : t'.R > \Longrightarrow list_{e_1} \\ \Sigma; A' \vdash e_2 : \tau \Longrightarrow list_{e_2} \\ m \in B\end{array}}{\Sigma; A, A' \vdash e_1.addMethod(m, e_2) : t.¡<< ..., m : \tau, ... >, super : t'.R > \Longrightarrow list_{e_1}, list_{e_2}}\ \mathrm{T-LChanMBd}$$

$$\frac{\begin{array}{c}\Sigma; A \vdash e_1 : t. << ..., m : \tau', ... >, super : t'.R > \Longrightarrow list_{e_1} \\ \Sigma; A' \vdash e_2 : \tau' \Longrightarrow \{\} \\ \tau' = \tau \to \tau'' \\ m \in B\end{array}}{\Sigma; A, A' \vdash e_1.addMethod(m, e_2) : t.¡<< ..., m : \tau', ... >, super : t'.R > \Longrightarrow list_{e_1}}\ \mathrm{T-ChanMBd}$$

$$\frac{\begin{array}{c}\Sigma; A \vdash e_1 : t_1.¡< B_1, super : t'.R_1 > \Longrightarrow list_{e_1} \\ \Sigma; A' \vdash e_2 : t_2. < B_2, super : t''.R_2 > \Longrightarrow list_{e_2}\end{array}}{\Sigma; A, A' \vdash e_1.delegate(e_2) : t_1.¡< B_1, super : (t_2. < B_2, super : t''.R_2 >)[t_1/t_2] > \Longrightarrow list_{e_1}, list_{e_2}}\ \mathrm{T-Deleg}$$

$$\frac{\Sigma; A \vdash e : t.¡< B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_e}{\Sigma; A \vdash change_linearity(e) : t. < B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_e}\ \mathrm{T-ChanLin}$$

$$\frac{\Sigma; A \vdash e_1 : \tau'[\to / \multimap]\tau'' \Longrightarrow list_{e_1} \quad \Sigma; A' \vdash e_2 : \tau' \Longrightarrow list_{e_2}}{\Sigma; A, A' \vdash e_1 e_2 : \tau'' \Longrightarrow list_{e_1}, list_{e_2}}\ \mathrm{T-Appl}$$

Figure 11: Static semantics of expressions. {} represents the empty list.

$$\frac{m \in B \quad B < ..., m : \tau, ... >}{mtype(m, t.[\mathrm{i}] < B, super : t'.R >) = \tau}$$

$$\frac{m \notin B}{mtype(m, t.[\mathrm{i}] < B, super : t'.R >) = mtype(m, t'.R)}$$

Figure 12: Rules for lookup methods type in a record type.

$$\frac{\forall L_i \in dom(\Sigma).\Sigma; \cdot \vdash S(L_i) : \Sigma(L_i) \Longrightarrow list_i}{\Sigma; \cdot \vdash S : \Sigma \Longrightarrow concat(list_i)} \ \mathrm{T-Store}$$

$$\frac{\Sigma(l) = t'.R \quad \forall (m_i, e_i).\Sigma; \cdot \vdash e_i : \tau_i \Longrightarrow list_{e_i}}{\Sigma; \cdot \vdash < l, ((m_1, e_1), ..., (m_n, e_n)) > : t.[\mathrm{i}] < \overline{\tau_i}, super : t'.R > \Longrightarrow concat(list_{e_i})} \ \mathrm{T-Odescr}$$

Figure 13: Static semantic of Store

■

Next, we define two lemmas that are useful in ensuring that linear methods and objects remain unaliased as the program executes.

Lemma 2
if $\Sigma; A \vdash e : \tau \Longrightarrow list_e$ *and* $\Sigma' \geq_l \Sigma$ *and* $l \notin list_e$ *then* $\Sigma'; A \vdash e : \tau \Longrightarrow list_e$

This lemma is used to proof that if the old list $(list_{e_1}, list_{e_2}, list_S)$ of linear objects has no duplicates and part of that list $(list_{e_1}, list_S)$ has changed $(list_{e'_1}, list'_S)$ because of an evaluation rule then the modified list $(list_{e'_1}, list_{e_2}, list'_S)$ has no duplicates. This is because if there are no duplicates in the bigger list there could not possibly be duplicates in the smaller one.
Proof: The proof follows by induction on the typing rule of e, $\Sigma; A \vdash e : \tau \Longrightarrow list_e$. ■

Lemma 3
For any rule, $e, S \rightarrow e', S'$, *where* $\Sigma; \cdot \vdash e : \tau \Longrightarrow list_e$, $\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$ *and no duplicates* $list_e, list_S$, *and for* $\Sigma' \geq_l \Sigma$ *and* $\Sigma'; \cdot \vdash e' : \tau \Longrightarrow list_{e'}$, $\Sigma'; \cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$ *and no duplicate in* $list_{S'}, list_{e'}$ *then* $\{list_{S'}\} \cup \{list_{e'}\} \subseteq \{list_S\} \cup \{list_e\} \cup \{l\}$.

Proof: By rule induction on the derivation $e, S \rightarrow e', S'$ ■

Theorem 4 (Preservation)
If $\Sigma; \cdot \vdash e : \tau \Longrightarrow list_e$ *and* $\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$ *there are no duplicates in* $list_e, list_S$ *and* $e, S \rightarrow e', S'$ *then for some* $\Sigma' \geq_l \Sigma$ *and memory* S' *we have* $\Sigma'; \cdot \vdash e' : \tau \Longrightarrow list_{e'}$ *and* $\Sigma'; \cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$ *and there are no duplicates in* $list_{e'}, list_{S'}$.

Proof: By rule induction on the derivation of $e, S \rightarrow e', S'$.

Case(Rule T-Invk).

$$\frac{e, S \rightarrow e', S'}{e.m, S \rightarrow e'.m, S'}$$

$e, S \rightarrow e', S'$	Subderivation
$\Sigma; \cdot \vdash e.m : \tau'' \Longrightarrow list_e$	Assumption
$\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$	Assumption
No duplicate $list_e, list_S$	Assumption
$\Sigma; \cdot \vdash e : t. < B, super : t'.R >= \tau' \Longrightarrow list_e$	By inversion
$mtype(m, t. < B, super : t'.R >) = \tau' \rightarrow \tau''$	By inversion
$\Sigma' \geq_l \Sigma$, $\Sigma'; \cdot \vdash e' : t. < B, super : t'.R > \Longrightarrow list_{e'}$	By i.h.
$\Sigma'; \cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$ and no duplicates in $list_{e'}, list_{S'}$	By i.h.
$\Sigma'; \cdot \vdash e'.m : \tau'' \Longrightarrow list_{e'}$	By rule

Same proof for the T-LInvk rule except the method invokated is not deleted.

Case.

$$\frac{mbody(S[L], m) = v \quad v \text{ is nonlinear}}{L.m, S \to vL, S}$$

$\Sigma; \cdot \vdash L.m : \tau'' \Longrightarrow \{\}$	Assumption
$\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$	Assumption
No duplicate $list_S$	Assumption
$\Sigma; \cdot \vdash L : t. < B, super : t'.R >= \tau' \Longrightarrow \{\}$	By inversion
$\Sigma; \cdot \vdash M(m) : \tau' \to \tau'' \Longrightarrow \{\}$	By inversion of T-Odescr
$\Sigma; \cdot \vdash vL : \tau'' \Longrightarrow \{\}$	By rule

Case.

$$\frac{S[L] = (l, ((m_1, v_1), ..., (m, v), ...)) \quad v \text{ is linear} \quad S' = S[L \to (l, ((m_1, v_1), ...))]}{L.m, S \to vL, S'}$$

$\Sigma; \cdot \vdash L.m : \tau'' \Longrightarrow L$	Assumption
$\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$	Assumption
No duplicate $list_S, L$	Assumption*
$\Sigma; \cdot \vdash L : t.¡< (..., m : \tau), super : t'.R >\Longrightarrow L$	By inversion
$\tau' = t.¡< (..., [m : \tau]), super : t'.R > \Sigma; \cdot \vdash M(m) : \tau \Longrightarrow list_m$	By inversion of T-Odescr

Supposing that m has a linear type otherwise the proof will be similar to the one above.

let $\Sigma' = \Sigma[L \to t.¡< (...), super : t'.R >]$	
$L \notin list_v$	By assumption*
$\Sigma'; \cdot \vdash v : \tau' - \circ \tau'' \Longrightarrow list_v$	By lemma 2
$\Sigma'; \cdot \vdash L : t.¡< (...), super : t'.R >= \tau' \Longrightarrow L$ /' By rule	
$\Sigma'; \cdot \vdash vL : \tau'' \Longrightarrow list_v, L$	By rule
$\forall m \in (...).\Sigma'; \cdot \vdash M(m) : \tau \Longrightarrow list_m$ /' By Assumption* and lemma 2	
$\Sigma'; \cdot \vdash S'(L) : \Sigma'(L) \Longrightarrow list_L - list_m$	**
$\Sigma'; \cdot \vdash S' : \Sigma' \Longrightarrow list_S - list_m$	By rule and **

We have to prove that $\Sigma'; \cdot \vdash [L/this]v : ([t''/t]\tau)[t''.|¡< (...)[t''/t], super : t'.R > ./t] \Longrightarrow L, list_m$. We do it using the Subsitution property. We have that $\Sigma'; this : t'' \vdash m : \tau[t''/t] \Longrightarrow list_m$ and $\Sigma'; \cdot \vdash L : t''.¡< B[t''/t], super : t'.R >\Longrightarrow L$ and so from substitution property we have that $\Sigma'; \cdot \vdash [L/this]v : ([t''/t]\tau)[t''.¡< (...)[t''/t], super : t'.R > /t''] \Longrightarrow L, list_m$

No duplicate $list_S - list_m, L, list_m$ È By Assumption*

Case(Rule T-Clone).

$$\frac{e, S \to e', S'}{clone(e), S \to clone(e'), S'}$$

$e, S \to e', S'$	Subderivation
$\Sigma; \cdot \vdash clone(e) : t.¡< B, super : t'. < B', super : t''.R'' >>\Longrightarrow list_e$	Assumption
$\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$	Assumption
No duplicate $list_e, list_S$	Assumption
$\Sigma; \cdot \vdash e : t.[¡] < B, super : t'. < B', super : t''.R'' >>\Longrightarrow list_e$	By inversion
$\forall m \in B.\Sigma; \cdot \vdash M(m) : \tau' \to \tau'' \Longrightarrow \{\}$	By inversion
$\Sigma' \geq_l \Sigma, \Sigma'; \cdot \vdash e' : t.[¡] < B, super : t'. < B', super : t''.R'' >>\Longrightarrow list_{e'}$	By i.h.
$\Sigma'; \cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$	By i.h.
No duplicate $list_{e'}, list_{S'}$	By i.h.
$\forall m \in B.\Sigma'; \cdot \vdash M(m) : \tau' \to \tau'' \Longrightarrow \{\}$	By lemma 2
$\Sigma'; \cdot \vdash clone(e') : t.¡< B, super : t'. < B', super : t''.R'' >>\Longrightarrow list_{e'}$	By rule

Case.

$$\frac{S[L] = (l, \overline{M}) \quad L' \notin dom(S) \quad S' = S[L' \to (l, \overline{M})]}{clone(L), S \to L', S'}$$

$\Sigma;\cdot \vdash clone(L) : t.¡< B, super : t'. < B', super : t''.R'' >>\Longrightarrow [L]$ Assumption
$\Sigma;\cdot \vdash S : \Sigma \Longrightarrow list_S$ Assumption
No duplicate $[L], list_S$ Assumption**
$\Sigma;\cdot \vdash L : t.[¡] < B, super : t'. < B', super : t''.R'' >>\Longrightarrow [L]$ By inversion
$\forall m \in B.\Sigma;\cdot \vdash M(m) : \tau' \to \tau'' \Longrightarrow \{\}$ By inversion
let $\Sigma' = \Sigma[L' \to t.¡< B, super : t'. < B', super : t''.R'' >>]$ then
$\Sigma';\cdot \vdash L' : t.¡< B, super : t'. < B', super : t''.R'' >>\Longrightarrow L'$ By rule
$\forall l \in dom(\Sigma).\Sigma(l) = \Sigma'(l)$ By definition of $\Sigma' \geq_l \Sigma$
$\forall l \in dom(S)$ and $\forall m \in S(l)$ then
$\Sigma';\cdot \vdash M(m) : \tau \Longrightarrow list_m$ By lemma 2
$\Sigma';\cdot \vdash S'(L') : \Sigma'(L') \Longrightarrow \{\}$ By rule and lemma 2
$\Sigma';\cdot \vdash S' : \Sigma' \Longrightarrow list_S$ By rule
No duplicate $list_S, L'$ By assumption** & L' $\notin \Sigma$

Case(Rule T-AddM).

$$\frac{e_1, S \to e_1', S'}{e_1.addMethod((m, e_2)), S \to e_1'.addMethod((m, e_2))}$$

$e_1, S \to e_1', S'$ Subderivation
$\Sigma;\cdot \vdash e_1.addMethod((m, e_2)) : t.¡<< B, m : \tau >, super : t'.R >$
$\Longrightarrow list_{e_1}, list_{e_2}$ Assumption
$\Sigma;\cdot \vdash S : \Sigma \Longrightarrow list_S$ Assumption
No duplicate $list_{e_1}, list_{e_2}, list_S$ Assumption
$\Sigma;\cdot \vdash e_1 : t.¡< B, super : t'.R >\Longrightarrow list_{e_1}$ By inversion
$\Sigma;\cdot \vdash e_2 : \tau \Longrightarrow list_{e_2}$ By inversion
$\Sigma' \geq_l \Sigma$, $\Sigma';\cdot \vdash e_1' : t.¡< B, super : t'.R >\Longrightarrow list_{e_1'}$ By i.h.
$\Sigma';\cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$ By i.h.
No duplicate $list_{e_1'}, list_{S'}$ By i.h.
$l \in list_(e_1'), list_{S'}$

No duplicate in $list_{e_1'}, list_{e_2}, list_{S'}$ because if $list_{e_2}, liste_1, list_S$ has no duplicate then from lemma 3 $list_{e_2}, list_{e_1'}, list_{S'}$ has no duplicate.

$l \notin list_{e_2}$
$\Sigma';\cdot \vdash e_2 : \tau \Longrightarrow list_{e_2}$ By lemma 2
$\Sigma';\cdot \vdash e_1'.addMethod((m, e_2)) : t.¡<< B, m : \tau >, super : t'.R >$
$\Longrightarrow list_{e_1'}, list_{e_2}$ By rule & Lemma 2

$$\frac{e_2, S \to e_2', S'}{e_1.addMethod((m, e_2)), S \to e_1.addMethod((m, e_2'))}$$

Symmetric to the previous case.

$$\frac{S[L] = (l, \overline{M}) \quad S' = S[L \to (l, (\overline{M}, (m, v)))]}{L.addMethod((m, v)), S \to L, S'}$$

$\Sigma;\cdot \vdash L.addMethod((m, v)) : t.¡<< B, m : \tau >, super : t'.R >\Longrightarrow L, list_v$ Assumption
$\Sigma;\cdot \vdash S : \Sigma \Longrightarrow list_S$ Assumption
No duplicate $L, list_v, list_S$ Assumption*
$\Sigma;\cdot \vdash L : t.¡< B, super : t'.R >\Longrightarrow L$ By inversion
$\Sigma;\cdot \vdash v : \tau \Longrightarrow list_v$ By inversion
let $\Sigma' = \Sigma[L \to t.¡<< B, m : \tau >, super : t'.R >]$
then $\Sigma';\cdot \vdash L : t.¡<< B, m : \tau. >, super : t'.R >\Longrightarrow L$ By rule
$L, list_v$ has no duplicate $\longrightarrow L \notin list_v$
$\Sigma';\cdot \vdash v : \tau \Longrightarrow list_v$ By rule & lemma 2
$\forall l \in dom(\Sigma).\Sigma(l) = \Sigma'(l)$ By definition of $\Sigma' \geq_L \Sigma$
$\forall l \in \{dom(S) - L\}$ and $\forall m \in S(l)$ then
$\Sigma';\cdot \vdash M(m) : \tau \Longrightarrow list_m$ By lemma 2
$\Sigma';\cdot \vdash S' : \Sigma' \Longrightarrow list_S, list_v$ By rule
No duplicate $L, list_S, list_v$ By Assumption*

The cases for **T-LChanMBd** and **T-ChanMBd** are similar to **T-Addm**.

Case(T-Deleg).

$$\frac{e_1, S \to e_1', S'}{e_1.delegate(e_2), S \to e_1'.delegate(e_2), S'}$$

$e_1, S \to e_1', S'$	Subderivation
$\Sigma;\cdot \vdash e_1.delegate(e_2) : t_1.¡< B_1, super : t_2. < B_2, super : t''.R_2 > [t_1/t_2] >\Longrightarrow list_e$	Assumption
$\Sigma;\cdot \vdash S : \Sigma \Longrightarrow list_S$	Assumption
No duplicate $list_e, list_S$	Assumption
$\Sigma;\cdot \vdash e_1 : t_1.¡< B_1, super : t'.R_1 >\Longrightarrow list_{e_1}$	By inversion
$\Sigma;\cdot \vdash e_2 : t_2. < B_2, super : t''.R_2 >\Longrightarrow list_{e_2}$	By inversion
$\Sigma' \geq_l \Sigma, \Sigma';\cdot \vdash e_1' : t_1.¡< B_1, super : t'.R_1 >\Longrightarrow list_{e_1'}$	By i.h.
$\Sigma';\cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$	By i.h.
No duplicate $list_{e_1'}, list_{S'}$	By i.h.
$l \in list_{e_1'}, list_{S'}$	By definition of $\Sigma' \geq_l \Sigma$

No duplicate in $list_{e_1'}, list_{e_2}, list_{S'}$ because if $list_{e_2}, liste_1, list_S$ has no duplicate then from lemma 3 $list_{e_2}, list_{e_1'}, list_{S'}$ has no duplicate.
$l \notin list_{e_2}$

$\Sigma';\cdot \vdash e_2 : t_2. < B_2, super : t''.R_2 >\Longrightarrow list_{e_2}$	By lemma 2
$\Sigma';\cdot \vdash e_1'.delegate(e_2) : t_1.¡< B_1, super : t_2.R_2 >\Longrightarrow list_{e_1'}, list_{e_2}$	By rule

Case.

$$\frac{S[L_1] = (l_1, \overline{M'}) \quad S[L_2] = (l_2, \overline{M''}) \quad S' = S[L_1 \to (L_2, \overline{M'})]}{L_1.delegate(L_2), S \to L_1, S'}$$

$\Sigma;\cdot \vdash L_1.delegate(L_2) : t_1.¡< B_1, super : t_2. < B_2, super : t''.R_2 >>\Longrightarrow L_1$	Assumption
$\Sigma;\cdot \vdash S : \Sigma \Longrightarrow list_S$	Assumption
No duplicate $L_1, list_S$	Assumption*
$\Sigma;\cdot \vdash L_1 : t_1.¡< B_1, super : t'.R_1 >\Longrightarrow L_1$	By inversion
$\Sigma;\cdot \vdash L_2 : t_2. < B_2, super : t''.R_2 >\Longrightarrow \{\}$	By inversion
let $\Sigma' = \Sigma[L_1 \to t_1.¡< B_1, super : t_2. < B_2, super : t''.R_2 >>]$	
then $\Sigma';\cdot \vdash L_1 : t_1.¡< B_1, super : t_2.R_2 >\Longrightarrow L_1$	By rule
$\Sigma';\cdot \vdash L_2 : t_2. < B_2, super : t''.R_2 >\longrightarrow \{\}$	By lemma 2
$\forall l \in dom(\Sigma).\Sigma(l) = \Sigma'(l)$	By definition of $\Sigma' \geq_L \Sigma$
$\forall l \in \{dom(S) - L_1\}$ and $\forall m \in S(l)$ then	
$\Sigma';\cdot \vdash M(m) : \tau \Longrightarrow list_m$	By Assumption* and lemma 2
$\Sigma';\cdot \vdash S' : \Sigma' \Longrightarrow list_S$	By rule
No duplicate $L_1, list_S$	By assumption*

Case(T-Appl).

$$\frac{e_1, S \to e_1', S'}{apply(e_1, e_2), S \to apply(e_1', e_2), S'}$$

$e_1, S \to e_1', S'$	Subderivation
$\Sigma;\cdot \vdash apply(e_1, e_2) : \tau'' \Longrightarrow list_{e_1}, list_{e_2}$	Assumption
$\Sigma;\cdot \vdash S : \Sigma \Longrightarrow list_S$	Assumption
No duplicate $list_{e_1}, list_{e_2}, list_S$	Assumption
$\Sigma;\cdot \vdash e_1 : \tau'[\to / -\circ]\tau'' \Longrightarrow list_{e_1}$	By inversion
$\Sigma;\cdot \vdash e_2 : \tau'' \Longrightarrow list_{e_2}$	By inversion
$\Sigma' \geq_l \Sigma, \Sigma';\cdot \vdash e_1' : \tau'[\to / -\circ]\tau'' \Longrightarrow list_{e_1'}$	By i.h.
$\Sigma';\cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$	By i.h.
No duplicate $list_{e_1'}, list_{S'}$	By i.h.

No duplicate in $list_{e_1'}, list_{e_2}, list_{S'}$ because if $list_{e_2}, liste_1, list_S$ has no duplicate then from lemma 3 $list_{e_2}, list_{e_1'}, list_{S'}$ has no duplicate.

$l \in list_{e_1'}, list_{S'} \longrightarrow l \notin list_{e_2}$ (*)
$\Sigma'; \cdot \vdash e_2 : \tau' \Longrightarrow list_{e_2}$ By lemma 2
$\Sigma'; \cdot \vdash apply(e_1', e_2) : \tau'' \Longrightarrow list_{e_1'}, list_{e_2}$ By rule

$$\frac{e_2, S \to e_2', S'}{apply(e_1, e_2), S \to apply(e_1, e_2'), S'}$$

Symmetric to the previous case.

Case.

$$\frac{}{([¡]\lambda x : \tau'.e_0)v, S \to [v/x]e_0, S}$$

$\Sigma; \cdot \vdash ([¡]\lambda x : \tau'.e_0)v : \tau'' \Longrightarrow list_{e_0}, list_v$ Assumption
$\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$ Assumption
No duplicate $list_{e_0}, list_v, list_S$ Assumption*
$\Sigma; \cdot \vdash [¡]\lambda x : \tau'.e_0 : \tau'[\to / -\circ]\tau'' \Longrightarrow list_{e_0}$ By inversion
$\Sigma; \cdot \vdash v : \tau' \Longrightarrow list_v$ By inversion
$\Sigma; \cdot \vdash [v/x]e_0 : \tau'' \Longrightarrow list_{e_0}, list_v$ By Substitution Property

Case(T-ChanLin).

$$\frac{e, S \to e', S'}{change_linearity(e), S \to change_linearity(e'), S'}$$

$e, S \to e', S'$ Subderivation
$\Sigma; \cdot \vdash change_linearity(e) : t. < B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_e$ Assumption
$\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$ Assumption
No duplicate $list_e, list_S$ Assumption
$\Sigma; \cdot \vdash e : t.¡< B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_e$ By inversion
$\Sigma' \geq_l \Sigma, \Sigma'; \cdot \vdash e' : t.¡< B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_{e'}$ By i.h.
$\Sigma'; \cdot \vdash S' : \Sigma' \Longrightarrow list_{S'}$ By i.h.
No duplicate $list_{e'}, list_{S'}$ By i.h.
$\Sigma'; \cdot \vdash change_linearity(e') : t. < B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_{e'}$ By rule

Case.

$$\frac{}{change_linearity(L), S \to L, S}$$

$\Sigma; \cdot \vdash change_linearity(L) : t. < B, super : t'. < B', super : t''.R'' >> \Longrightarrow L$ Assumption
$\Sigma; \cdot \vdash S : \Sigma \Longrightarrow list_S$ Assumption
No duplicate $L, list_S$ Assumption*
$\Sigma; \cdot \vdash L : t.¡< B, super : t'. < B', super : t''.R'' >> \Longrightarrow L$ By inversion
let $\Sigma' = \Sigma[L \to t.¡< B, super : t'. < B', super : t''.R'' >>]$
then $\Sigma'; \cdot \vdash L : t. < B, super : t'. < B', super : t''.R'' >> \Longrightarrow \{\}$ By rule
$\forall l \in dom(\Sigma).\Sigma(l) = \Sigma'(l)$ By definition of $\Sigma' \geq_L \Sigma$
$\forall l \in \{dom(S) - L\}$ and $\forall m \in S(l)$ then
$\Sigma'; \cdot \vdash M(m) : \tau \Longrightarrow list_m$ By Assumption* and lemma 2
$\Sigma'; \cdot \vdash S : \Sigma' \Longrightarrow list_S$ By rule

■

Progess.. This asserts that the computation of closed well-typed expressions will never get stuck. The critical observation behind the proof of the progress theorem is that a value of function type will indeed be a function and a value of object type be an object. We state these critical properties as an inversion lemmas, because they are not immediately syntactically obvious.

Lemma 5 (Value inversion)
(i) If $\Sigma; \cdot \vdash v : t.R \Longrightarrow list_v$ *then* $v = L$.

(ii) If $\Sigma;\cdot \vdash v : \tau'[\multimap / -\circ]\tau'' \Longrightarrow list_v$ *then* $v = [¡]\lambda x : \tau'.e_0$.

Proof: We distinguish cases on v value and then apply inversion to the given typing judgment.

Property (i).

Case: $v = L$.. We are done because v = L.

Case: $v = [¡]\lambda x : \tau'.e_0$.. Then we would have $\Sigma;\cdot \vdash [¡]\lambda x : \tau'.e_0 : t.R$, which is impossible by inspection of the typing rules.

Property (ii).

Case: $v = L$.. Then we would have $\Sigma;\cdot \vdash L : \tau'[\multimap / -\circ]\tau''$, which is impossible by inspection of the typing rules.

Case: $v = [¡]\lambda x : \tau'.e_0$.. We are done because $v = [¡]\lambda x : \tau'.e_0$. ■

Now we can prove the progress theorem.

Theorem 6 (Progress)
If $\Sigma;\cdot \vdash e : \tau \Longrightarrow list_e$ *and* $\Sigma;\cdot \vdash S : \Sigma \Longrightarrow list_S$ *then either*

(i) e *is a value* v, *or*

(ii) $e, S \to e', S'$ *for some* S' *and* e'.

Proof: By induction on the derivation of the typing judgment, analyzing all possible cases.

Case(T-Loc).

$$\frac{\Sigma(L) = t.R}{\Sigma;\cdot \vdash L : t.R \Longrightarrow [L]}$$

Then L value.

Case(T-[Non]Linear Method).

$$\frac{\Sigma; A, x : \tau' \vdash e_0 : \tau'' \Longrightarrow list_{e_0}}{\Sigma; A \vdash ([¡]\lambda x : \tau'.e_0) : \tau' -\circ \tau'' \Longrightarrow list_{e_0}}$$

Then $[¡]\lambda x : \tau'.e_0$ value.

Case(T-Var).

$$\frac{}{\Sigma; x : \tau \vdash x : \tau \Longrightarrow \{\}}$$

This case is impossible since the context can not be empty. Same for T-Kill and T-Copy.

Case(T-Clone).

$$\frac{\begin{array}{c}\Sigma;\cdot \vdash e : t.[¡] < B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_e \\ \forall m \in B.\Sigma;\cdot \vdash M(m) : \tau' \to \tau''\end{array}}{\Sigma;\cdot \vdash clone(e) : t.¡< B, super : t'. < B', super : t''.R'' >> \Longrightarrow list_e}$$

Either $e, S \to e', S'$ for e' and S' or e is a value v	By i.h.
$e, S \to e', S'$	First subcase
$clone(e), S \to clone(e'), S'$	By rule
e is a value v	Second subcase
$e = L$	By value inversion
$clone(L), S \to L', S'$	By rule

Case(T-[L]Invk).

$$\frac{\begin{array}{c}\Sigma; A \vdash e : t. < B, super : t'.R > \Longrightarrow list_e \\ mtype(m, < B, super : t'.R >) = \tau' \to \tau'' = \tau\end{array}}{\Sigma; A \vdash e.m : \tau[t. < B, super : t'.R > /t] \Longrightarrow list_e}$$

Either $e, S \to e', S'$ for e' and S' or e is a value v — By i.h.
$e, S \to e', S'$ — First subcase
$e.m, S \to e'.m, S'$ — By rule
e value — Second subcase
$e = L$ — By value inversion
$L.m, S \to [L/this]v, S'$ — By rule

The same proof for **T-LInvk**

Case(T-AddM).

$$\frac{\Sigma; A', this : t.[/¡] < B, super : t'.R >\vdash e_2 : \tau'[\to / - \circ]\tau'' \Longrightarrow list_{e_2} \quad m \notin B \quad \Sigma; A \vdash e_1 : t.¡< B, super : t'.R > \Longrightarrow list_{e_1}}{\Sigma; A, A' \vdash e_1.addMethod((m, e_2)) : t << B, m : \tau'[\to / - \circ]\tau'' >, super : t'.R > \Longrightarrow list_{e_1}, list_{e_2}}$$

Either $e_1, S \to e_1', S'$ for e_1' and S' or e_1 is a value v — By i.h.
$e_1, S \to e_1', S'$ — First subcase
$e_1.addMethod((m, e_2)), S \to e_1'.addMethod((m, e_2)), S'$ — By rule
e_1 value — Second subcase
Either $e_2, S \to e_2', S'$ for e_2' and S' or e_2 is a value v — By i.h.
$e_2, S \to e_2', S'$ — First subsubcase
$e_1.addMethod((m, e_2)), S \to e_1.addMethod((m, e_2')), S'$ — By rule
e_2 is a value v — Second subsubcase
$e_1 = L$ — By value inversion
$e_2 = [¡]\lambda x : \tau'.e_0$ — By value inversion
$L.addMethod((m, v)), S \to L, S'$ — By rule

The proof for **T-LChanMBd** and **T-ChanMBd** is similar to the proof above.

Case(T-Delegate).

$$\frac{\Sigma; A \vdash e_1 : t_1.¡< B_1, super : t'.R_1 > \Longrightarrow list_{e_1} \quad \Sigma; A' \vdash e_2 : t_2.R_2 \Longrightarrow list_{e_2}}{\Sigma; A, A' \vdash e_1.delegate(e_2) : t_1.¡< B_1, super : t_2.R_2[t_1/t_2] > \Longrightarrow list_{e_1}, list_{e_2}}$$

Either $e_1, S \to e_1', S'$ for e_1' and S' or e_1 is a value v — By i.h.
$e_1, S \to e_1', S'$ — First subcase
$e_1.delegate(e_2), S \to e_1'.delegate(e_2), S'$ — By rule
e_1 value — Second subcase
Either $e_2, S \to e_2', S'$ for e_2' or e_2 value — By i.h.
$e_2, S \to e_2', S'$ — First subsubcase
$e_1.delegate(e_2), S \to e_1.delegate(e_2'), S'$ — By rule
e_2 value — Second subsubcase
$e_1 = L$ — By value inversion
$e_2 = L$ — By value inversion
$L_1.addMethod(L_2), S \to L_1, S'$ — By rule

Case(T-ChanLin).

$$\frac{\Sigma; A \vdash e : t.¡< B, super : t'. < B', super : t''.R'' >>\Longrightarrow list_e}{\Sigma; A \vdash change_linearity(e) : t. < B, super : t'. < B', super : t''.R'' >>\Longrightarrow list_e}$$

Either $e, S \to e', S'$ for e' and S' or e is a value v — By i.h.
$e, S \to e', S'$ — First subcase
$change_linearity(e), S \to change_linearity(e'), S'$ — By rule
e value — Second subcase
$e = L$ — By value inversion
$change_linearity(L), S \to L, S$ — By rule

Case(T-Appl).

$$\frac{\Sigma; A \vdash e_1 : \tau'[\rightarrow / - \circ]\tau'' \Longrightarrow list_{e_1} \quad \Sigma; A' \vdash e_2 : \tau' \Longrightarrow list_{e_2}}{\Sigma; A, A' \vdash \mathit{apply}(e_1, e_2) : \tau'' \Longrightarrow list_{e_1}, list_{e_2}}$$

Either $e_1, S \rightarrow e_1', S'$ for e_1' and S' or e_1 is a value v	By i.h.
$e_1, S \rightarrow e_1', S'$	First subcase
$\mathit{apply}(e_1, e_2), S \rightarrow \mathit{apply}(e_1', e_2), S'$	By rule
e_1 value	Second subcase
Either $e_2, S \rightarrow e_2', S'$ for e_2' or e_2 value	By i.h.
$e_2, S \rightarrow e_2', S'$	First subsubcase
$\mathit{apply}(e_1, e_2), S \rightarrow \mathit{apply}(e_1, e_2'), S'$	By rule
e_2 value	Second subsubcase
$e_1 = [\mathfrak{i}]\lambda x : \tau'.e_0$	By value inversion
$e_2 = L$	By value inversion
$([\mathfrak{i}]\lambda x : \tau'.e_0)(v), S \rightarrow [v/x]e_0, S$	By rule

■

4. Related work

This section summarizes related work in language foundations, aliasing, and state-based method dispatch. SELF [15] was the first prototype-based language and also defined mechanisms for dynamic modifications of object definitions. We largely adopt that expressiveness but make it statically checkable.

Abadi and Cardelli use prototype-based object calculi to study issues of subtyping, quantification, and the typing of the receiver object *self*. Fisher and Mitchell describe a delegation-based object calculus with subtyping and type inference [9]. Compared to these systems, our work focuses on the orthogonal issue of ensuring that dynamic type changes to a linear object are safe.

Our work builds on Philip Wadler's linear type system [16], which in turn builds on a foundational linear logic developed by Girard [11]. The concept of linear types in [16] is used more for resources that should not be duplicated. Resources like files should have linear type in order to not accidentally duplicate or discard them. In contrast, our system uses linear types to allow the client to safely change the type of an object. In addition, we show how linear first-class functions like those presented in [16] can be naturally used in an object-oriented system.

Predicate classes [5] and its more general form, predicate dispatch [8] support method dispatch based on predicates over the run-time state of the object. When a message is sent in these systems, the predicates of all relevent methods are evaluated, and the method chosen is the one with the most specific predicate that evaluates to true. Dynamic inheritance and dynamic method modification is a complimentary way to get similar behavior: instead of dispatching indirectly based on the state of an object, the state is encoded through the dispatch hierarchy. These mechanisms are probably complimentary, with each appropriate in different situations; one advantage of our approach is that it can change the type of an object, rather than just which method is selected at run time.

Typestates were initially introduced by [14] for procedural programming languages. [6] defines a resource-controlling system for such languages based on keys. Keys can optionally be parameterized with typestates. This class of systems is formally modeled in [12] as refinement types that layer additional, changing resources on a conventional static type system. All these approaches do not consider inheritance and effectively only allow linear types. Thus they are unsuitable for object-oriented languages.

The State design pattern in [10] allows implementing different behavior for a method depending on the main object's state. However, there is no way of statically restricting the available methods for a state. [7] defined a model for tracking typestates in object-oriented languages. In particular, they address the issue of typestates in the presence of subtyping. In our work, objects have a dynamically changing type instead of a changing typestate layered on top of a fixed type.

5. Conclusions

EGO offers a prototype-based language that has expressiveness, simplicity and a static typechecker. The expressiveness follows from dynamic inheritance, adding methods, changing method bodies, and even changing method types dynamically. Its simplicity follows from the lack of the class concept, from the concept of cloning instead of instantiation, and from the unification of fields and methods.

EGO imposes restrictions on the programmer in order to control SELF's "power of simplicity". These are loose enough to allow interesting programs using EGO's dynamic features. But these restrictions are also strong enough to ensure EGO's static type safety. Its static typechecker provides a safer and more efficient paradigm than SELF: EGO programs will only contain valid method invocations.

In future work, we plan to investigate adding more advanced object-oriented language features to the system, including multiple inheritance, parametric polymorphism, and multiple dispatch. Allowing subtyping for non-linear objects is easy, but more research will have to be devoted to finding ways of supporting subtyping along with dynamic inheritance. Recent developments in typestate systems may provide a path forward here [7].

Acknowledgments

This work was supported in part by the High Dependability Computing Program from NASA Ames cooperative agreement NCC-2-1298, NSF grant CCR-0204047, and the Army Research Office grant number DAAD19-02-1-0389 entitled "Perpetually Available and Secure Information Systems."

References

[1] M. Abadi, L. Cardelli. A theory of objects. Springer, 1996.

[2] J. Aldrich, V. Kostandinov, C. Chambers. Alias Annotations for Program Understanding. Proc. Object-Oriented Programming, Systems, Languages, and Applications, November 2002.

[3] A. Bejleri. A type checked prototype-based model with linearity. Draft senior thesis, published as Carnegie Mellon Technical Report CMU-ISRI-04-142, December 2004.

[4] J. Boyland. Alias burying: Unique variables without reads. Journal : Software—Practice and Experience 31(6):533-553, May 2001.

[5] C. Chambers. Predicate classes. Proc. European Conference on Object-Oriented Programming, 1993.

[6] R. DeLine, M. Fähndrich. Enforcing High-Level Protocols in Low-Level Software. Proc. Programming Language Design and Implementation, June 2001.

[7] R. DeLine, M. Fähndrich. Typestates for Objects. Proc. European Conference on Object-Oriented Programming, 2004.

[8] M. D. Ernst, C. Kaplan, C. Chambers. Predicate Dispatching: A Unified Theory of Dispatch. Proc. European Conference on Object-Oriented Programming, 1998.

[9] K. Fisher, J. C. Mitchell. A Delegation-based Object Calculus with Subtyping. Proc. Fundamentals of Computation Theory, 1995.

[10] E. Gamma, R. Helm, R. Johnson, J. Vlissides. Design Patterns: Elements of Reusable Object-Oriented Software. Addison-Wesley, October 1994.

[11] J.-Y. Girard. Linear logic. Theoretical Computer Science 50(1):1-102, 1987.

[12] Y. Mandelbaum, D. Walker, R. Harper. An effective theory of type refinements. Proc. International Conference on Functional Programming, 2003.

[13] R. Milner, M.Tofte, R. Harper, D. MacQueen. The Definition of Standard ML (Revised). MIT Press, 1997.

[14] R. E. Strom, S. Yemini. Typestate: A programming language concept for enhancing software reliability. IEEE Trans. Software Engineering 12(1):157-171, January 1986.

[15] D. Ungar, R. B. Smith. Self: The power of simplicity. Proc. Object-Oriented Programming Systems, Languages, and Applications, 1987.

[16] P. Wadler. Linear types can change the world! In M. Broy and C. Jones, editors, Programming Concepts and Methods, North Holland, 1990.