

Trust the clones

Sophia Drossopoulou¹ and James Noble²

¹ Department of Computing,
Imperial College, London
`scd@doc.ic.ac.uk`

² School of Mathematics, Statistics, and Computer Science,
Victoria University of Wellington,
Wellington, New Zealand
`kjx@mcs.vuw.ac.nz`

Abstract. Most object-oriented programming languages provide some way to clone objects — to produce a new object that is a copy of an existing object. Typically this is either a shallow clone that clones only one object, or a deep clone that clones objects recursively. In practice, programmers have to write custom cloning methods for each of their classes, and this boilerplate code is tedious to write and requires care to write correctly. Inspired by ownership types, we propose a system that annotates classes to express how their instances should be cloned. Furthermore, we show how these annotations can be used to generate cloning methods for these classes.

1 Introduction

Background Ownership types were first suggested in 1998; their aim was to characterise aliases, and thus to control the topology of objects on the heap [19]. Since then, they have attracted tremendous interest; they have been developed in several brands and variations, and they have been put to several different uses, e.g. memory management, encapsulation, effect systems, avoidance of race conditions, locations, parallel programming, program verification etc.

One example given in the seminal ECOOP paper [19], was cloning, whereby all the objects “inside” a certain other object would automatically be copied when the enclosing object was copied.

The aim of supporting cloning, although appealing, has — to our knowledge — not been further pursued in the context of static ownership types. Interestingly, this aim was recently pursued in the static analysis world [14], whereby copy policies are expressed by annotating code so as to specify the maximally allowed sharing between an object and its clone, and a type and effects system checks whether the copy policy is adhered to.

In this paper we propose the opposite approach: the copy code is *generated* out of *clone annotations* given with the types of fields, and thus, adherence to the copy policy is automatically guaranteed.

Overview of our approach The key problem we aim to address is determining which objects to copy. Generally, the set of objects which have to be cloned when cloning an object o does not solely depend on o — it also depends on the object which started the the cloning process — the “originator” of the cloning.

Consider the object structure shown in Fig. 1, where the small boxes represent objects of a class as written in the box, and the labelled arrows represent fields. For example, we have an object of class `Union`, with a field called `students` which points to an object of class `StudentList`.

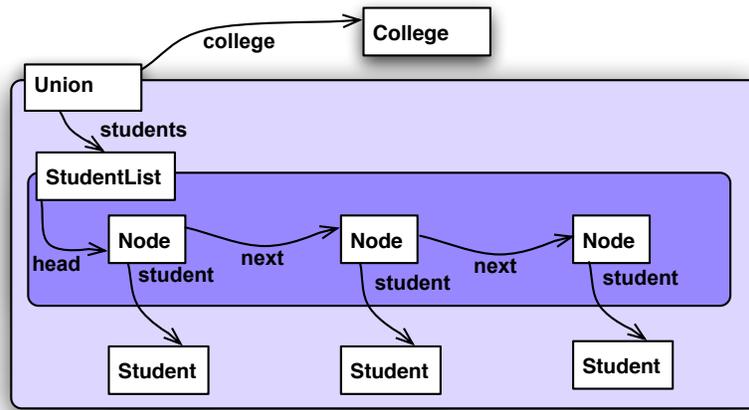


Fig. 1. An example of cloning domains

Cloning a `Node` object when the originator is itself requires only the current object (the `Node`) to be cloned. Cloning a `Node` as a result of cloning a `StudentList` requires cloning all the successors of that node too, so that all the `Nodes` in the list are cloned. On the other hand, cloning the `StudentList` should not clone the `Students` pointed to by the `Nodes` in the list — however, if the `StudentList` is cloned as part of cloning the `Union` object, then the `Students` should also be cloned. Therefore, we say that the `Node` objects belong to the “cloning domain” of the `StudentList` object, as well as that of the `Union` object, while the `StudentList` and the `Student` objects belong to the cloning domain of the `Union` object.

The two cloning domains are depicted graphically in Fig. 1, through the two shaded rounded boxes. The cloning domains introduce a hierarchy among objects: when e.g. o belongs to the cloning domain of o' and o' belongs to the cloning domain of o'' , then o belongs to the cloning domain of o'' .

The hierarchy shown in Fig. 1 is very similar to that shown in various works on ownership types. In both cases, objects are put into boxes, and the boxes belong to objects. The difference is in the meaning of these boxes. In some cases, the boxes guarantee encapsulation[24], while in others they guarantee

domination[8], or restrict who may modify the objects[16], or guarantee common de-allocation[1]. In our current work, the boxes guarantee that cloning the object which owns the box will also clone all objects in that particular box.

Note that the diagram in Fig. 1 presents a “ghost view” of the objects; the boxes, *i.e.*, the cloning domain information is *not* available at runtime. All the information we have is that `StudentList` will have a pointer field called `head` that points to its nodes; those `Nodes` have two pointer fields, one called `next` that points to the next `Node` and another called `student` that points to its `Student`.

To clone an object, we have to answer this question: for each object we reach, considering each of its fields in turn, *without* the runtime topological ownership information shown in Fig. 1: should we clone the object to which the field refers?

Organisation of our paper In Sect. 2 we describe our solution — clone annotated types — and give an example. In Sect. 3 we describe how clone annotated types can be used to decide whether or not a particular pointer should be cloned, and then in Sect. 4 we show how to generate the appropriate cloning methods. In Sect. 5 we describe the guarantees provided by our approach. In Sect. 6 we discuss related work and conclude.

2 Cloning Annotations

We propose to answer the question by introducing *cloning annotations* on field types. These annotations follow the tradition of ownership types [6, 8]. Every class is defined so as to take one or more formal *cloning parameters*, c . Thus, cloning annotated types have the form $C\langle c_1, \dots, c_n \rangle$, where C is a class, and c_1, \dots, c_n are the cloning parameters. The type $C\langle c_1, \dots, c_n \rangle$ expresses that the object has class C , and that it will be cloned whenever the object standing for c_1 is cloned; the remaining copy annotations, c_2, \dots, c_n may be used to annotate fields from that object and so describe when they are to be cloned.

The class definition determines the scope of these parameters — similarly to generic parameters. As in traditional ownership types, these parameters stand for objects. Again as in traditional ownership type systems, types are formed by classes followed by *actual cloning parameters*, ca , which may be any cloning parameter which is in scope, or `this`. The syntax is shown in Fig. 2.

$$\begin{array}{l}
 \textit{ClassDecl} ::= \text{class } C\langle \bar{c} \rangle \\
 \qquad \qquad \qquad \{ \overline{\textit{FieldDecl}} \ \overline{\textit{MethDecl}} \} \\
 \textit{FieldDecl} ::= \textit{FieldType } f \\
 \textit{FieldType} ::= C\langle \bar{ca} \rangle \\
 c \qquad \qquad ::= \textit{Identifier} \qquad \text{clone parameters} \\
 ca \qquad \qquad ::= c \mid \text{this} \qquad \text{clone arguments}
 \end{array}$$

Fig. 2. Syntax extracts

Note that the syntax of clone arguments does not allow any annotations which indicate that the cloning domain is unknown; thus, we do not support annotations such as `norep` from [10], or `any` from [9] or existentially bound arguments as in [2].

We use the term *cloning domain* of an object o to describe all the objects which have to be cloned when o is cloned. The first cloning parameter, c_1 , in the annotation of a field f determines that the object pointed at by f belongs to the cloning domain of c_1 : thus f must be cloned whenever c_1 is cloned. The remaining cloning parameters are used to determine cloning for the fields of f .

```
class Node<c1, c2>{
    Node<c1,c2> next;
    Student<c2> student;
}
class StudentList<c>{
    Node<this,c> head;
}
class Union<c>{
    StudentList<this> students;
    College<c> college;
}
```

Fig. 3. Lists example

Fig. 3 shows cloning annotated class declarations that match the structure from Fig. 1. When a `Node` is cloned, then a new object of class `Node` needs to be created, and its fields need to be initialised according to those of the old object. When a `StudentList` is cloned, then a new object of class `StudentList` needs to be created, and *all* accessible `Nodes` will have to be cloned — this is denoted by the `head` field’s first cloning parameter being `this`. Finally, when a `Union` is cloned, then the `StudentList` will be cloned, all accessible `Nodes` and all accessible `Students` will have to be cloned.

3 Cloning Domains and Cloning Paths

Cloning domains have no runtime representation. In order to clone objects, we can only access their subordinate objects via *cloning paths* traced through objects’ fields.

Inspecting Fig. 1 and Fig. 3 we can see how cloning paths lead to objects in different cloning domains. Consider paths beginning from a `StudentList` object. The objects at the paths `this.head`, or `this.head.next`, or `this.head.next.next` belong to the cloning domain of that `StudentList` object, but the object reachable through `this.head.next.student` does not. For paths starting at a `Union` object, objects reachable at `this.students`, or at `this.students.head`,

or through `this.students.head.next`, and finally also at `this.students.head.next.next.student` belong to the cloning domain of the `Union` object.

Note that cloning domains are nested, and thus object may belong to more than one cloning domain. For example, for a `Union` object, `this.students`, `this.students.head` and `this.students.head.next` all belong to the `Union`'s cloning domain. On the other hand, the objects at `this.students.head` and `this.students.head.next` also belong to the cloning domain of the `StudentList` at the `this.students`, as well as to the `Union`'s cloning domain. Furthermore, it is possible for an object to be reachable from only one other object, but not belong to its cloning domain. For example, `this.head.next` is only reachable via `this.head`, but it does not belong to `this.head`'s cloning domain.

We now formalise the notion of cloning paths and cloning domains. We first define paths and path types as follows:

$p ::= \mathbf{this} \mid p.f$	paths
$PT ::= C\langle \overline{cap} \rangle$	path types
$cap ::= ca \mid p$	actual path cloning parameters

Path types express types which are relative to the current object `this`, and where paths can be used as actual cloning parameters. For example `Node(this.students, this)` describes an object of class `Node`, which belongs to the cloning domain of the object at `this.students`, and whose `student` field belongs to the cloning domain of `this`.

In Fig. 4 we define a judgment of the form $C \vdash p : D\langle cap_1, \dots, cap_n \rangle$ which says that from an object of class `C`, the path `p` leads to an object in the cloning domain of `cap1`. Furthermore, the actual path cloning parameters `cap1`, ... `capn` may be used to characterize the cloning domain of `p.f`, where `f` is a field of `p`.

$$\frac{\text{class } C\langle c_1, \dots, c_n \rangle \{ \dots \}}{C \vdash \mathbf{this} : C\langle c_1, \dots, c_n \rangle}$$

$$\frac{\text{class } D\langle c_1, \dots, c_m \rangle \{ \dots E\langle ap_1, \dots, ap_n \rangle f \dots \}}{C \vdash \mathbf{this}.\bar{f} : D\langle cap_1, \dots, cap_m \rangle}$$

$$\frac{}{C \vdash \mathbf{this}.\bar{f}.f : E\langle ap_1, \dots, ap_n \rangle [cap_1, \dots, cap_m / c_1, \dots, c_m] [\mathbf{this}.\bar{f} / \mathbf{this}]}$$

Fig. 4. Path types for paths

The first rule says that in class `C` the path `this` has the type given by the cloning parameters of the class `C`. This gives, for example, that $\text{Union} \vdash \mathbf{this} : \text{Union}\langle c \rangle$.

The second rule says that from class `C` the path `this. \bar{f} .f` has the type of `f` as given in `f`'s declaration, but where `c1`, ... `cm`, the cloning parameters of the class declaring field `f`, are replaced by `cap1`, ... `capm`, the actual cloning parameters for `this. \bar{f}` . Moreover, any occurrence of `this` in the type of `f` is replaced by `this. \bar{f}` . The first replacement is standard in the literature, and the second replacement is novel. This rule gives $\text{Union} \vdash \mathbf{this}.\text{students} : \text{StudentList}\langle \mathbf{this} \rangle$ and $\text{Union} \vdash \mathbf{this}.\text{students}.\text{head} : \text{Node}\langle \mathbf{this}.\text{students}, \mathbf{this} \rangle$.

Here are some more examples of the path type judgment:

$$\begin{aligned} \text{Node} &\vdash \text{this.next.next} : \text{Node}\langle c1, c2 \rangle \\ \text{StudentList} &\vdash \text{this.head.next.next} : \text{Node}\langle \text{this}, c \rangle \\ \text{StudentList} &\vdash \text{this.head.student} : \text{Student}\langle c \rangle \\ \text{Union} &\vdash \text{this.students.head} : \text{Node}\langle \text{this.students}, \text{this} \rangle \\ \text{Union} &\vdash \text{this.students.head.student} : \text{Student}\langle \text{this} \rangle \\ \text{Union} &\vdash \text{this.college} : \text{College}\langle c \rangle \end{aligned}$$

Note that the judgments only talk about paths, and makes no requirement about consistent views of the heap. So far, nothing precludes two paths p_1 and p_2 , which are aliases at runtime, but which have different path types, eg $C \vdash p_1 : D\langle \text{cap}_1, \dots, \text{cap}_n \rangle$ and $C \vdash p_2 : D\langle \text{cap}'_1, \dots, \text{cap}'_n \rangle$ and $\text{cap}_1 \neq \text{cap}'_1$. However, as we see in section 5.3, we require consistent views on the heap in order to obtain completeness, *i.e.*, that all objects which should be cloned, are, indeed, cloned.

We can now characterise for a given class C the set of paths to the objects in their cloning domain, $ClnDom(C)$. This set consists of **this**, and all these paths $p.f$ whose path-type has as first argument another path p' (remember that all paths start at **this**), *i.e.*, $C \vdash p.f : D\langle p', \dots \rangle$, and which lie exclusively within the cloning domain of C , *i.e.*, $p \in ClnDom(C)$.

$$ClnDom(C) = \{ \text{this} \} \cup \{ p.f \mid C \vdash p.f : D\langle p', \dots \rangle \text{ for a class } D, \text{ and a path } p', \text{ and where } p \in ClnDom(C) \}$$

Thus, we obtain that

$$\begin{aligned} \text{this.next.next} &\notin ClnDom(\text{Node}) \\ \text{this.head.next.next} &\in ClnDom(\text{StudentList}) \\ \text{this.college} &\notin ClnDom(\text{Union}) \\ \text{this.students.head.next.next.student} &\in ClnDom(\text{Union}) \end{aligned}$$

The requirement that the complete path should lie within the cloning domain (i.e that $p \in ClnDom(C)$) is relevant for the case where the cloning-owners are not dominators. For example, if the **College** object had a field **f** to some object which lied under the **StudentList** object, then **this.college.f** would not belong to $ClnDom(\text{Union})$.

4 Generating cloning methods

We have seen that the cloning behaviour of an object depends on the originator of the cloning action. For example, cloning a **Node** behaves differently when called as a result of cloning a **Union** or cloning a **StudentList**. This means that a single `clone()` method defined on each object cannot clone that object correctly in all contexts.

One approach would be to write a cloning method for every class that not only clones the object itself, but also clones every object within its cloning domain. Such a clone method would have to navigate through the appropriate paths and duplicate the appropriate objects. This approach is highly non-modular, since it exposes the internals of potentially every class.

Instead, we propose an overloaded parametric `clone(...)` method for each class, with additional Boolean actual parameters to describe whether or not its additional cloning domains should be cloned. For a class $C\langle c_1, \dots, c_n \rangle$, we generate one overloaded parametric method, `clone(Boolean s1, ... Boolean sn, Map m)`. The value of s_i determines whether objects from the cloning domain c_i are to be cloned too. The original `clone` method can then be reimplemented to call the parametric cloning method, requesting only the object itself is cloned, by passing **false** to the additional cloning parameters.

The `Map m` formal parameter to the cloning method is to ensure that we clone the current object only if it has not already been cloned. We use the conventional solution of maintaining a table mapping original to cloned objects — something like a Smalltalk `IdentityDictionary` or Java 1.4's `IdentityHashMap` is ideal. If the object is not yet in the map, then the method will call the corresponding `clone` methods for all fields which need to be cloned, and will make aliases to all other fields — otherwise we simply use the clone stored in the map.

We show some examples of cloning methods in subsection 4.1 and then define the general case in subsection in 4.2.

4.1 A clone method for our example

For class `Node`, the basic `clone` method is called when the originator is the node itself, and simply delegates to the parametric cloning method. This method will be part of the interface of the cloning library.

```
Node clone() {
    this.clone(false, false, new IdentityHashMap())
}
```

The parametric cloning method distinguishes whether the cloning domains of the clone parameters are to be cloned too, and accepts a `Map` to ensure each object is only cloned once. This method is internal to our system, and will not be visible outside the cloning library.

```
Node clone(Boolean s1, Boolean s2, Map m) {
    Object n = m.get(this);
    if ( n != null) then {
        return (Node)n;
    } else {
        Node clone=new Node();
        m.put(this,clone);
        clone.next= s1 ? this.next.clone(s1,s2,m) : this.next;
        clone.student= s2 ? this.student.clone(s2,m) : this.student;
        return clone;
    }
}
```

4.2 Code generation for clone

In Fig. 5 we show the generation of the basic cloning method for a class C . It calls the parametric cloning method on itself, where all the boolean parameters are set to `false`. This indicates that except for the cloning domain of the current object, no other cloning domains are “active”.

In Fig. 6 we show the generation of the parametric cloning method for C .

As a first step, in the first four lines of the method, we need to determine whether the object is to be cloned. If the object is the outcome of a previous cloning action, *i.e.* is in the lookup table m , then nothing happens. Otherwise, a new object of that class is created, and entered into the table.

```
C clone(){
    return this.clone(false1, ...falseq, newIdentityHashMap());
}
```

where q is the number of clone parameters of class C

Fig. 5. The function `clone()` for class C

```
C clone(Boolean s1...Boolean sq, Map m){
    Object o = m.get(this)
    if o ≠ null then
        return (C)o;
    else{
        C o' = new C();
        m.put(this, o');
        o'.f1 = s1,1 ? this.f1.clone(s1,1...s1,k1, m) : this.f1;
        ... = ...
        o'.fn = sn,1 ? this.fn.clone(sn,1...sn,kn, m) : this.fn;
        return o';
    }
}
```

where

$\{f_1, \dots, f_n\}$ are the fields defined in class C

and where, for all $i \in 1..n$:

$(fType(C(s_1 \dots s_n), f_i))[true/this] = C_i \langle s_{i,1} \dots s_{i,k_i} \rangle$

for some classes C_1, \dots, C_n .

Fig. 6. The function `clone(s1 ... sn, m)` for class C

As a second step, we need to initialize the fields of the newly created object, and determine which further objects have to be cloned. We assume that $\{f_1, \dots, f_n\}$ are all the fields defined in class C . Consider any field f_i . We calculate the type of f_i , and because anything that belongs to the cloning domain of the current object also belongs to the cloning domain of the owner

of the current object, we replace any occurrence of `this` by `true`. This gives that $C_i\langle s_{i,1}\dots s_{i,k_i}\rangle = fType(C\langle s_1\dots s_n\rangle, f_i)[\text{true}/\text{this}]$ for some C_i . We find the cloning parameter corresponding to $s_{i,1}$: If $s_{i,1}$ is `true`, then the object at f_i has to be cloned and the result has to be assigned to the field f_i . If $s_{i,1}$ is `false`, then the object at f_i should not be cloned, but the field f_i of the new object needs to alias its value. This is why we emit the conditional expression:

```
o'.f_i = s_{i,1} ? this.f_i.clone(s_{i,1}\dots s_{i,k_i}, m) : this.f_i.
```

5 Guarantees of cloning

We now discuss the properties of our cloning operation, namely

- termination,
- soundness, *i.e.*, the objects which are cloned are those which are reachable through paths in the cloning domain,
- completeness, *i.e.*, if the system also guarantees a consistent view of the heap then all objects from the cloning domain will be cloned.

We do not present an operational semantics, because our system works for any imperative object oriented language with the standard meaning for field read and write, conditional expressions and method calls. Because the system works on the basis of statically known paths and their types, the runtime representation of objects does not need to be enhanced with cloning domain information - we do not even need any ghost information for the purposes of the formal argument.

Nevertheless, we expect an operational semantics of form $H, \phi, \text{expr} \rightsquigarrow H', v$, where H and H' are heaps, ϕ is a stack frame, `expr` is an expression, v stands for values, that values include `null` the booleans `true` and `false`, and addresses; and that addresses are represented as ι , ι' etc. Furthermore, we expect that $H(\iota, f)$ returns the value of field f of the object at ι , that H stores the class for each object, and that the stack frames ϕ map identifiers to values. Finally, we expect that expressions include field read, field write, conditional expressions, and that these have the standard meaning.

The guarantees we describe in this section are applicable to richer languages, such that *e.g.*, support exceptions, or re-entrant method calls, or in fact, any further sequential control features. This is so, because our work makes guarantees about the effect of the code produced by our approach. This code only uses the language features listed above, and does not call any user-defined functions. The guarantees about the effects of the code remain valid regardless about the features used in the context that may be calling the generated `clone` methods, provided that this context is in the sequential setting.

In the following, we distinguish between the original call of the `clone()` method — without any arguments — and the subsequent recursive calls of the `clone(...)` methods which have at least one argument.

5.1 Termination

Lemma 1. *For all heaps H and variables $x \in \text{dom}(\phi)$ there exists an address ι and a heap H' , such that $H, \phi, x.\text{clone}() \rightsquigarrow H', \iota$.*

Proof Sketch We can show that during execution of `clone()` the values of fields in the objects in the original heap H do not change, and that all receivers of any of the recursive calls of the `clone(...)` methods were accessible from the originator object, ι , through a path p in the original heap H .

Therefore, the transitive calls of the `clone` method do not clone objects from the original heap H more than once, and do not clone any of the newly created objects. This gives a finite upper bound to the possible number of calls to the recursively called `clone(...)` methods. Since these methods do not contain any iteration, the fact that the number of possible recursive calls is finite, guarantees termination.

5.2 Soundness

With respect to the new heap, the cloning operation $H, \phi, \mathbf{x.clone}() \rightsquigarrow H', \iota$, makes the following four guarantees:

1. It creates a new object for the object \mathbf{x} ,
2. H' , the new heap, is homomorphic to H , the old heap,
3. In H' there is a self-contained part that corresponds to the old heap,
4. In H' there is a part that corresponds to the part in the old heap, which had to be copied according to the $ClnDom$ function.

In order to express the last point formally, we define the set of objects that should be copied, as follows:

$$ToCopy : Addr \times Heap \rightarrow Power(Addr)$$

$$ToCopy(\iota)_H = \{ H(\iota.\bar{\mathbf{f}}) = \iota' \mid \mathbf{this}.\bar{\mathbf{f}} \in ClnDom(\mathbf{C}), \text{ and } \mathbf{C} \text{ the class of } \iota \text{ in } H \}$$

We can now express soundness of the `clone` operation as follows:

Lemma 2. *If $H, \phi, \mathbf{x.clone}() \rightsquigarrow H', \iota$, then there exists a mapping $\alpha : dom(H') \rightarrow dom(H)$ such that*

1. ι is new in H , and $\alpha(\iota) = \phi(\mathbf{x})$,
2. $\alpha(H'(\iota', \mathbf{f})) = H(\alpha(\iota'), \mathbf{f})$ for all $\iota' \in dom(H')$ and fields \mathbf{f} .
3. $\alpha|_{dom(H)}$ is the identity function, and $\alpha|_{dom(H') \setminus dom(H)}$ is injective.³
4. $\alpha(dom(H') \setminus dom(H)) \subseteq ToCopy(\phi(x))_H$.⁴

Proof Sketch: We can show that during execution of `clone` for the newly created `Map` object \mathbf{m} , that the range of \mathbf{m} is the original heap H , and that any object being mapped comes from the new heap H' , *i.e.*, for all objects ι', ι'' , if $\mathbf{m.get}(\iota'') = \iota'$ then $\iota'' \notin dom(H)$, and $\iota' \in dom(H)$. Furthermore, we can show that the lookup $\mathbf{m.get}(\dots)$ is injective, *i.e.*, for all objects ι', ι'' , if $\mathbf{m.get}(\iota') = \mathbf{m.get}(\iota'')$ then $\iota' = \iota''$.

³ **Notation:** For an f which is a mapping from A to B , and for $A' \subseteq A$, we use the term $f|_{A'}$ to the function f as restricted to the domain A' .

⁴ **Notation:** For an f which is a mapping from A to B , and for $A' \subseteq A$, we use the term $f(A')$ to describe the image of f from A' .

Because of injectivity, we can construct the mapping $\alpha(\iota')$ as the inverse of \mathbf{m} in the cases where ι' is the image of an original object in H (ie $\mathbf{m.get}(\iota') \neq \mathbf{null}$), and the identity otherwise.

Parts 1 and 3 follow from the construction of α and from the properties of the `Map` object mentioned above.

Part 2 follows from the body of the methods `clone(...)`.

For Part 4, we show that for any recursive application of any of the `clone(...)` methods, all the objects cloned so far are reachable from \mathbf{x} through a path which lies completely within $ClnDom(\mathbf{C})$, where \mathbf{C} is the class of the object at \mathbf{x} .

The latter follows from the following two facts

- All receivers of one of the recursive calls of the `clone(...)` methods are reachable from $\phi(\mathbf{x})$ through a path that lies within $ClnDom(\mathbf{C})$. This fact can be shown by induction on the recursive calls as follows:
 - The base case is easy.
 - For the inductive step, consider an object ι'' which executes a method `clone(...)` as part of the n -th recursive call, and which calls a method `clone(...)` on an object ι''' .

Then, by construction of the `clone(...)` method, we know that the former call has the shape `clone(true, $\mathbf{b}_2, \dots, \mathbf{b}_q$)` and $q \geq 2$, that the latter call has the shape `clone($\mathbf{b}'_1, \dots, \mathbf{b}'_r$)` and $r \geq 2$, and that $\mathbf{b}'_1 = \mathbf{true}$ such that $H(\iota'', f) = \iota'''$, and where $fType(\mathbf{C}' < \mathbf{true}, \mathbf{b}_2, \dots, \mathbf{b}_q >, f) = \mathbf{D} < \mathbf{b}''_1, \dots, \mathbf{b}''_r >$ for some $\mathbf{b}''_1, \dots, \mathbf{b}''_r$, s.t. $\mathbf{b}''_1, \dots, \mathbf{b}''_r[\mathbf{true}/\mathbf{this}] = \mathbf{b}'_1, \dots, \mathbf{b}'_r$, and where \mathbf{C}' is the class of the object ι'' .

From the inductive hypothesis, we obtain that there exists a path from $\phi(x)$ to ι'' , which lies within $ClnDom(\mathbf{C})$, i.e., that $\mathbf{C} \vdash p : \mathbf{C}' < \mathbf{p}', \dots >$ for some path p' . Furthermore, the fact that $\mathbf{b}'_1 = \mathbf{true} = \mathbf{b}''_1[\mathbf{true}/\mathbf{this}]$ gives that $fType(\mathbf{C}' < \mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_q >, f) = \mathbf{D} < \mathbf{ca}_1, \dots, \mathbf{ca}_r >$ such that $\mathbf{ca}_1 = \mathbf{this}$, or that there exists an m such that $\mathbf{ca}_1 = \mathbf{c}_m$ and $\mathbf{b}_m = \mathbf{true}$. In the first case, by definition of the judgment, we obtain that $\mathbf{C} \vdash p.f : \mathbf{D} < \mathbf{p}', \dots >$, while in the second case we obtain $\mathbf{C} \vdash p.f : \mathbf{D} < \mathbf{p}, \dots >$. In either case, the path $p.f$ lies within $ClnDom(\mathbf{C})$.

- Only the objects that received one of the recursive calls of `clone(...)` may be have been cloned.

5.3 Completeness

Note, that Lemma 2 only guarantees that all cloned objects come from $ToCopy(\iota)_H$, and does not guarantee that *all* objects from $ToCopy(\iota)_H$ have been cloned. However, we can obtain this stronger property, if the system guarantees consistent views of the heap.

We first define consistent views of the heap to say that different paths leading to the same object must have the same path type.

Definition 1. We say that a type system gives consistent views of the heap, if for all classes C , paths p_1 and p_2 , and for all heaps H and stack frames ϕ which may arise through execution of well-typed expressions, if p_1 and p_2 are aliases in H and ϕ , and $\phi(\mathbf{this})$ is an object of class C , then

$$\Gamma \vdash p_1 : T_1 \text{ and } \Gamma \vdash p_2 : T_2 \text{ imply } T_1 = T_2$$

For example, if we had a system with consistent views of the heap, and if in our Fig. 1 the object of class `StudentList` contained a field called `last` which pointed to the same object as `this.head.next.next`, then, for typing environments Γ_1 and Γ_2 , such that $\Gamma_1(\mathbf{this}) = \mathbf{StudentList}$ and $\Gamma_2(\mathbf{this}) = \mathbf{Union}$, we would also have that $\Gamma_1 \vdash \mathbf{this.last} : \mathbf{Node} < \mathbf{this}, c >$ — this follows because of $\Gamma_1 \vdash \mathbf{this.head.next.next} : \mathbf{Node} < \mathbf{this}, c >$. Furthermore, we would also have that $\Gamma_2 \vdash \mathbf{this.students.last.student} : \mathbf{Student} < \mathbf{this} >$ — this follows because of $\Gamma_2 \vdash \mathbf{this.students.head.next.next.student} : \mathbf{Student} < \mathbf{this} >$.

The requirement of consistent views guarantees that the underlying type system restricts field assignment so that paths pointing to separate cloning domains cannot be aliases. Such requirements are satisfied by any sound ownership type systems, *e.g.*, as early as the system in [8].

Consistency in the views of the heap implies consistency in the path types:

Lemma 3. If a type system gives consistent views of the heap, then, for all classes C , paths p_1 and p_2 , and for all heaps H and stack frames ϕ which may arise through execution of well-typed expressions, if p_1 and p_2 are aliases in H and ϕ , and $\phi(\mathbf{this})$ is an object of class C , then

$$C \vdash p_1 : PT_1 \text{ and } C \vdash p_2 : PT_2 \text{ imply } PT_1 = PT_2$$

Continuing our example from above, consistent views of the heap would give us that $\mathbf{Union} \vdash \mathbf{this.students.last} : \mathbf{Node} < \mathbf{this.students}, \mathbf{this} >$, because $\mathbf{Union} \vdash \mathbf{this.students.head.next.next} : \mathbf{Node} < \mathbf{this.students}, \mathbf{this} >$.

Lemma 4. Assume that the type system gives consistent views of the heap. If $H, \phi, x.\mathbf{clone}() \rightsquigarrow H', \iota$, then the mapping $\alpha : \text{dom}(H') \rightarrow \text{dom}(H)$ described in lemma 2 has the property that $\alpha(\text{dom}(H') \setminus \text{dom}(H)) = \text{ToCopy}(\phi(x))_H$

Proof Sketch: We show by induction on n , that when considering the breadth first view of recursive applications of the `clone(...)` methods at depth n , then all objects which are reachable from x through a path of length at most n which lies within $\text{ClnDom}(C')$, have been cloned – with C' the class of the object at x .

In order to prove the above, we strengthen the assertion to say that for all n :

1. For all paths p and p' , where $p \in \text{ClnDom}(C')$, and p has length at most n such that $C' \vdash p : D < p_1, \text{cap}_2 \dots \text{cap}_m >$, if $H(\phi(x), p) = \iota$, then within the breadth first view of recursive applications of the `clone(...)` methods at depth n there is a call of `clone(true, b2, .. bm)` on ι , where $b_i = \mathbf{true}$ if cap_i has the form `this.f`, and `false` otherwise.

2. For any calls of `clone(true, b2, ..bm)` on an object ι within the breadth first view of recursive applications of the `clone(...)` methods at depth n , there is a path \mathbf{p} from $\phi(\mathbf{x})$ to ι , such that $\mathcal{C}' \vdash \mathbf{p} : \mathbf{D} \langle \mathbf{p}_1, \mathbf{cap}_2 \dots \mathbf{cap}_m \rangle$, and $\mathbf{b}_i = \mathbf{true}$ if \mathbf{cap}_i has the form `this. \bar{f}` , and `false` otherwise.

We now sketch the proof of the proposition from above:

- The base case is trivial.
- For the inductive step, we look at the two assertions separately:
 1. Consider any object ι' which is reachable from $\phi(\mathbf{x})$ through a path $\mathbf{p.f}$ of length $n + 1$ such that $\mathbf{p.f} \in \mathit{ClmDom}(\mathcal{C}')$, and $\mathcal{C}' \vdash \mathbf{p.f} : \mathbf{D} \langle \mathbf{p}', \mathbf{cap}_2 \dots \mathbf{cap}_m \rangle$ for some further path \mathbf{p}' . Therefore, there exists a further object ι'' , reachable from $\phi(\mathbf{x})$ through \mathbf{p} , such that $\mathcal{C}' \vdash \mathbf{p} : \mathbf{E} \langle \mathbf{cap}'_1, \mathbf{cap}'_2 \dots \mathbf{cap}'_r \rangle$, and that $H(\iota'', f) = \iota'$, and \mathbf{cap}'_1 has the form of a path.
By application of the inductive hypothesis part 1., we know that when considering the breadth first view of the recursive applications of the `clone(...)` methods at depth n , the object ι'' will have received the recursive method call `clone(b1, b2, ...br)` such that \mathbf{b}_i is `true`, if \mathbf{cap}'_i has itself the form of a path, and otherwise \mathbf{b}_i is `false`.
We now inspect the first lines of the code of the `clone(...)` method. At the point of execution of the method call `clone(b1, b2, ...br)` on ι' it is possible that the object ι' already had been entered in the map table \mathbf{m} . By inspection of the `clone(...)` methods, this implies that ι' had already received the method call `clone(b'1, b'2, ...b'r)` for some $\mathbf{b}'_1, \mathbf{b}'_2, \dots \mathbf{b}'_r$. By application of the inductive hypothesis part 2, we know that there exists a path \mathbf{p}'' from $\phi(\mathbf{x})$ to ι'' , such that $\mathcal{C}' \vdash \mathbf{p}'' : \mathbf{E} \langle \mathbf{p}''_1, \mathbf{cap}''_2 \dots \mathbf{cap}''_m \rangle$, and $\mathbf{b}'_i = \mathbf{true}$ if \mathbf{cap}''_i has the form `this. \bar{f}` , and `false` otherwise.
By application of the consistent view property and lemma 3, we obtain that $\mathbf{b}'_i = \mathbf{b}_i$. We now inspect the remaining lines of the code of the `clone(...)` method. These will give that the object ι' will receive the method call `clone(true, b''2, ...b''m)` where \mathbf{b}'_i is `true` if \mathbf{cap}_i has itself the form of a path, and otherwise \mathbf{b}'_i is `false`.
 2. This assertion is proven by induction on the depth of method calls of `clone(...)`, and analysis of the body of the method.

6 Related Work and Conclusions

In this paper we have sketched a way of *specifying* cloning policies by annotating the fields of objects, and then *deriving* code for cloning which satisfies the cloning policy. Our lemma 2 guarantees that the derived cloning code produces a faithful copy of the cloned object, and that cloned objects are within the domain expressed by the types, while lemma 4 guarantees that if the types give consistent views of the heap, then *all* objects within the domain expressed by the types will be cloned. We have not yet dealt with subclasses, but we do not expect this to be challenging, because the cloning method will be bound dynamically to the

class of the receiver. We plan to develop a detailed formal system and its proofs in further work.

Our work was inspired by Jensen et al.’s secure cloning policies [14]. Unlike the secure cloning policies, our work *generates* cloning methods, rather than simply checking them, and is based on ownership types, rather than a series of ad-hoc annotations. Compared with most ownership systems [5, 10, 9] our system is *descriptive* [3] in that it captures object relationships, rather than *prescriptively* ensuring invariants over the heap. We do not require or enforce heap properties (*e.g.*, clone/owners as dominators or owners as modifiers), nor even that the paths support a consistent view of the heap. In fact, owners as dominators could be too strong a requirement for cloning purposes, as it would unnecessarily forbid many programs, *e.g.* cloning a list with iterators. A consistent view of the heap may be enforced by the type system in the usual way [5]. We want to investigate in how far such a consistent view is the usual idiom, or whether there are situations where consistency is too restrictive from the practical perspective [15].

The relationship between ownership and cloning was first identified in a *dynamic* ownership setting, where every object knows its owner at runtime [17]. These ownership-based clones are called “sheep clones”, as they are intermediate between shallow (single object) clones and deep (fully recursive) clones. Cloning, and programming models and languages based on cloning, have of course been studied more generally [4, 18, 20, 23]. Ownership type techniques have also been applied to other problems closely related to cloning, in particular object initialisation [11, 22], as cloning creates and initialises a new object — and ownership transfer [7], as transfer can be modelled by cloning an object then deleting the original. Cloning objects is also similar to comparing objects [12, 13] and generating hash codes [21].

As well as completing the formal model, in further work we plan to investigate whether it is possible to derive more efficient versions of the `clone` function, so as to avoid the use of the `Map` object when possible. We also want to support generics, and existential (unknown) clone parameters.

Acknowledgments

We thank the FTfJP’11 reviewers for their constructive criticism on an earlier version of this paper, the FoVEOOS’11 reviewers for their feedback, and Paley Li for general discussions.

References

1. C. Andreae, Y. Coady, C. Gibbs, J. Noble, J. Vitek, and T. Zhao. Scoped Types and Aspects for Real-Time Java. In European Conference on Object-Oriented Programming (ECOOP), 2006.
2. N. Cameron. *Existential Types for Variance — Java Wildcards and Ownership Types*. PhD thesis, Imperial College London, 2009.

3. N. Cameron, S. Drossopoulou, J. Noble, and M. Smith. Multiple Ownership. In *Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 441–460. ACM, 2007.
4. K. Campbell, J. McWhir, W. Ritchie, and I. Wilmot. Sheep cloned by nuclear transfer from a cultured cell line. *Nature*, 380:64–66, Mar. 1996.
5. D. Clarke. *Object Ownership and Containment*. PhD thesis, University of New South Wales, 2001.
6. D. Clarke and S. Drossopoulou. Ownership, Encapsulation and the Disjointness of Types and Effects. In *OOPSLA*, pages 292–310, Seattle, Washington, USA, November 2002. ACM Press.
7. D. Clarke and T. Wrigstad. External uniqueness is unique enough. In *ECOOP*, pages 176–200, 2003.
8. D. G. Clarke, J. M. Potter, and J. Noble. Ownership types for flexible alias protection. In *OOPSLA*, volume 33(10), pages 48–64. ACM Press, 1998.
9. W. Dietl, S. Drossopoulou, and P. Müller. Generic Universe Types. In *ECOOP*, volume 4609 of *LNCS*, pages 28–53. Springer, 2007.
10. W. Dietl and P. Müller. Universes: Lightweight ownership for JML. *JOT*, 4(8):5–32, October 2005.
11. M. Fähndrich and S. Xia. Establishing object invariants with delayed types. In *OOPSLA*, pages 337–350. ACM Press, 2007.
12. P. Grogono and P. Chalin. Copying, sharing, and aliasing. In *Proceedings of the Colloquium on Object Orientation in Data bases and Software Engineering (COODBSE'94)*, Montreal, Quebec, May 1994.
13. P. Grogono and M. Sakkinen. Copying and comparing: Problems and solutions. In *ECOOP*, pages 226–250, 2000.
14. T. Jensen, F. Kirchner, and D. Pichardie. Secure the Clones – Static Enforcement of Policies for Secure Object Copying. In *ESOP*, 2011.
15. P. Li. Sheep cloning for ownership. In *OOPSLA Doctoral Consortium*, 2011. To Appear.
16. P. Müller. *Modular Specification and Verification of Object-Oriented Programs*, volume 2262 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.
17. J. Noble, D. Clarke, and J. Potter. Object ownership for dynamic alias protection. In *TOOLS Pacific*, 1999.
18. J. Noble and B. Foote. Attack of the clones. In *Proceedings of the 2002 Australasian Conference on Pattern Languages of Programs*, volume 13 of *CRPIT*, pages 99–144, 2002.
19. J. Noble, J. Potter, and J. Vitek. Flexible alias protection. In *ECOOP*, Brussels, Belgium, July 1998.
20. J. Noble, A. Taivalaari, and I. Moore, editors. *Prototype-Based Programming: Concepts, Languages, Applications*. Springer-Verlag, 1997.
21. D. Rayside, Z. Benjamin, R. Singh, J. P. Near, A. Milicevic, and D. Jackson. Equality and hashing for (almost) free: Generating implementations from abstraction functions. In *ICSE*, pages 342–352, 2009.
22. A. J. Summers and P. Müller. Freedom before commitment : A lightweight type system for object initialisation. In *OOPSLA*, 2011. To Appear.
23. D. Ungar and R. B. Smith. SELF: the Power of Simplicity. *LISP and Symbolic Computation*, 4(3), June 1991.
24. J. Vitek and J. Bokowski. Confined types for java. *Software Partice and Experience*, 2001.