GC-Safe Interprocedural Unboxing
Extended Version

Leaf Petersen and Neal Glew

Intel Labs, Santa Clara CA
leaf.petersen@intel.com
neal.glew@intel.com

Abstract. Modern approaches to garbage collection (*GC*) require information about which variables and fields contain GC-managed pointers. Interprocedural flow analysis can be used to eliminate otherwise unnecessary heap allocated objects (*unboxing*), but must maintain the necessary GC information. We define a core language which models compiler correctness with respect to the GC, and develop a correctness specification for interprocedural unboxing optimizations. We prove that any optimization which satisfies our specification will preserve GC safety properties and program semantics, and give a practical unboxing algorithm satisfying this specification.

1 Introduction

Precise garbage collection (*GC*) for managed languages is usually implemented by requiring the compiler to keep track of meta-data indicating which variables and fields contain GC-managed references and which should be ignored by the garbage collector. We refer to this information as the *traceability* of a field or variable: a field or variable is traceable if it should be treated as a pointer into the heap by the garbage collector.

In order to maintain this information in the presence of polymorphic functions (including subtype polymorphism), many languages and compilers use a *uniform object representation* in which every source level object is represented at least initially by a heap allocated object. All interprocedural use of native (non-heap) data therefore occurs only through fields of objects. This is commonly referred to as *boxing*, objects represented in this way are referred to as *boxed*, and projecting out of the uniform representation is referred to as *unboxing*. Boxing imposes substantial performance penalties for many reasons: the additional overhead of the allocation and projection is substantial, arrays of boxed objects exhibit poor locality, and the additional memory pressure can cause bottlenecks in the hardware. In this paper, we show how to use the results of interprocedural flow analyses (reaching definitions) to implement an interprocedural unboxing optimization while preserving the meta-data necessary for precise garbage collection.

In the following sections, we define a high-level core language that captures the essential aspects of GC meta-data and GC safety. We then give a high-level

specification of what a reasonable flow analysis on this language must compute, and define a notion of a general unboxing optimization for this language. We give a specification for when such an optimization is acceptable for a given flow analysis. We show that the optimization induced by an acceptable unboxing preserves the semantics of the original program, including GC safety. Finally, we construct an algorithm closely based on one in use in our compiler and prove that it produces an unboxing that satisfies our correctness specification, and hence that it preserves the semantics of the program (including GC safety).

All the lemmas and theorems in the paper have been proven. Proofs appear in Appendix A.

2 GC safety

Consider the following program (using informal notation), where `box` denotes a boxing operation that wraps its argument in a heap-allocated structure, and `unbox` denotes its elimination form that projects out the boxed item from the box:

$$\texttt{let}\ f\ =\ \lambda x.(\texttt{box}\, x)\ \texttt{in}\, \texttt{unbox}(\texttt{unbox}(f\ (\texttt{box}\, 3)))$$

The only definition reaching the variable x is the boxed machine integer 3. Information from an interprocedural analysis can be used to rewrite this program to eliminate the boxing as follows:

$$\texttt{let}\ f\ =\ \lambda x.x\ \texttt{in}\, f\ 3$$

In the second version of this program, the traceability of the values reaching x has changed: whereas in the original program all values reaching x are represented as heap allocated pointers, in the second program all values reaching x are represented as machine integers. From the standpoint of a garbage collector, a garbage collection occuring while x is live must treat x as a root in the first program, and must ignore x in the second program. There are numerous approaches to communicating this information to the garbage collector. For example, some implementations choose to dynamically tag values in such a way as to allow the garbage collector to distinguish pointers from non-pointers by inspection. Such an implementation might steal a low bit from the machine integer representation to allow the machine integer 3 to be distinguished from a heap pointer.

Another very commonly used approach (particularly in more recent systems) is to require the compiler to statically annotate the program with garbage collection meta-data such that at any garbage collection point the garbage collector can reconstruct exactly which live variables are roots. Typically, this takes the form of annotations on variables and temporaries indicating which contain heap-pointers (the roots) and which do not (the non-roots), along with information at every allocation site indicating which fields of the allocated object contain traceable data. It is this approach that we target in this paper.

The requirement that the compiler be able to annotate program variables with a single static traceability constrains the compiler's ability to rewrite programs in that it must do so in a way that preserves the correctness of the GC

meta-data of the program. Consider an extension of the previous example.

$$\texttt{let}\ f\ =\ \lambda x.x\ \texttt{in}\,\texttt{unbox}((f\ f)\ (\texttt{box}\,3))$$

Assuming that functions are represented as heap-allocated objects then each variable in this program can be assigned a traceability, since all objects passed to f are heap references. However, an attempt to unbox this program as with the previous example results in f being applied to both heap references (f) and non-heap references (3).

$$\texttt{let}\ f\ =\ \lambda x.x\ \texttt{in}(f\ f)\ 3$$

As this example shows, the concerns of maintaining garbage-collector meta-data constrain optimization[1] in ways not apparent in a GC-ignorant semantic model.

2.1 A core language for GC safety

In order to give a precise account of interprocedural unboxing, we begin by defining a core language capturing the essential features of GC safety. The motivation for the idiosyncracies of this language lies in the requirements of the underlying model of garbage collection. We assume that pointers cannot be intrinsically distinguished from non-pointers, and hence must be tracked by the compiler. In our implementation, the compiler intermediate language under consideration is substantially more low-level: a control-flow graph based, static single assignment intermediate representation. We believe however that all of the key issues are captured faithfully in this higher-level representation.

Traceabilities t	$::=\ \texttt{b} \mid \texttt{r}$	Terms m	$::=\ x \mid \lambda x^i{:}t.e \mid e_1\ e_2 \mid \texttt{box}_t\ e \mid$
Term variables	x, y, z		$\texttt{unbox}\ e \mid \rho(e)$
Constants	c	Values v	$::=\ c \mid \langle \rho, \lambda x^i{:}t.e \rangle \mid \langle v^i{:}t \rangle$
Labels i	$::=\ 0, 1, \ldots$	Environments ρ	$::=\ x_1^{i_1}{:}t_1 = {v_1}^{j_1}, \ldots, x_n^{i_n}{:}t_n = {v_n}^{j_n}$
Labeled Terms e	$::=\ m^i \mid v^i$	States M	$::=\ (\rho, e)$

Fig. 1. Syntax

Figure 1 defines the syntax of our core language. The essence of the language is largely that of the standard untyped lambda calculus with an explicit environment semantics, extended with a form of degenerate type information we call *traceabilities*. Traceabilities describe the GC status of variables: the traceability `b` (for bits) indicates something that should be ignored by the garbage collector, while the traceability `r` (for reference) indicates a GC-managed pointer.

[1] It is worth noting that a serious compiler might be expected to duplicate the body of f in this simple example thereby eliminating this constraint and allowing the unboxing optimization to be more effective.

The traceability `b` is inhabited by an unspecified set of constants c while the traceability `r` is inhabited by functions (anticipating their implementation by heap-allocated closures) and by boxed objects. Anticipating the needs of the flow analysis, we label each term, value, and variable binding site with an integer label. We do not assume that labels or variables are unique within a program.

Expressions e consist of labeled terms m^i and labeled values v^i. The terms m consist of variables, functions, applications, box introductions, box eliminations, and frames. Variable binding sites are decorated with traceability information ($\lambda x^i{:}t.e$), as are box introductions ($\texttt{box}_t\, e$). We represent heap allocation in the language via the $\texttt{box}_t\, e$ term, which corresponds to allocating a heap cell containing the value for e. The traceability t gives the meta-data with which the heap-cell will be tagged, allowing the garbage collector to trace the cell. Objects can be projected out of an allocated object by the $\texttt{unbox}\, e$ operation. Frames $\rho(e)$ are discussed below.

Values consist of either constants, closures, or heap-allocated boxes. We distinguish between the introduction form ($\texttt{box}_t\, e$) and the value form ($\langle v^i{:}t \rangle$) for allocated objects. The introduction form corresponds to the allocation instruction, whereas the value form corresponds to the allocated heap value. This distinction is key for the formulation of GC safety and the dynamic semantics. For the purposes of the dynamic semantics we also distinguish between functions ($\lambda x^i{:}t.e$) and the heap allocated closures that represent them at runtime ($\langle \rho, \lambda x^i{:}t.e \rangle$).

For notational convenience, we will sometimes use the notation $v_{\texttt{b}}$ to indicate that a value v is a non-heap-allocated value (i.e. a constant c), and $v_{\texttt{r}}$ to indicate that a value v is a heap-allocated value (i.e. either a lambda value or a boxed value). If t is a traceability meta-variable, then we use v_t to indicate that v is a value of the same traceability as t. In examples, we use a derived `let` expression, taking it to be syntactic sugar for application in the usual manner. Environments ρ map variables to values. The term $\rho(e)$ executes e in the environment ρ rather than the outer environment – all of the free variables of e are provided by ρ. The nested set of these environments at any point can be thought of as the activation stack frames of the executing program. The traceability annotations on variables in the environments play the role of stack frame GC meta-data, indicating which slots of the frame are roots (traceability `r`). The environments buried in closures ($\langle \rho, \lambda x^i{:}t.e \rangle$) similarly provide the traceabilities of values reachable from the closure, and hence provide the GC meta-data for tracing through closures. While we do not make the process of garbage collection explicit, it should be clear how to extract the appropriate set of GC roots from the environment and any active frames.

This core language contains the appropriate information to formalize a notion of GC safety consisting of two complementing pieces. First we define a dynamic semantics in which reductions that might lead to undefined garbage-collector behavior are explicitly undefined. Programs that takes steps in this semantics do not introduce ill-formed heap objects. Secondly, we define a notion of a traceable program: one in which all heap values have valid GC meta-data. Reduction steps in the semantics can then be shown to maintain the traceability property. The

GC correctness criteria for a compiler optimization then is that the optimization map traceable programs to traceable programs, and that it not introduce new undefined behavior.

2.2 Operational semantics

We choose to use an explicit environment semantics rather than a standard substitution semantics since this makes the GC meta-data for stack frames and closures explicit in the semantics. Thus a machine state (ρ, e) supplies an environment ρ for e that provides the values of the free variables of e during execution. Environments contain traceability annotations on each of the variables mapped by the environment.

$$\frac{x^i{:}t = v^j \in \rho}{(\rho, x^k) \longmapsto (\rho, v^j)} \qquad \frac{}{(\rho, (\lambda x^i{:}t.e)^j) \longmapsto (\rho, \langle \rho, \lambda x^i{:}t.e \rangle^j)}$$

$$\frac{t = t'}{(\rho, (\mathtt{box}_t\ {v_{t'}}^i)^j) \longmapsto (\rho, \langle {v_{t'}}^i{:}t \rangle^j)} \qquad \frac{(\rho, e_1) \longmapsto (\rho, e_1')}{(\rho, (e_1\ e_2)^i) \longmapsto (\rho, (e_1'\ e_2)^i)}$$

$$\frac{(\rho, e_2) \longmapsto (\rho, e_2')}{(\rho, (v^i\ e_2)^j) \longmapsto (\rho, (v^i\ e_2')^j)} \qquad \frac{t = t'}{(\rho, (\langle \rho', \lambda x^i{:}t.e \rangle^j\ {v_{t'}}^k)^l) \longmapsto (\rho, (\rho', x^i{:}t = v_{t'}{}^k)(e)^l)}$$

$$\frac{(\rho, e) \longmapsto (\rho, e')}{(\rho, (\mathtt{box}_t\ e)^i) \longmapsto (\rho, (\mathtt{box}_t\ e')^i)}$$

$$\frac{(\rho, e) \longmapsto (\rho, e')}{(\rho, (\mathtt{unbox}\ e)^i) \longmapsto (\rho, (\mathtt{unbox}\ e')^i)} \qquad \frac{}{(\rho, (\mathtt{unbox}\ \langle v^i{:}t \rangle^j)^k) \longmapsto (\rho, v^i)}$$

$$\frac{(\rho', e) \longmapsto (\rho', e')}{(\rho, \rho'(e)^i) \longmapsto (\rho, \rho'(e')^i)} \qquad \frac{}{(\rho, \rho'(v^i)^j) \longmapsto (\rho, v^i)}$$

Fig. 2. Operational Semantics

Reduction in this language is for the most part fairly standard. We deviate somewhat in that we explicitly model the allocation of heap objects as a reduction step—hence there is an explicit reduction mapping a lambda term $\lambda x^i{:}t.e$ to an allocated closure $\langle \rho, \lambda x^i{:}t.e \rangle$, and similarly for boxed objects and values. More notably, beta-reduction is restricted to only permit construction of a stack frame when the meta-data attached to the parameter variable is appropriate for the actual argument value. This captures the requirement that stack frames have correct meta-data for the garbage collector. In actual practice, incorrect meta-data for stack frames leads to undefined behavior (since incorrect meta-data may

cause arbitrary memory corruption by the garbage collector)—similarly here in the meta-theory we leave the behavior of such programs undefined. In a similar fashion, we only define the reduction of the allocation operation to an allocated value ($\mathtt{box}_t\, v_{t'} \longmapsto \langle v_{t'}{:}t \rangle$) when the operation meta-data is appropriate for the value (i.e. $t = t'$).

It is important to note that this semantics does not model a dynamically checked language, in which there is an explicit check of the meta-data associated with these reductions. The point is simply that the semantics only specifies how programs behave when these conditions are met—in all other cases the behavior of the program is undefined.

2.3 Traceability

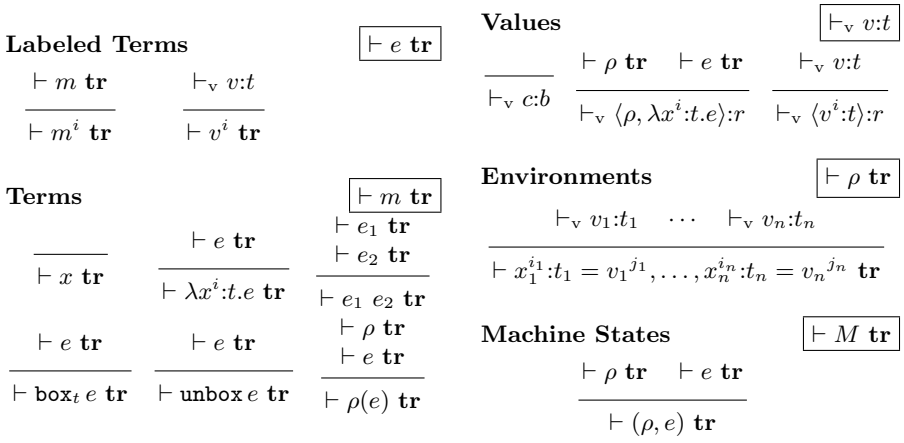

Fig. 3. Traceability

The operational semantics ensures that no reduction step introduces mis-tagged values. In order to make use of this, we define a judgment for checking that a program does not have a mis-tagged value in the first place. Implicitly this judgement defines what a well-formed heap and activation stack looks like; however, since our heap and stack are implicit in our machine states, it takes the form of a judgement on terms, values, environments, and machine states.

The value judgement $\vdash_{\mathrm{v}} v{:}t$ asserts that a value v is well-formed, and has traceability t. In this simple language, this corresponds to having the meta-data on the environment of each lambda value be consistent and the meta-data on each boxed value be consistent with the traceability of the object nested in the box. An environment is consistent, $\vdash \rho$ **tr**, when the annotation on each variable agrees with the traceability of the value it is bound to. Since we cannot check

the consistency of general terms with the first-order information available, the term judgement $\vdash e$ **tr** and machine state judgement $\vdash M$ **tr** simply check that all values and environments (and hence stack frames) contained in the term or machine state are well-formed.

The key result for traceability is that it is preserved under reduction. That is, if a traceable term takes a well-defined reduction step, then the resulting term will be traceable.

Lemma 1 (Preservation of traceability). *If* $\vdash M$ **tr** *and* $M \longmapsto M'$ *then* $\vdash M'$ **tr**.

There is of course no corresponding progress property for our notion of traceability, since programs can go wrong. Compiler optimizations are simply responsible for ensuring that they do not introduce new ways to go wrong.

3 Flow analysis

Our original motivation for this work was to apply interprocedural analysis to the problem of eliminating unnecessary boxing in programs. There is a vast body of literature on interprocedural analysis and optimization, and it is generally fairly straightforward to use these approaches to obtain information about what terms flow to what use sites. This paper is not intended to provide any contribution to this body of work, which we will broadly refer to as *flow analysis*. Instead, we focus on how to use the results of such a generic analysis to implement an unboxing optimization that preserves GC safety.

In order to do this, we must provide some framework for describing what information a flow analysis must provide. For the purposes of our unboxing optimization, we are interested in finding (inter-procedurally) for every $(\texttt{unbox}\, v^j)^i$ operation the set of $(\texttt{box}_t\, e)^k$ terms that could possibly reach v. Under appropriate conditions, we can then eliminate both the box introductions and the box elimination, thereby improving the program. The core language defined in Section 2 provides labels serving as proxies for the terms and variables on which they occur – the question above can therefore be re-stated as finding the set of labels k that reach the position labeled with j.

More generally, following previous work we begin by defining an abstract notion of analysis. We say that an analysis is a pair (C, ϱ). Binding environments ϱ simply serve to map variables to the label of their binding sites. The mappings are, as usual, global for the program. Consequently, a given environment may not apply to alpha-variants of a term. We do not require that labels be unique within a program—as usual however, analyses will be more precise if this is the case. Variables are also not required to be unique (since reduction may duplicate terms and hence binding sites). However, duplicate variable bindings in a program must be labeled consistently according to ϱ or else no analysis of the program can be acceptable according to our definition. This can always be avoided by alpha-varying or relabeling appropriately.

A cache C is a mapping from labels to sets of shapes. Shapes are given by the grammar:

$$\text{Shapes: } s ::= c^i \mid (i{:}t \to j)^k \mid (\texttt{box}_t\ i)^j$$

The idea behind shapes is that each shape provides a proxy for a set of terms that might flow to a given location, describing both the shape of the values that might flow there and the labels of the sub-components of those values. For example, for an analysis (C, ϱ), $c^i \in \mathrm{C}(k)$ indicates that (according to the analysis) the constant c, labeled with i, might flow to a location labeled with k. Similarly, if $(i{:}t \to j)^k \in \mathrm{C}(l)$, then the analysis specifies that among the values flowing to locations labeled with l might be lambdas labeled with k, whose parameter variable is labeled with i and annotated with t and whose bodies are labeled with j. If $(\texttt{box}_t\ k)^i \in \mathrm{C}(l)$ then among the values that might flow to l (according to the analysis) are boxed values labeled with i, with meta-data t and whose bodies are labeled by some j such that $\mathrm{C}(j) \subseteq \mathrm{C}(k)$.

It is important to note that the shapes in the cache may not correspond exactly to the terms in the program, since reduction may change program terms (e.g. by instantiating variables with values). However, reduction does not change the outer shape and labeling of values—it is this reduction invariant information that is captured by shapes.

Clearly, not every choice of analysis pairs is meaningful for program optimization. While in general it is reasonable (indeed, unavoidable) for an analysis to overestimate the set of terms associated with a label, it is unacceptable for an analysis to underestimate the set of terms that flow to a label—most optimizations will produce incorrect results, since they are designed around the idea that the analysis is telling them everything that could possibly flow to them. In order to capture the notion of when an analysis pair gives a suitable approximation of the flow of values in a program we follow the general spirit of Nielson et al. [6], and define a notion of an *acceptable analysis*. That is, we give a declarative specification that gives sufficient conditions for specifying when a given analysis does not underestimate the set of terms flowing to a label, without committing to a particular analysis. We arrange the subsequent meta-theory such that our results apply to any analysis that is *acceptable*. In this way, we completely decouple our optimization from the particulars of how the analysis is computed.

Our acceptable-analysis relation is given in Figure 4 – the judgement $\mathrm{C}; \varrho \vdash (\rho, e)$ determines that an analysis pair (C, ϱ) is *acceptable* for a machine state (ρ, e), and similarly for the environment and expression forms of the judgement. We use the notation $\mathrm{lbl}(e)$ to denote the outermost label of e: that is, i where e is of the form m^i or v^i. The acceptability judgement generally indicates for each syntactic form what the flow of values is. For example, in the application rule, the judgment insists that for every lambda value that flows to the applicand position, the set of shapes associated with the parameter of that lambda is a super-set of the set of shapes associated with the argument of the application; and that the set of shapes associated with the result of the lambda is a sub-set of the set of shapes associated with the application itself.

$\boxed{\mathrm{C};\varrho \vdash e}$

$$\frac{\mathrm{C}(\varrho(x)) \subseteq \mathrm{C}(i)}{\mathrm{C};\varrho \vdash x^i} \qquad \frac{\varrho(x) = j \quad \mathrm{C};\varrho \vdash e \quad (j{:}t \to \mathrm{lbl}(e))^i \in \mathrm{C}(i)}{\mathrm{C};\varrho \vdash (\lambda x^j{:}t.e)^i}$$

$$\frac{\begin{array}{c}\mathrm{C};\varrho \vdash e_1 \quad \mathrm{C};\varrho \vdash e_2 \\ \forall (k{:}t \to l)^j \in \mathrm{C}(\mathrm{lbl}(e_1)) : \\ \mathrm{C}(\mathrm{lbl}(e_2)) \subseteq \mathrm{C}(k) \wedge \mathrm{C}(l) \subseteq \mathrm{C}(i)\end{array}}{\mathrm{C};\varrho \vdash (e_1\ e_2)^i} \qquad \frac{\mathrm{C};\varrho \vdash e \quad (\mathtt{box}_t\ j)^i \in \mathrm{C}(i) \quad \mathrm{C}(\mathrm{lbl}(e)) \subseteq \mathrm{C}(j)}{\mathrm{C};\varrho \vdash (\mathtt{box}_t\ e)^i}$$

$$\frac{\begin{array}{c}\mathrm{C};\varrho \vdash e \\ \forall (\mathtt{box}_t\ k)^j \in \mathrm{C}(\mathrm{lbl}(e)) : \mathrm{C}(k) \subseteq \mathrm{C}(i)\end{array}}{\mathrm{C};\varrho \vdash (\mathtt{unbox}\ e)^i} \qquad \frac{\mathrm{C};\varrho \vdash \rho \quad \mathrm{C};\varrho \vdash e \quad \mathrm{C}(\mathrm{lbl}(e)) \subseteq \mathrm{C}(i)}{\mathrm{C};\varrho \vdash \rho(e)^i}$$

$$\frac{c^i \in \mathrm{C}(i)}{\mathrm{C};\varrho \vdash c^i} \qquad \frac{\varrho(x) = j \quad \mathrm{C};\varrho \vdash \rho \quad \mathrm{C};\varrho \vdash e \quad (j{:}t \to \mathrm{lbl}(e))^i \in \mathrm{C}(i)}{\mathrm{C};\varrho \vdash \langle \rho, \lambda x^j{:}t.e \rangle^i}$$

$$\frac{\mathrm{C};\varrho \vdash v^j \quad (\mathtt{box}_t\ k)^i \in \mathrm{C}(i) \quad \mathrm{C}(j) \subseteq \mathrm{C}(k)}{\mathrm{C};\varrho \vdash \langle v^j{:}t \rangle^i}$$

$\boxed{\mathrm{C};\varrho \vdash \rho}$

$$\frac{\forall 1 \le k \le n : \varrho(x_k) = i_k \wedge \mathrm{C}(j_k) \subseteq \mathrm{C}(i_k) \wedge \mathrm{C};\varrho \vdash {v_k}^{j_k}}{\mathrm{C};\varrho \vdash x_1^{i_1}{:}t_1 = {v_i}^{j_1}, \ldots, x_n^{i_n}{:}t_n = {v_n}^{j_n}}$$

$\boxed{\mathrm{C};\varrho \vdash M}$

$$\frac{\mathrm{C};\varrho \vdash \rho \quad \mathrm{C};\varrho \vdash e}{\mathrm{C};\varrho \vdash (\rho, e)}$$

Fig. 4. Acceptable Analysis

Given this definition, we can show that the acceptability relation is preserved under reduction.

Lemma 2 (Many-step reduction preserves acceptability). *If* $\mathrm{C};\varrho \vdash M$ *and* $M \longmapsto^* M'$ *then* $\mathrm{C};\varrho \vdash M'$.

4 Unboxing

The goal of the unboxing optimization is to use the information provided by a flow analysis to replace a boxed object with the contents of the box. Doing

so may change the traceability, since the object in the box may not be a GC-managed reference. Moreover, the object in the box may itself be a candidate for unboxing; consequently, determining the traceability of boxed objects depends on exactly which objects are unboxed. Function parameters may be instantiated with objects from multiple different definition sites, some of which may be unboxed and some of which may not.

Consider again the first example from Section 1, written out with explicit GC information and labels:

$$\texttt{let}\ f^0{:}\texttt{r}\ =\ (\lambda x^1{:}\texttt{r}.(\texttt{box}_\texttt{r}\, x^2)^3)^4\ \texttt{in}\,(\texttt{unbox}\,(\texttt{unbox}\,(f^5\ (\texttt{box}_\texttt{b}\, 3^6)^7)^8)^9)^{10}$$

It is fairly easy to see that this program is unboxable. The binding site for x is only reached by the term labeled with 7 (the outer box introduction), and hence there should be no problems with changing its traceability annotation. Each box elimination is reached only by a single box introduction, and hence the box/unbox pairs in this program should be eliminable, yielding an optimized program:

$$\texttt{let}\ f^0{:}\texttt{r}\ =\ (\lambda x^1{:}\texttt{b}.x^2)^4\ \texttt{in}\,(f^5\ 3^6)^8$$

Notice that in order to rewrite the program, we have had to change the traceability annotation at the binding site for x, since we have eliminated the box introduction on its argument. This constraint is imposed on us by the need to keep the GC information consistent. If we choose (perhaps because of limitations on the precision of the analysis, or perhaps because of other constraints) to only eliminate the innermost box/unbox pair, then we must similarly adjust the traceability annotation on the remaining box introduction (labeled with 3).

$$\texttt{let}\ f^0{:}\texttt{r}\ =\ (\lambda x^1{:}\texttt{b}.(\texttt{box}_\texttt{b}\, x^2)^3)^4\ \texttt{in}\,(\texttt{unbox}\,(f^5\ 3^6)^8)^9$$

Not all programs can be consistently rewritten in this manner. If we consider again the second example from Section 1, we see an example of a program in which we must forgo optimization if we wish to preserve GC safety.

$$\texttt{let}\ f^0{:}\texttt{r}\ =\ (\lambda x^1{:}\texttt{r}.x^2)^3\ \texttt{in}\,(\texttt{unbox}\,((f^4\ f^5)^6\ (\texttt{box}_\texttt{b}\, 3^7)^8)^9)^{10}$$

It is easy to see that any acceptable analysis must include the function labeled with 3 and the boxed term labeled with 8 in the set of terms reaching the binding site for x, labeled with 1. We might naively attempt to eliminate the box/unbox pair as follows:

$$\texttt{let}\ f^0{:}\texttt{r}\ =\ (\lambda x^1{:}?.x^2)^3\ \texttt{in}\,((f^4\ f^5)^6\ 3^7)^9$$

Unfortunately, there is no consistent choice of traceability annotation for the binding site for x. If we choose `b` as the traceability annotation then after reduction we arrive at a state that has no defined reduction:

$$((\langle\epsilon, \lambda x^1{:}\texttt{b}.x^2\rangle^3\ \langle\epsilon, \lambda x^1{:}\texttt{b}.x^2\rangle^3)^6\ 3^7)^9$$

The first application leads to undefined behavior, since the traceability of the argument value does not match the traceability annotation on the parameter variable. If we had instead chosen `r` as the traceability annotation, then one further reduction would still lead us to undefined behavior.

$$\left(\langle \epsilon, \lambda x^1{:}\mathtt{r}.x^2 \rangle^3 \ 3^7\right)^9$$

The requirement to preserve GC information imposes two burdens on us then: we must provide some mechanism for assigning new GC meta-data when we optimize the program, and we must also ensure that we do not optimize the program in a way that does not admit a consistent assignment of such meta-data. In the rest of this section, we first develop a framework for specifying an unboxing assignment regardless of any correctness concerns, and then separately define a judgement specifying when such an assignment is a reasonable one.

4.1 The unboxing optimization

We can divide the problem of specifying an unboxing into two sub-parts: choosing the particular box/unbox pairs that are valid to eliminate and assigning new traceability annotations to terms that are affected. An unboxing then is a pair (T, Υ), where Υ is a set of labels, and T is a partial function from labels to traceabilities. The unboxed set Υ is the set of labels to be unboxed, and the traceability map T specifies new traceabilities for labels affected by the unboxing. The fact that T is a partial function is essential for several reasons. On a technical level, we do not require that labels be unique in a program. Consequently, it is possible that there is no consistent choice for a specific label. More importantly, requiring that T be a total function would put unsatisfiable requirements on the flow analysis. For example, a program that allocates a mis-tagged object after going into an infinite loop is GC safe according to our specification since the bad allocation is never reached. Requiring the analysis to find a consistent traceability map for such a program is equivalent to requiring it to solve the halting problem, since it must statically prove that the set of values dynamically reaching the mis-tagged allocation site is empty. By allowing T to be a partial function, we allow for necessary imprecision in the analysis. Also of importance is the need to allow for relative imprecision in the analysis. In order to achieve faster compile times, we may choose to use less precise analyses that potentially over-approximate the set of terms reaching a use point. Consequently, even if a consistent traceability assignment exists, we may not have sufficiently precise information to construct it.

An unboxing pair defines a total function mapping labeled terms to labeled terms, as shown in Figure 5. For notational convenience, we take $\mathrm{T}(i) = t$ as asserting that i is in the domain of T, and that its image is t. We also say that: $\mathrm{T}(i) \geq t$ if and only if $\mathrm{T}(i) = t$ or $\mathrm{T}(i)$ is undefined; and $\underline{\mathrm{T}}(i, t) = \mathrm{T}(i)$ if T defined at i, otherwise t.

An important observation about the unboxing optimization as we have defined it is unlike many previous interprocedural approaches (Section 6), it only

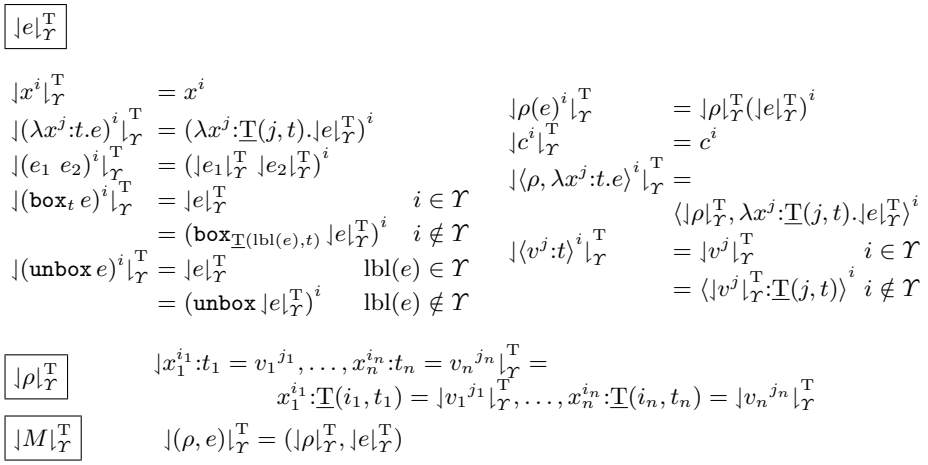

Fig. 5. Unboxing

improves programs and never introduces instructions or allocation. This is easy to see, since the unboxing function only removes boxes (which allocate and have an instruction cost), and unboxes (which have an instruction cost) and never introduces any new operations at all.

4.2 Acceptable unboxings

While any choice of (T, Υ) defines an unboxing, not every unboxing pair is reasonable in the sense that it defines a semantics preserving optimization. Just as we defined a notion of acceptable analysis in Section 3, we will define a judgement that captures sufficient conditions for ensuring correctness of an unboxing, without specifying a particular method of choosing such an unboxing. By using analyses of different precisions or choosing different optimization strategies we may end up with quite different choices of unboxings; however, so long as they satisfy our notion of acceptability we can be sure that they will preserve correctness.

Informally, we can eliminate a box introduction if certain criteria are met. Firstly, we must be able to eliminate all of the unbox operations that it reaches. Secondly, we must be able to find a consistent traceability assignment covering each intermediate variable or field that it reaches, given all of the rest of our unboxing choices. We can eliminate an unbox operation if we can eliminate all of the box operations that reach it. Finally, we must also impose coherence requirements on traceability assignments. For every variable whose binding-site label occurs in the domain of T, we require that its new traceability assignment agree with the traceability assignment of all of its reaching definitions. Similarly, for every box introduction (or value form) that is not itself unboxed, we require that the traceability assignment for its contents agree with the traceability assignment for every reaching definition in the flow analysis.

This informal description is made precise in Figure 6. We use the notation $i \overset{\mathrm{T},\Upsilon}{\simeq} j$ to indicate when an unboxing *agrees* at two labels i and j.

$i \overset{\mathrm{T}}{\simeq} j$ iff either $\mathrm{T}(i) = \mathrm{T}(j)$ (both defined) or $\mathrm{T}(i)$ and $\mathrm{T}(j)$ undefined
$i \overset{\Upsilon}{\simeq} j$ iff either $i, j \in \Upsilon$ or $i, j \notin \Upsilon$
$i \overset{\mathrm{T},\Upsilon}{\simeq} j$ iff $i \overset{\mathrm{T}}{\simeq} j$ and $i \overset{\Upsilon}{\simeq} j$

An unboxing pair (T, Υ) is acceptable relative to an analysis (C, ϱ) for a program M (judgement $\mathrm{C} \vdash M \downarrow\downarrow (\mathrm{T}, \Upsilon)$) if the unboxing is *cache consistent* (judgement $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$), and *consistent* (judgement $\mathrm{T}, \Upsilon \vdash M$).

$\boxed{\mathrm{T}, \Upsilon \vdash e}$

$$\frac{}{\mathrm{T}, \Upsilon \vdash x^i} \qquad \frac{\mathrm{T}, \Upsilon \vdash e \quad \mathrm{T}(j) \geq \mathtt{r}}{\mathrm{T}, \Upsilon \vdash (\lambda x^i{:}t.e)^j} \qquad \frac{\mathrm{T}, \Upsilon \vdash e_1 \quad \mathrm{T}, \Upsilon \vdash e_2}{\mathrm{T}, \Upsilon \vdash (e_1\ e_2)^i} \qquad \frac{i \in \Upsilon \quad \mathrm{T}(i) = \mathrm{T}(\mathrm{lbl}(e)) \quad \mathrm{T}, \Upsilon \vdash e}{\mathrm{T}, \Upsilon \vdash (\mathtt{box}_t\ e)^i} \qquad \frac{i \notin \Upsilon \quad \mathrm{T}(i) \geq \mathtt{r} \quad \mathrm{T}, \Upsilon \vdash e}{\mathrm{T}, \Upsilon \vdash (\mathtt{box}_t\ e)^i}$$

$$\frac{\mathrm{T}, \Upsilon \vdash e}{\mathrm{T}, \Upsilon \vdash (\mathtt{unbox}\ e)^i} \qquad \frac{\mathrm{T}, \Upsilon \vdash \rho \quad \mathrm{T}, \Upsilon \vdash e}{\mathrm{T}, \Upsilon \vdash \rho(e)^i} \qquad \frac{\mathrm{T}(i) \geq \mathtt{b}}{\mathrm{T}, \Upsilon \vdash c^i} \qquad \frac{\mathrm{T}, \Upsilon \vdash \rho \quad \mathrm{T}, \Upsilon \vdash e \quad \mathrm{T}(j) \geq \mathtt{r}}{\mathrm{T}, \Upsilon \vdash \langle \rho, \lambda x^i{:}t.e \rangle^j}$$

$$\frac{j \in \Upsilon \quad \mathrm{T}(j) = \mathrm{T}(i) \quad \mathrm{T}, \Upsilon \vdash v^i}{\mathrm{T}, \Upsilon \vdash \langle v^i{:}t \rangle^j} \qquad \frac{j \notin \Upsilon \quad \mathrm{T}(j) \geq r \quad \mathrm{T}, \Upsilon \vdash v^i}{\mathrm{T}, \Upsilon \vdash \langle v^i{:}t \rangle^j}$$

$\boxed{\mathrm{T}, \Upsilon \vdash \rho}$ $$\frac{\forall 1 \leq k \leq n : \underline{\mathrm{T}}(i_k, t_k) = \underline{\mathrm{T}}(j_k, t_k) \wedge \mathrm{T}, \Upsilon \vdash {v_k}^{j_k}}{\mathrm{T}, \Upsilon \vdash x_1^{i_1}{:}t_1 = {v_1}^{j_1}, \ldots, x_n^{i_n}{:}t_n = {v_n}^{j_n}}$$

$\boxed{\mathrm{T}, \Upsilon \vdash M}$ $$\frac{\mathrm{T}, \Upsilon \vdash \rho \quad \mathrm{T}, \Upsilon \vdash e}{\mathrm{T}, \Upsilon \vdash (\rho, e)}$$

$\boxed{\mathrm{C} \vdash (\mathrm{T}, \Upsilon)}$ $$\frac{\forall i : s \in \mathrm{C}(i) \implies i \overset{\mathrm{T},\Upsilon}{\simeq} \mathrm{lbl}(s)}{\mathrm{C} \vdash (\mathrm{T}, \Upsilon)}$$

$\boxed{\mathrm{C} \vdash M \downarrow\downarrow (\mathrm{T}, \Upsilon)}$ $$\frac{\mathrm{C} \vdash (\mathrm{T}, \Upsilon) \quad \mathrm{T}, \Upsilon \vdash M}{\mathrm{C} \vdash M \downarrow\downarrow (\mathrm{T}, \Upsilon)}$$

Fig. 6. Consistent and acceptable unboxing

Cache consistency $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$ encapsulates the requirement that an unbox can only be eliminated if all of the reaching definitions of its target are unboxed. It requires agreement between the label of the target of the unbox and the labels of everything in the cache of the target. The results from Section 3 ensure that under any evaluation, any term reaching the unbox is in the cache of the original

target label, and hence that the unboxing approximation takes into account a sufficient set of terms[2].

Cache consistency does not put any constraints on the actual choice of traceabilities in T. The consistency judgement $(\mathrm{T}, \Upsilon \vdash M)$ ensures that the traceability map T encodes choices that are compatible with the actual labeled terms in M, given a particular choice of terms to unbox Υ.

For environments, the consistency judgement insists that the traceability map assign consistent traceabilities to values and the variables to which they are bound. In this way we can ensure that the result of unboxing an environment still provides good traceability information for the garbage collector.

The term consistency judgement for the most part only requires that the traceability map be consistent with the labeled values. Variable uses incur no constraints, and neither do applications nor unbox operations (beyond requiring the consistency of their sub-terms). For constants c^i, we require that the traceability assignment T, if defined at i, maps i to `b`. That is, we require that the traceability assignment for i is consistent with the actual term inhabiting the label. Functions have a similar requirement: the traceability assignment for their label, if present, must be `r` since functions are represented by heap allocated closures. In the value form the closed over environment must be consistent as well.

The only particularly interesting rules are those covering the boxed introduction form and isomorphically the boxed value form. There are two cases: one for when the boxed value is selected for unboxing (that is, its label is in Υ), and one for when it has not been selected for unboxing.

If the term is not to be unboxed ($j \notin \Upsilon$), then the consistency rule requires that its traceability assignment (if any) be `r`. In the case that the term is to be unboxed ($j \in \Upsilon$) this is not required since the unboxed value may not in fact end up having traceability `r`. Instead, we require that the traceability map have an assignment both for the label of the box (j), and for the label of the contents of the box (i), and that it assign the same traceability to both. This requirement may be somewhat unexpected at first. The intuition behind it is that the end result of unboxing will replace the outer box by the inner boxed value; therefore we wish to treat the boxed value as having the same traceability as its contents.

Our goal is to show that the unboxing function induced by any acceptable unboxing is in some sense correct as an optimization. The first part of this is to show that unboxing preserves traceability.

Theorem 1 (Consistent unboxings preserve traceability).

- *If* $\mathrm{T}, \Upsilon \vdash v^i$ *and* $\vdash_{\mathrm{v}} v^i{:}t$ *then* $\vdash_{\mathrm{v}} \lfloor v^i \rfloor^{\mathrm{T}}_{\Upsilon}{:}\underline{\mathrm{T}}(i,t)$.
- *If* $\mathrm{T}, \Upsilon \vdash e$ *and* $\vdash e$ **tr** *then* $\vdash \lfloor e \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**.
- *If* $\mathrm{T}, \Upsilon \vdash \rho$ *and* $\vdash \rho$ **tr** *then* $\vdash \lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**.
- *If* $\mathrm{T}, \Upsilon \vdash M$ *and* $\vdash M$ **tr** *then* $\vdash \lfloor M \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**.

[2] See the cache refinement lemma in Appendix A for more detail.

Theorem 1 tells us that if we have a traceable program, then the result of unboxing it is still traceable. The second step to showing correctness is to show that unboxing does not introduce new undefined behavior.

Theorem 2 (Coherence).

- *If* $\mathrm{C}; \varrho \vdash M$, $\mathrm{C} \vdash M \downharpoonright\downharpoonright (\mathrm{T}, \Upsilon)$, *and* $M \longmapsto^* (\rho, v^i)$ *then* $\lfloor M \rfloor^{\mathrm{T}}_{\Upsilon} \longmapsto^* (\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor v^i \rfloor^{\mathrm{T}}_{\Upsilon})$.
- *If* $\mathrm{C}; \varrho \vdash M$, $\mathrm{C} \vdash M \downharpoonright\downharpoonright (\mathrm{T}, \Upsilon)$, *and* $M \longmapsto \cdots$ *then* $\lfloor M \rfloor^{\mathrm{T}}_{\Upsilon} \longmapsto \cdots$.

Theorem 2 shows that if two terms are related by reduction, then their images under the unboxing function are also related by the many step reduction relation given that the unboxing pair is acceptable; and that if a term diverges under reduction, then its image under the unboxing function also diverges. In other words, for an acceptable analysis and an acceptable unboxing, the induced unboxing function preserves the semantics of the original program up to elimination of boxes. Since the semantics of the core language only defines reduction steps that preserve GC safety, this theorem implies that the image of a GC safe program under unboxing is also GC safe.

5 Construction of An Acceptable Unboxing

The previous section gives a declarative specification for when an unboxing pair (T, Υ) is correct but does not specify how such a pair might be produced. In this section we give a simple algorithm for constructing an acceptable unboxing given an arbitrary acceptable flow analysis.

The idea behind the algorithm is that given a program and an acceptable flow analysis for it, we use the results of the flow analysis to construct the connected components of the inter-procedural flow graph of the program. Each connected component initially defines its own equivalence class. For each equivalence class, we then compute the least upper bound of the traceabilities of all of the introduction forms of all of the elements of the component except the box introductions. Box introductions are left initially unconstrained, since we intend to eliminate them. If the least upper bound is well-defined, then the equivalence class can potentially be eliminated. We then consider each box introduction in turn and attempt to eliminate it by combining the respective equivalence classes of the box and its contents. This is possible whenever doing so will not over-constrain the resulting combined equivalence class. When all possible boxes have been eliminated, the algorithm terminates. In the rest of the section, we make this informal algorithm concrete and show that the choice of unboxing that it produces is in fact acceptable.

For the purposes of this section we ignore environments and the intermediate forms $\rho(e)$, $\langle \rho, \lambda x^i{:}t.e \rangle^j$ and $\langle v^i{:}t \rangle^j$. These constructs are present in the language solely as mechanisms to discuss the dynamic semantics—in this sense they can be thought of as intermediate terms, rather than source terms. It is straightforward to incorporate these into the algorithm if desired.

Given a flow analysis (C, ϱ), we define the induced undirected flow graph $\mathcal{FG}$ as an undirected graph with a node for every label in C, and edges as follows.

- For every label i and every shape $s \in \mathrm{C}(i)$, we add an edge between i and $\mathrm{lbl}(s)$.
- For every box introduction in the program $(\mathtt{box}_t\, e)^i$ and every shape in the cache $(\mathtt{box}_t\, j)^i \in \mathrm{C}(i)$ we add an edge between $\mathrm{lbl}(e)$ and j.

The first set of edges simply connects up each program point with all of its reaching definitions. The second set of edges is added to simplify the proofs in the pathological case that e has no reaching definitions (and hence the box itself is dynamically dead and unreachable): in the usual case where values reach e then the definition of an acceptable analysis implies that these edges are already present.

We define equivalence classes of labels as disjoint sets of labels in the usual way. The function $\mathcal{EC}$ maps labels i to the disjoint set containing i. We extend traceabilities t to a complete flat lattice $\hat{t}$ with a top element $\top$, a bottom element $\bot$ and the usual least upper bound function on $\hat{t}$.

$$\begin{array}{llll} \top \sqcup\ t\ = \top & \bot \sqcup\ t\ = t & & \mathtt{b} \sqcup \mathtt{r} = \top \\ t\ \sqcup \top = \top & t\ \sqcup \bot = t & t \sqcup t = t & \mathtt{r} \sqcup \mathtt{b} = \top \end{array}$$

We initialize the mapping $\mathcal{EC}$ by finding the connected components of the induced undirected flow-graph $\mathcal{FG}$, and initializing $\mathcal{EC}(i)$ with the connected component containing i. As the algorithm proceeds, two equivalence classes may be collapsed into a single equivalence class requiring an updated mapping $\mathcal{EC}$.

We maintain a set of equivalence classes $\mathcal{CN}$ consisting of current candidates for unboxing. When equivalence classes are collapsed, the elements of $\mathcal{CN}$ are adjusted appropriately, as will be shown.

We maintain a set of labels Υ, which is an unboxing set in the sense of Section 4. The set Υ at all times contains all of the labels already selected for unboxing, and is initially empty.

We maintain an extended traceability map $\mathcal{T}$ that maps equivalence classes to extended traceabilities $\hat{t}$. For notational convenience we define $\mathcal{T}_{\mathcal{EC}}$ to be the derived function mapping labels to the extended traceabilities of their equivalence classes: $\mathcal{T}_{\mathcal{EC}}(k) = \mathcal{T}(\mathcal{EC}(k))$. The derivation of a standard traceability map T from an extended traceability map $\mathcal{T}$ is then given as follows.

$$\begin{array}{ll} \mathrm{T}(k) = \mathcal{T}_{\mathcal{EC}}(k) & \bot < \mathcal{T}_{\mathcal{EC}}(k) < \top \\ \mathrm{T}(k) = r & \mathcal{T}_{\mathcal{EC}}(k) = \bot \\ \mathrm{T}(k) = \text{undefined} & \mathcal{T}_{\mathcal{EC}}(k) = \top \end{array}$$

The general idea is that an equivalence class is mapped by the $\mathcal{T}$ function to the least upper bound of the traceabilities of all of the reaching definitions of all of the labels in the equivalence class. An equivalence class containing no reaching definitions will be unconstrained – for technical reasons we choose an arbitrary traceability ($\mathtt{r}$) for such classes. An equivalence class containing definitions with

inconsistent traceabilities will have no defined traceability in the induced mapping.

During the algorithm traceability constraints imposed by box introductions in the candidate set are left out of the initial mapping and hence must be added back in before computing the induced traceability map. We write $\mathcal{T}^{\mathcal{CN}}$ for the extended traceability map obtained by adding in the delayed constraints for each equivalence class, and $\mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}$ for the extension of this to individual labels given by $\mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}(i) = \mathcal{T}^{\mathcal{CN}}(\mathcal{EC}(i))$.

$$\begin{array}{l} \mathcal{T}^{\mathcal{CN}}(\mathcal{EC}(k)) = \mathcal{T}(\mathcal{EC}(k)) \qquad \text{if } \mathcal{EC}(k) \notin \mathcal{CN} \\ \mathcal{T}^{\mathcal{CN}}(\mathcal{EC}(k)) = \mathcal{T}(\mathcal{EC}(k)) \sqcup \mathtt{r} \text{ if } \mathcal{EC}(k) \in \mathcal{CN} \end{array}$$

Note that by definition, if labels i and j are in the same equivalence class ($\mathcal{EC}(i) = \mathcal{EC}(j)$), then the traceability map T induced by an extended traceability map $\mathcal{T}$ agrees on i and j.

We define the immediate extended traceability of a labeled term $itr(e)$ as follows.

$$\begin{array}{llll} itr(c^i) & = b & itr((\mathtt{box}_t\, e)^i) & = r \\ itr((\lambda x^i{:}t.e)^j) & = r & itr(e) & = \bot \text{ otherwise} \end{array}$$

The algorithm starts with an empty unboxing set Υ. The candidate set $\mathcal{CN}$ is initialized by including $\mathcal{EC}(i)$ for each $(\mathtt{box}_t\, e)^i$ in the program. The extended traceability map is initialized by setting for each equivalence class S:

$$\mathcal{T}(S) = \bigsqcup_{i \in S,\ s \in \mathrm{C}(i),\ s \neq (\mathtt{box}_t\, k)^j} itr(s)$$

That is, we take the extended traceability associated with the equivalence class S to be least upper bound of the immediate traceabilities of all of the elements of the cache of all the labels in the equivalence class, except those which are box introductions. The practical effect of this is to make the extended traceability of every label be the least upper bound of the traceability of every introduction form in its connected component (again, excepting boxes). An equivalence class that is unconstrained ($\bot$) either counts only box introductions among its definitions, or contains no definitions at all and hence is uninhabited (this can arise because of unreachable code).

The result of the initialization phase is a $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ quadruple, which induces an unboxing pair (T, Υ) where $\mathrm{T} = \mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}$. It can be shown that the unboxing pair induced in this manner is acceptable.

Lemma 3 (The initial unboxing is acceptable). *If* $\mathrm{C}; \varrho \vdash e$ *then the unboxing quadruple* $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ *computed by the algorithm in this section induces an unboxing* (T, Υ) *(where* $\mathrm{T} = \mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}$*) such that* $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$ *and* $\mathrm{T}, \Upsilon \vdash e$.

The unboxing pair created by the initial phase is acceptable, but does no unboxing. The second phase of the algorithm proceeds by incrementally moving equivalence classes from the candidate set $\mathcal{CN}$ to the unboxing set Υ, while maintaining the invariant that at every step $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ define an acceptable

(and increasingly useful) unboxing. Equivalence classes of boxes that get chosen for unboxing are collapsed into the same equivalence class as the contents of the box. We use the notation $\mathcal{EC}' = \cup_{\underline{i},\underline{j}}\mathcal{EC}$ to stand for combining the equivalence classes for i and j to get a new equivalence class in the usual way.

For the unboxing steps, we consider in turn each $(\mathtt{box}_t\ e)^i$ in the program. Let T be the traceability map induced by $\mathcal{T}$. The principal selection criterium for choosing which things to unbox is that $\mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(\mathrm{lbl}(e)) < \top$. The idea is that under the assumption that no further unboxing is done, combining the equivalence classes for i and $\mathrm{lbl}(e)$ will not over-constrain the resulting equivalence class, and hence that the final induced traceability map will be well-defined at i and $\mathrm{lbl}(e)$. The extended traceability $\mathcal{T}_{\mathcal{EC}}(i)$ is the extended traceability associated with i under the assumption that $\mathcal{EC}(i)$ is unboxed, while $\mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(\mathrm{lbl}(e))$ is the extended traceability associated with $\mathrm{lbl}(e)$ under the assumption that no further unboxing is done. If i is either unconstrained, or constrained to something compatible with $\mathrm{lbl}(e)$, then it is safe to unbox it.

Formally, if we have that $\mathcal{EC}(i) \in \mathcal{CN}$, $\mathcal{EC}(\mathrm{lbl}(e)) \neq \mathcal{EC}(i)$, and $\mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(\mathrm{lbl}(e)) < \top$ then we select i for elimination. We then take the new unboxing to be the updated quadruple $(\mathcal{T}', \Upsilon', \mathcal{CN}', \mathcal{EC}')$ where:

$$
\begin{array}{lll}
\mathcal{EC}' & = \cup_{\mathrm{lbl}(e),i}\mathcal{EC} & \\
\mathcal{CN}' & = (\overline{\mathcal{CN} - \{\mathcal{EC}(\mathrm{lbl}(e)), \mathcal{EC}(i)\}}) \cup \{\mathcal{EC}'(i)\} & \text{if } \mathcal{EC}(\mathrm{lbl}(e)) \in \mathcal{CN} \\
 & = (\mathcal{CN} - \{\mathcal{EC}(\mathrm{lbl}(e)), \mathcal{EC}(i)\}) & \text{if } \mathcal{EC}(\mathrm{lbl}(e)) \notin \mathcal{CN} \\
\Upsilon' & = \Upsilon \cup \{\mathcal{EC}(i)\} & \\
\mathcal{T}'(s) & = \mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathcal{T}_{\mathcal{EC}}(\mathrm{lbl}(e)) & \text{if } s = \mathcal{EC}'(i) \\
 & = \mathcal{T}(s) & \text{otherwise}
\end{array}
$$

For $\mathcal{EC}'$, we repartition the graph so that the equivalence classes for the box and its contents are combined into a single equivalence class. We remove the two original equivalence classes from the candidate set, and if the contents of the box was a candidate for unboxing we add back in the new equivalence class, which is the union of the two original classes. All of the elements of the original equivalence class of the box introduction are added to the unbox set. The extended traceability map is updated to map the new equivalence class (including both i and $\mathrm{lbl}(e)$) to the extended traceability of the contents of the box.

If the conditions for unboxing i are not satisfied, then we take $\mathcal{CN}' = \mathcal{CN} - \{\mathcal{EC}(i)\}$ and take $\mathcal{T}'(\mathcal{EC}(i)) = \mathcal{T}(\mathcal{EC}(i)) \sqcup r$ and leave the rest of the data structures unchanged. Since we only consider each box introduction in the program once, the algorithm terminates.

Lemma 4 (Unboxing steps preserve acceptability). *If $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ define an acceptable unboxing as constructed by the initial phase of the algorithm and maintained by the unboxing phase, then the $(\mathcal{T}', \Upsilon', \mathcal{CN}', \mathcal{EC}')$ quadruple produced by a single step of the algorithm above also define an acceptable unboxing.*

Lemma 4 states that each step of the unboxing phase of the algorithm preserves the property that the induced unboxing pair is acceptable. Consequently, the algorithm terminates with an acceptable unboxing.

Theorem 3 (The algorithm produces an acceptable unboxing). *If* $\mathrm{C}; \varrho \vdash e$ *then the algorithm defined in this section produces a quadruple* $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ *such that* $\mathrm{C} \vdash e \downharpoonleft\downharpoonright (\mathrm{T}, \Upsilon)$ *where* $\mathrm{T} = \mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}$.

This construction demonstrates that the specification defined in Section 4 is a useful one in the sense that it is satisfiable. While the algorithm defined here is unlikely to be optimal, it has proved very effective in our compiler: on floating-point intensive benchmarks we have measured an order of magnitude reduction in allocation and substantial performance and parallel scalability gains.

6 Related Work

This paper provides a modular approach to showing correctness of a realistic compiler optimization that rewrites the structure of program data structures in significant ways. Our approach uses an arbitrary inter-procedural reaching definitions analysis to eliminate unnecessary heap allocation in an intermediate representation in which object representation has been made explicit. Our optimization can be staged freely with other optimizations. Unlike any previous work that we are aware of, we account for correctness with respect to the meta-data requirements of the garbage collector. For presentational purposes, we have restricted our attention to the core concern of GC safety, but additional issues such as value size, dynamic type tests, etc. are straightforward to incorporate.

There has been substantial previous work addressing the problem of unboxing. Peyton Jones [3] introduced an explicit distinction between boxed and unboxed objects to provide a linguistic account of unboxing, and hence to allow a high-level compiler to locally eliminate unboxes of syntactically apparent box introduction operations. Leroy [4] defined a type-driven approach to adding coercions into and out of specialized representations. The type driven translation represented monomorphic objects natively (unboxed, in our terminology), and then introduced wrappers to coerce polymorphic uses into an appropriate form. To a first-order approximation, instead of boxing at definition sites this approach boxes objects at polymorphic use sites. This style of approach has the problem that it is not necessarily beneficial, since allocation is introduced in places where it would not otherwise be present. This is reflected in the slowdowns observed on some benchmarks described in the original paper. This approach also has the potential to introduce space leaks. In a later paper [5] Leroy argued that a simple untyped approach gives better and more predictable results.

Henglein and Jørgensen [2] defined a formal notion of optimality for local unboxings and gave two different choices of coercion placements that satisfy their notion of optimality. Their definition of optimality explicitly does not correspond in any way to reduced allocation or reduced instruction count and does not seem to provide uniform improvement over Leroy's approach.

The MLton compiler [10] largely avoids the issue of a uniform object representation by completely monomorphizing programs before compilation. This approach requires whole-program compilation. More limited monomorphization

schemes could be considered in an incremental compilation setting. Monomorphization does not eliminate the need for boxing in the presence of dynamic type tests or reflection. Just in time compilers (e.g. for .NET) may monomorphize dynamically at runtime.

The TIL compiler [1, 9] uses intensional type analysis in a whole-program compiler to allow native data representations without committing to whole-program compilation. As with the Leroy coercion approach, polymorphic uses of objects require conditionals and boxing coercions to be inserted at use sites, and consequently there is the potential to slow down, rather than speed up, the program.

Serrano and Feeley [8] described a flow analysis for performing unboxing substantially similar in spirit to our approach. Their algorithm attempts to find a monomorphic typing for a program in which object representations have not been made explicit, which they then use selectively to choose whether to use a uniform or non-uniform representation for each particular object. Their approach differs in that they define a dedicated analysis rather than using a generic reaching definitions analysis. They assume a conservative garbage collector and hence do not need to account for the requirements of GC safety, and they do not prove a correctness result.

References

1. Harper, R., Morrisett, G.: Compiling polymorphism using intensional type analysis. In: Twenty-Second ACM Symposium on Principles of Programming Languages. pp. 130–141. San Francisco, CA (January 1995)
2. Henglein, F., Jørgensen, J.: Formally optimal boxing. In: Proceedings of the 21st ACM SIGPLAN-SIGACT symposium on Principles of programming languages. pp. 213–226. POPL '94, ACM, New York, NY, USA (1994)
3. Jones, S.L.P., Launchbury, J.: Unboxed values as first class citizens in a non-strict functional language. In: Proceedings of the 5th ACM conference on Functional programming languages and computer architecture. pp. 636–666. Springer-Verlag New York, Inc. (1991)
4. Leroy, X.: Unboxed objects and polymorphic typing. In: Proceedings of the 19th ACM SIGPLAN-SIGACT symposium on Principles of programming languages. pp. 177–188. POPL '92, ACM, New York, NY, USA (1992)
5. Leroy, X.: The effectiveness of type-based unboxing. Tech. rep., Boston College, Computer Science Department (1997)
6. Nielson, F., Nielson, H.R., Hankin, C.: Principles of Program Analysis. Springer-Verlag New York, Inc., Secaucus, NJ, USA (1999)
7. Petersen, L., Glew, N.: GC-safe interprocedural unboxing: Extended version. Tech. rep., http://leafpetersen.com/leaf/papers.html
8. Serrano, M., Feeley, M.: Storage use analysis and its applications. In: Proceedings of the first ACM SIGPLAN international conference on Functional programming. pp. 50–61. ICFP '96, ACM, New York, NY, USA (1996)
9. Tarditi, D., Morrisett, G., Cheng, P., Stone, C., Harper, R., Lee, P.: Til: a type-directed, optimizing compiler for ml. SIGPLAN Not. 39, 554–567 (April 2004)
10. Weeks, S.: Whole-program compilation in MLton. In: Proceedings of the 2006 workshop on ML. pp. 1–1. ML '06, ACM, New York, NY, USA (2006)

A Appendix: Proofs

A.1 Proofs for Section 2

Proof of Lemma 1

Proof. If $\vdash (\rho, e)$ **tr** then $\vdash \rho$ **tr** and $\vdash e$ **tr**. If $(\rho, e) \longmapsto (\rho, e')$ then the result follows if we can show $\vdash e'$ **tr**. The proof of that is by induction on the derivation of $(\rho, e) \longmapsto (\rho, e')$. Consider the cases for the last rule used to derive it (the cases are in the same order as in the figure):

- In this case, $e = x^k$ for some x and k, and $e' = v^j$ where $x^i{:}t = v^j \in \rho$ for some i, t, v, and j. Since $\vdash \rho$ **tr** then $\vdash_{\mathrm{v}} v{:}t$, so by the traceability rules $\vdash v^j$ **tr** as required.
- In this case, $e = (\lambda x^i{:}t.e'')^j$ for some x, i, t, e'', and j, and $e' = \langle \rho, \lambda x^i{:}t.e'' \rangle^j$. The first hypothesis is that $\vdash (\lambda x^i{:}t.e'')^j$ **tr**. There is only one rule to derive this judgement and that rule requires that $\vdash \lambda x^i{:}t.e''$ **tr**, which in turn can only be derived by one rule that requires that $\vdash e''$ **tr**. Then, and since $\vdash \rho$ **tr**, by the rules for traceability, $\vdash_{\mathrm{v}} \langle \rho, \lambda x^i{:}t.e'' \rangle{:}\mathtt{r}$, and by the traceability rules again $\vdash \langle \rho, \lambda x^i{:}t.e'' \rangle^j$ **tr**, as we are required to prove.
- In this case, $e = (\mathtt{box}_t\ {v_{t'}}^i)^j$ for some t, $v_{t'}$, i, and j, $e' = \langle {v_{t'}}^i{:}t \rangle^j$, and $t = t'$. The first hypothesis is that $\vdash (\mathtt{box}_t\ {v_{t'}}^i)^j$ **tr**. There is only one rule to derive this judgement and that rule requires that $\vdash \mathtt{box}_t\ {v_{t'}}^i$ **tr**, which in turn can only be derived by one rule that requires that $\vdash {v_{t'}}^i$ **tr**. There is only one rule to derive the latter judgement and it requires that $\vdash_{\mathrm{v}} v_{t'}{:}t''$ for some t''. By inspection of the rules for value traceability, we see that $t' = t''$. Since $t = t' = t''$, by the rules for traceability, $\vdash_{\mathrm{v}} \langle {v_{t'}}^i{:}t \rangle{:}\mathtt{r}$, and by the traceability rules again $\vdash \langle {v_{t'}}^i{:}t \rangle^j$ **tr**, as we are required to prove.
- In this case, $e = (e_1\ e_2)^i$ for some e_1, e_2, and i, $e' = (e_1'\ e_2)^i$ for some e_1', and $(\rho, e_1) \longmapsto (\rho, e_1')$ is a subderivation. The first hypothesis is $\vdash (e_1\ e_2)^i$ **tr**. There is only one rule to derive this judgement and that rule requires that $\vdash e_1\ e_2$ **tr**, which in turn can only be derived by one rule that requires both $\vdash e_1$ **tr** and $\vdash e_2$ **tr**. Thus, by the induction hypothesis, $\vdash e_1'$ **tr**. Then, by the rules for traceability, $\vdash e_1'\ e_2$ **tr**, and by the traceability rules again, $\vdash (e_1'\ e_2)^i$ **tr**, as we are required to prove.
- In this case, $e = (v^i\ e_2)^j$ for some v, i, e_2, and j, $e' = (v^i\ e_2')^j$ for some e_2', and $(\rho, e_2) \longmapsto (\rho, e_2')$ is a subderivation. The first hypothesis is $\vdash (v^i\ e_2)^j$ **tr**. There is only one rule to derive this judgement and that rule requires that $\vdash v^i\ e_2$ **tr**, which in turn can only be derived by one rule that requires both $\vdash v^i$ **tr** and $\vdash e_2$ **tr**. Thus, by the induction hypothesis, $\vdash e_2'$ **tr**. Then, by the rules for traceability, $\vdash v^i\ e_2'$ **tr**, and by the traceability rules again, $\vdash (v^i\ e_2')^j$ **tr**, as we are required to prove.
- In this case, $e = \left(\langle \rho', \lambda x^i{:}t.e'' \rangle^j\ {v_{t'}}^k\right)^l$ for some ρ', x, i, t, e'', j, $v_{t'}$, k, and l, $e' = (\rho', x^i{:}t = v_{t'}{}^k)(e'')^l$, and $t = t'$. The first hypothesis is that $\vdash \left(\langle \rho', \lambda x^i{:}t.e'' \rangle^j\ {v_{t'}}^k\right)^l$ **tr**. There is only one rule to derive that judgement

and that rule requires that $\vdash \langle \rho', \lambda x^i{:}t.e'' \rangle^j \ {v_{t'}}^k$ **tr**, which in turn can only be derived by one rule that requires both $\vdash \langle \rho', \lambda x^i{:}t.e'' \rangle^j$ **tr** and $\vdash {v_{t'}}^k$ **tr**. Both of these latter derivations can only be derived by one rule and those rules require that $\vdash_{\mathrm{v}} \langle \rho', \lambda x^i{:}t.e'' \rangle{:}t_1$ (1) and $\vdash_{\mathrm{v}} v_{t'}{:}t_2$ (2) for some t_1 and t_2. Judgement 1 can only be derived by one rule and that rule requires that $\vdash \rho'$ **tr** (3) and $\vdash e''$ **tr** (4). By (3) and (2) we can derive $\vdash \rho', x^i{:}t = {v_{t'}}^k$ **tr** (5). By (5) and (1) we can derive $\vdash (\rho', x^i{:}t = {v_{t'}}^k)(e'')$ **tr**, and by then $\vdash e'$ **tr**, as required.

- In this case, $e = (\mathtt{box}_t\, e'')^i$ for some t, e'', and i, $e' = (\mathtt{box}_t\, e''')^i$, and $(\rho, e'') \longmapsto (\rho, e''')$ is a subderivation. The first hypothesis is $\vdash (\mathtt{box}_t\, e'')^i$ **tr**. There is only one rule to derive this judgement and that rule requires that $\vdash \mathtt{box}_t\, e''$ **tr**, which in turn can only be derived by one rule that requires that $\vdash e''$ **tr**. Thus, by the induction hypothesis, $\vdash e'''$ **tr**. Then, by the rules for traceability, $\vdash \mathtt{box}_t\, e'''$ **tr**, and by the traceability rules again, $\vdash (\mathtt{box}_t\, e''')^i$ **tr**, as we are required to prove.
- In this case, $e = (\mathtt{unbox}\, e'')^i$ for some e'' and i, $e' = (\mathtt{unbox}\, e''')^i$, and $(\rho, e'') \longmapsto (\rho, e''')$ is a subderivation. The first hypothesis is $\vdash (\mathtt{unbox}\, e'')^i$ **tr**. There is only one rule to derive this judgement and that rule requires that $\vdash \mathtt{unbox}\, e''$ **tr**, which in turn can only be derived by one rule that requires that $\vdash e''$ **tr**. Thus, by the induction hypothesis, $\vdash e'''$ **tr**. Then, by the rules for traceability, $\vdash \mathtt{unbox}\, e'''$ **tr**, and by the traceability rules again, $\vdash (\mathtt{unbox}\, e''')^i$ **tr**, as we are required to prove.
- In this case, $e = (\mathtt{unbox}\, \langle v^i{:}t \rangle^j)^k$ for some t, v, i, j, and k, and $e' = v^i$. The first hypothesis is that $\vdash (\mathtt{unbox}\, \langle v^i{:}t \rangle^j)^k$ **tr**. There is only one rule to derive this judgement and that rule requires that $\vdash \mathtt{unbox}\, \langle v^i{:}t \rangle^j$ **tr**, which in turn can only be derived by one rule that requires that $\vdash \langle v^i{:}t \rangle^j$ **tr**. There is only one rule to derive this latter judgement and that rule requires that $\vdash_{\mathrm{v}} \langle v^i{:}t \rangle{:}t'$ for some t', which in turn can only be derived by one rule that requires that $\vdash_{\mathrm{v}} v{:}t$. Then, by the rules for traceability, $\vdash v^i$ **tr**, as we are required to prove.
- In this case, $e = \rho'(e'')^i$ for some ρ', e'' and i, $e' = \rho'(e''')^i$ for some e''', and $(\rho', e'') \longmapsto (\rho', e''')$. The hypothesis $\vdash e$ **tr** can only be derived in a certain way, unpacking that we see that $\vdash \rho'$ **tr** and $\vdash e''$ **tr**. Then by the induction hypothesis, $\vdash e'''$ **tr**. So applyling the rules, we derive that $\vdash \rho'(e''')$ **tr** and then $\vdash e'$ **tr**, as required.
- In this case, $e = \rho'(v^i)^j$ for some ρ', v, i, and j, and $e' = v^i$. The hypothesis, $\vdash e$ **tr** can only be derived in one way and unpacking that we see that $\vdash v^i$ **tr**, which is what we are required to prove.

A.2 Proofs for Section 3

Lemma 5 (Cache refinement under reduction). *If* $\mathrm{C}; \varrho \vdash \rho$, $\mathrm{C}; \varrho \vdash e_1$, *and* $(\rho, e_1) \longmapsto (\rho, e_2)$ *then* $\mathrm{C}(\mathrm{lbl}(e_1)) \supseteq \mathrm{C}(\mathrm{lbl}(e_2))$.

Proof. The proof is by induction on the derivation of $(\rho, e_1) \longmapsto (\rho, e_2)$. Consider the cases for the last rule used to it (the cases are in the same order as in the figure):

- (Variable instantiation.) In this case, $e_1 = x^k$, $e_2 = v^j$, and $x^i{:}t = v^j \in \rho$. The assumption $\mathrm{C}; \varrho \vdash \rho$ requires that $\mathrm{C}(j) \subseteq \mathrm{C}(i)$ and $\varrho(x) = i$. The assumption $\mathrm{C}; \varrho \vdash e_1$ requires that $\mathrm{C}(\varrho(x)) \subseteq \mathrm{C}(k)$. Thus $\mathrm{C}(j) \subseteq \mathrm{C}(k)$. Clearly, $\mathrm{lbl}(e_1) = k$ and $\mathrm{lbl}(e_2) = j$ and the result follows.
- (Lambda introduction.) In this case, clearly $\mathrm{lbl}(e_1) = \mathrm{lbl}(e_2)$ and the result immediately follows.
- (Box introduction.) In this case, clearly $\mathrm{lbl}(e_1) = \mathrm{lbl}(e_2)$ and the result immediately follows.
- (Application left.) In this case, clearly $\mathrm{lbl}(e_1) = \mathrm{lbl}(e_2)$ and the result immediately follows.
- (Application right.) In this case, clearly $\mathrm{lbl}(e_1) = \mathrm{lbl}(e_2)$ and the result immediately follows.
- (Application beta.) In this case, clearly $\mathrm{lbl}(e_1) = \mathrm{lbl}(e_2)$ and the result immediately follows.
- (Under box.) In this case, clearly $\mathrm{lbl}(e_1) = \mathrm{lbl}(e_2)$ and the result immediately follows.
- (Under unbox.) In this case, clearly $\mathrm{lbl}(e_1) = \mathrm{lbl}(e_2)$ and the result immediately follows.
- (Unbox beta.) In this case, $e_1 = (\mathtt{unbox}\,\langle v^i{:}t\rangle^j)^k$ and $e_2 = v^i$. The first hypothesis can be derived only by one rule and it requires that $\mathrm{C}; \varrho \vdash \langle v^i{:}t\rangle^j$ (1), and (2) for all $(\mathtt{box}_{t'}\, i')^{j'} \in \mathrm{C}(j)$ that $\mathrm{C}(i') \subseteq \mathrm{C}(k)$. Judgement 1 can only be derived by one rule and it requires that $\mathrm{C}; \varrho \vdash v^i$ (4), $(\mathtt{box}_t\, i'')^j \in \mathrm{C}(j)$ (5) for some i'', and $\mathrm{C}(i) \subseteq \mathrm{C}(i'')$ (6). Instantiating Fact 2 with Fact 5 we get that $\mathrm{C}(i'') \subseteq \mathrm{C}(k)$ (7). Combining Facts 6 and 7, $\mathrm{C}(i) \subseteq \mathrm{C}(k)$, as we are required to prove.
- (Under frame.) In this case, clearly $\mathrm{lbl}(e_1) = \mathrm{lbl}(e_2)$ and the result immediately follows.
- (Frame return.) In this case, $e_1 = \rho'(v^i)^j$ and $e_2 = v^i$. The assumption $\mathrm{C}; \varrho \vdash e_1$ requires that $\mathrm{C}(i) \subseteq \mathrm{C}(j)$. Since $\mathrm{lbl}(e_1) = j$ and $\mathrm{lbl}(e_2) = i$, the result is immediate.

Lemma 6 (Preservation of acceptability under reduction). *If* $\mathrm{C}; \varrho \vdash M$ *and* $M \longmapsto M'$ *then* $\mathrm{C}; \varrho \vdash M'$.

Proof. If $\mathrm{C}; \varrho \vdash (\rho, e)$ then $\mathrm{C}; \varrho \vdash \rho$ and $\mathrm{C}; \varrho \vdash e$. If $(\rho, e) \longmapsto (\rho, e')$ then the result follows if we show that $\mathrm{C}; \varrho \vdash e'$. The proof of the latter is by induction on the derivation of $(\rho, e) \longmapsto (\rho, e')$. Consider the cases for the last rule used to derive it (the cases are in the same order as in the figure):

- In this case $e = x^k$, $e' = v^j$, and $x^i{:}t = v^j \in \rho$. The assumption $\mathrm{C}; \varrho \vdash \rho$ requires that $\mathrm{C}; \varrho \vdash v^j$, which is what we need to prove.

– In this case $e = (\lambda x^i{:}t.e'')^j$ for some x, i, t, e'', and j, and $e' = \langle \rho, \lambda x^i{:}t.e'' \rangle^j$. The first hypothesis can only be derived by one rule and it requires that $\varrho(x) = i$, $\mathrm{C}; \varrho \vdash e''$, and $(i{:}t \to \mathrm{lbl}(e''))^j \in \mathrm{C}(j)$. Then, and nothing $\mathrm{C}; \varrho \vdash \rho$ by assumption, by the rules for acceptable analysis, $\mathrm{C}; \varrho \vdash \langle \rho, \lambda x^i{:}t.e'' \rangle^j$, as we are required to prove.

– In this case $e = (\mathtt{box}_t\, v^i)^j$ for some t, v, i, and j, and $e' = \langle v^i{:}t \rangle^j$. The first hypothesis can only be derived by one rule and it requires that $\mathrm{C}; \varrho \vdash v^i$, $(\mathtt{box}_t\, k)^j \in \mathrm{C}(j)$ for some k, and $\mathrm{C}(i) \subseteq \mathrm{C}(k)$. Then by the rules for acceptable analysis, $\mathrm{C}; \varrho \vdash \langle v^i{:}t \rangle^j$, as we are required to prove.

– In this case $e = (e_1\ e_2)^i$ for some e_1, e_2, and i, $e' = (e_1'\ e_2)^i$, and $(\rho, e_1) \longmapsto (\rho, e_1')$ is a subderivation. The first hypothesis can only be derived by one rule and it requires that $\mathrm{C}; \varrho \vdash e_1$ (1), $\mathrm{C}; \varrho \vdash e_2$ (2), and (3) for all $(i_2{:}t \to i')^{i_1} \in \mathrm{C}(\mathrm{lbl}(e_1))$, $\mathrm{C}(\mathrm{lbl}(e_2)) \subseteq \mathrm{C}(i_2)$ and $\mathrm{C}(i') \subseteq \mathrm{C}(i)$. By the induction hypothesis and Judgement 1, $\mathrm{C}; \varrho \vdash e_1'$ (4). By Lemma 5, $\mathrm{C}(\mathrm{lbl}(e_1')) \subseteq \mathrm{C}(\mathrm{lbl}(e_1))$ (5). Combining Facts 3 and 5, we see that (6) for all $(i_2{:}t \to i')^{i_1} \in \mathrm{C}(\mathrm{lbl}(e_1'))$, $\mathrm{C}(\mathrm{lbl}(e_2)) \subseteq \mathrm{C}(i_2)$ and $\mathrm{C}(i') \subseteq \mathrm{C}(i)$. Combining Facts 4, 2, and 6, and using the rules for acceptable analysis, we see that $\mathrm{C}; \varrho \vdash (e_1'\ e_2)^i$, as we are required to prove.

– In this case $e = (v^j\ e_2)^i$ for some v, j, e_2, and i, $e' = (v^j\ e_2')^i$, and $(\rho, e_2) \longmapsto (\rho, e_2')$ is a subderivation. The first hypothesis can only be derived by one rule and it requires that $\mathrm{C}; \varrho \vdash v^j$ (1), $\mathrm{C}; \varrho \vdash e_2$ (2), and (3) for all $(i_2{:}t \to i')^{i_1} \in \mathrm{C}(j)$, $\mathrm{C}(\mathrm{lbl}(e_2)) \subseteq \mathrm{C}(i_2)$ and $\mathrm{C}(i') \subseteq \mathrm{C}(i)$. By the induction hypothesis and Judgement 1, $\mathrm{C}; \varrho \vdash e_2'$ (4). By Lemma 5, $\mathrm{C}(\mathrm{lbl}(e_2')) \subseteq \mathrm{C}(\mathrm{lbl}(e_2))$ (5). Combining Facts 3 and 5, we see that (6) for all $(i_2{:}t \to i')^{i_1} \in \mathrm{C}(j)$, $\mathrm{C}(\mathrm{lbl}(e_2')) \subseteq \mathrm{C}(i_2)$ and $\mathrm{C}(i') \subseteq \mathrm{C}(i)$. Combining Facts 1, 4, and 6, and using the rules for acceptable analysis, we see that $\mathrm{C}; \varrho \vdash (v^j\ e_2')^i$, as we are required to prove.

– In this case $e = (\langle \rho', \lambda x^i{:}t.e'' \rangle^j\ v^k)^l$ for some x, i, t, e'', v, k, and l, and $e' = (\rho', x^i{:}t = v^k)(e'')^l$. The first hypothesis can only be derived by one rule and it requires that $\mathrm{C}; \varrho \vdash \langle \rho', \lambda x^i{:}t.e'' \rangle^j$ (1) and $\mathrm{C}; \varrho \vdash v^k$ (2), and (3) for all $(i'{:}t' \to l')^{j'} \in \mathrm{C}(j)$ that $\mathrm{C}(k) \subseteq \mathrm{C}(i')$ and $\mathrm{C}(l') \subseteq \mathrm{C}(l)$. Judgement 1 can only be derived by one rule and it requires that $\varrho(x) = i$ (4), $\mathrm{C}; \varrho \vdash e''$ (5), and $(i{:}t \to \mathrm{lbl}(e''))^j \in \mathrm{C}(j)$ (6). Instantiating Fact 3 with Fact 6, $\mathrm{C}(k) \subseteq \mathrm{C}(i)$ (7) and $\mathrm{C}(\mathrm{lbl}(e'')) \subseteq \mathrm{C}(l)$ (8). Since $\mathrm{C}; \varrho \vdash \rho$, (4), (7), and (2), we can derive $\mathrm{C}; \varrho \vdash \rho', x^i{:}t = v^k$ (9). By (9), 5), and (8), we can derive $\mathrm{C}; \varrho \vdash e'$, as required.

– In this case $e = (\mathtt{box}_t\, e_1)^i$ for some t, e_1, and i, $e' = (\mathtt{box}_t\, e_2)^i$ for some e_2, and $(\rho, e_1) \longmapsto (\rho, e_2)$ is a subderivation. The first hypothesis can only be derived by one rule and it requires that $\mathrm{C}; \varrho \vdash e_1$ (1), $(\mathtt{box}_t\, j)^i \in \mathrm{C}(i)$ (2) for some j, and $\mathrm{C}(\mathrm{lbl}(e_1)) \subseteq \mathrm{C}(j)$ (3). By the induction hypothesis and Judgement 1, $\mathrm{C}; \varrho \vdash e_2$ (4). By Lemma 5, $\mathrm{C}(\mathrm{lbl}(e_2)) \subseteq \mathrm{C}(\mathrm{lbl}(e_1))$ (5). Combining Facts 3 and 5 gives $\mathrm{C}(\mathrm{lbl}(e_2)) \subseteq \mathrm{C}(j)$ (6). Then by Facts 4, 2,

and 6, and using the rules for acceptable analysis, $\mathrm{C};\varrho \vdash (\mathtt{box}_t\, e_2)^i$, as we are required to prove.

- In this case $e = (\mathtt{unbox}\, e_1)^i$ for some e_1 and i, $e' = (\mathtt{unbox}\, e_2)^i$ for some e_2, and $(\rho, e_1) \longmapsto (\rho, e_2)$ is a subderivation. The first hypothesis can only be derived by one rule and it requires that $\mathrm{C};\varrho \vdash e_1$ (1) and (2) for all $(\mathtt{box}_t\, k)^j \in \mathrm{C}(\mathrm{lbl}(e_1))$, $\mathrm{C}(k) \subseteq \mathrm{C}(i)$. By Judgement (1) and the induction hypothesis, $\mathrm{C};\varrho \vdash e_2$ (3). By Lemma 5, $\mathrm{C}(\mathrm{lbl}(e_2)) \subseteq \mathrm{C}(\mathrm{lbl}(e_1))$ (4). Combining Facts 4 and 2 we get that (5) for all $(\mathtt{box}_t\, k)^j \in \mathrm{C}(\mathrm{lbl}(e_2))$, $\mathrm{C}(k) \subseteq \mathrm{C}(i)$. Combining Facts 4 and 5, by the rules for acceptable analysis, $\mathrm{C};\varrho \vdash (\mathtt{unbox}\, e_2)^i$, as we are required to prove.
- In this case $e = (\mathtt{unbox}\, \langle v^i{:}t\rangle^j)^k$ for some t, v, i, j, and k, and $e' = v^i$. The first hypothesis can only be derived by one rule and it requires that $\mathrm{C};\varrho \vdash \langle v^i{:}t\rangle^j$, which in turn can only be derived by one rule that requires that $\mathrm{C};\varrho \vdash v^i$, as we are required to prove.
- In this case $e = \rho'(e'')^i$, $e' = \rho'(e''')^i$, and $(\rho', e'') \longmapsto (\rho', e''')$ is a subderivation. Assumption $\mathrm{C};\varrho \vdash e$ requires that $\mathrm{C};\varrho \vdash \rho'$ (1), $\mathrm{C};\varrho \vdash e''$ (2), and $\mathrm{C}(\mathrm{lbl}(e'')) \subseteq \mathrm{C}(i)$ (3). By (1), (2), and the induction hypothesis, $\mathrm{C};\varrho \vdash e'''$ (4). By Lemma 5, $\mathrm{C}(\mathrm{lbl}(e''')) \subseteq \mathrm{C}(\mathrm{lbl}(e''))$ (5). Combining (3) and (5), $\mathrm{C}(\mathrm{lbl}(e''')) \subseteq \mathrm{C}(i)$ (6). Using (1), (4), and (6) we derive $\mathrm{C};\varrho \vdash \rho'(e''')^i$, as required.
- In this case $e = \rho'(v^i)^j$ and $e' = v^i$. The assumption $\mathrm{C};\varrho \vdash e$ unpacks to requiring that $\mathrm{C};\varrho \vdash v^i$, as required.

Proof of Lemma 2

Proof. The proof is by a straightforward induction on the length of the reduction sequence and Lemma 6.

A.3 Proofs for Section 4

Proof of Theorem 1

Proof. – Suppose $\mathrm{T}, \Upsilon \vdash v^i$ and $\vdash_{\mathrm{v}} v^i{:}t$:

- If $v = c$ then we're done.
- If $v^i = \langle \rho, \lambda x^m{:}t.e\rangle^j$ then by inversion of the $\mathrm{T}, \Upsilon \vdash v^i$ assumption we have that $\mathrm{T}(j) \geq \mathtt{r}$, so $\underline{\mathrm{T}}(j, \mathtt{r}) = \mathtt{r}$. By induction $\vdash \lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}$ $\mathbf{tr}$ and $\vdash \lfloor e\rfloor^{\mathrm{T}}_{\Upsilon}$ $\mathbf{tr}$, and by applying the traceability rule for value lambdas we get $\vdash_{\mathrm{v}} \lfloor\langle \rho, \lambda x^m{:}t.e\rangle^j\rfloor^{\mathrm{T}}_{\Upsilon}{:}\underline{\mathrm{T}}(j, \mathtt{r})$.
- If $v^i = \langle v^n{:}t'\rangle^j$ then:
  - If $j \in \Upsilon$ then $\lfloor\langle v^n{:}t'\rangle^j\rfloor^{\mathrm{T}}_{\Upsilon} = \lfloor v^n\rfloor^{\mathrm{T}}_{\Upsilon}$. By assumption we have $\vdash_{\mathrm{v}} v^n{:}t'$ and $\mathrm{T}, \Upsilon \vdash v^n$, so by induction we have $\vdash_{\mathrm{v}} \lfloor v^n\rfloor^{\mathrm{T}}_{\Upsilon}{:}\underline{\mathrm{T}}(n, t')$. By the premises of the rule for acceptable unboxing we also have that $\underline{\mathrm{T}}(n, t') = \mathrm{T}(n) = \mathrm{T}(j) = \underline{\mathrm{T}}(j, t)$, so we're done.

* If $j \notin \Upsilon$ then $\lfloor \langle v^n{:}t' \rangle^j \rfloor^{\mathrm{T}}_{\Upsilon} = \langle \lfloor v^n \rfloor^{\mathrm{T}}_{\Upsilon}{:}\underline{\mathrm{T}}(n,t') \rangle^j$. By assumption we have $\vdash_{\mathrm{v}} v^n{:}t'$ so by induction, we have $\vdash_{\mathrm{v}} \lfloor v^n \rfloor^{\mathrm{T}}_{\Upsilon}{:}\underline{\mathrm{T}}(n,t')$, so by the definition of traceability we have $\vdash_{\mathrm{v}} \langle \lfloor v^n \rfloor^{\mathrm{T}}_{\Upsilon}{:}\underline{\mathrm{T}}(n,t') \rangle^j{:}r$. By the premises of the acceptable unboxing rule, we have that $\mathrm{T}(j) \geq \mathtt{r}$, so $\mathtt{r} = \underline{\mathrm{T}}(j,\mathtt{r})$ and we're done.

– Suppose $\mathrm{T}, \Upsilon \vdash e$ and $\vdash e$ **tr**:
  • If $e = v^i$ then by the above arguments $\vdash_{\mathrm{v}} \lfloor v^i \rfloor^{\mathrm{T}}_{\Upsilon}{:}\underline{\mathrm{T}}(i,t)$, so $\vdash \lfloor v^i \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**.
  • If $e = x^i$ then the result follows immediately.
  • If $e = (\lambda x^i{:}t.e')^j$ then $\lfloor e \rfloor^{\mathrm{T}}_{\Upsilon} = (\lambda x^i{:}\underline{\mathrm{T}}(i,t).\lfloor e' \rfloor^{\mathrm{T}}_{\Upsilon})^j$. By induction $\vdash \lfloor e' \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**, so by definition $\vdash e$ **tr**.
  • If $e = (e_1\ e_2)^i$ then we apply the induction hypthesis to the two subexpressions and reconstruct the derivation with the application rule.
  • If $e = (\mathtt{box}_t\, e_1)^j$ and $\mathrm{lbl}(e_1) = i$ then:
    * If $j \in \Upsilon$ then $\lfloor e \rfloor^{\mathrm{T}}_{\Upsilon} = \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon}$. By inverting the original assumptions we get $\mathrm{T}, \Upsilon \vdash e_1$ and $\vdash e_1$ **tr**, so by induction $\vdash \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**.
    * If $j \notin \Upsilon$ then $\lfloor e \rfloor^{\mathrm{T}}_{\Upsilon} = (\mathtt{box}_{\underline{\mathrm{T}}(i,t)} \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon})^j$. By inverting the original assumptions we get $\mathrm{T}, \Upsilon \vdash e_1$ and $\vdash e_1$ **tr**, so by induction $\vdash \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**. By applying the box traceability rule, we get $\vdash (\mathtt{box}_{\underline{\mathrm{T}}(i,t)} \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon})^j$ **tr** and hence $\vdash \lfloor (\mathtt{box}_t\, e_1)^j \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**
  • If $e = (\mathtt{unbox}\, e_1)^j$ and $\mathrm{lbl}(e_1) = i$ then:
    * If $i \in \Upsilon$ then $\lfloor e \rfloor^{\mathrm{T}}_{\Upsilon} = \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon}$. By inverting the original assumptions we get $\mathrm{T}, \Upsilon \vdash e_1$ and $\vdash e_1$ **tr**, so by induction $\vdash \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**.
    * If $i \notin \Upsilon$ then $\lfloor e \rfloor^{\mathrm{T}}_{\Upsilon} = (\mathtt{unbox} \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon})^j$. By inverting the original assumptions we get $\mathrm{T}, \Upsilon \vdash e_1$ and $\vdash e_1$ **tr**, so by induction $\vdash \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**. By applying the unbox traceability rule, we get $\vdash (\mathtt{unbox} \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon})^j$ **tr** and hence $\vdash \lfloor (\mathtt{unbox}\, e_1)^j \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**
  • If $e = \rho(e')$ then the assumptions give $\mathrm{T}, \Upsilon \vdash \rho$, $\mathrm{T}, \Upsilon \vdash e'$, $\vdash \rho$ **tr**, and $\vdash e'$ **tr**. By the induction hypothesis, $\vdash \lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr** and $\vdash \lfloor e' \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**, from which $\vdash \lfloor e \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr** follows.

– Suppose $\mathrm{T}, \Upsilon \vdash \rho$ and $\vdash \rho$ **tr**.
Let $\rho = x_1^{i_1}{:}t_1 = {v_i}^{j_1}, \ldots, x_n^{i_n}{:}t_n = {v_n}^{j_n}$. By the assumptions, $\underline{\mathrm{T}}(i_k,t_k) = \underline{\mathrm{T}}(j_k,t_k)$, $\mathrm{T}, \Upsilon \vdash {v_k}^{j_k}$, and $\vdash_{\mathrm{v}} v_k{:}t_k$ for all $1 \leq k \leq n$. By induction, $\vdash_{\mathrm{v}} {v_k}^{j_k}{:}\underline{\mathrm{T}}(j_k,t_k)$ for all $1 \leq k \leq n$. So by the rules $\vdash x_1^{i_1}{:}\underline{\mathrm{T}}(i_1,t_1) = \lfloor {v_1}^{j_1} \rfloor^{\mathrm{T}}_{\Upsilon}, \ldots, x_n^{i_n}{:}\underline{\mathrm{T}}(i_n,t_n) = \lfloor {v_n}^{j_n} \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**, which is $\vdash \lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**.

– Suppose $\mathrm{T}, \Upsilon \vdash M$ and $\vdash M$ **tr**. Let $M = (\rho, e)$. Then by the assumptions, $\mathrm{T}, \Upsilon \vdash \rho$, $\mathrm{T}, \Upsilon \vdash e$, $\vdash \rho$ **tr**, and $\vdash e$ **tr**. Then by induction, $\vdash \lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr** and $\vdash \lfloor e \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**. So by the rules, $\vdash (\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e \rfloor^{\mathrm{T}}_{\Upsilon})$ **tr**, which is $\vdash \lfloor M \rfloor^{\mathrm{T}}_{\Upsilon}$ **tr**.

Lemma 7 (Inhabitance). *If* $\mathrm{C}; \varrho \vdash v^k$ *then* $\exists s \in \mathrm{C}(k)$ *s.t.* $\mathrm{lbl}(s) = k$.

Proof. By inspection of the acceptable analysis and acceptable instantiation rules.

Lemma 8 (Agreement). *If* $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$ *then* $\forall i, j : \phi \subset \mathrm{C}(i) \subseteq \mathrm{C}(j) \implies i \overset{\mathrm{T},\Upsilon}{\simeq} j$.

Proof. Since $\mathrm{C}(i) \subseteq \mathrm{C}(j)$ and since $\mathrm{C}(i)$ is not empty, there is at least one s in both. By definition of $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$ then, i and j each agree with $\mathrm{lbl}(s)$ and hence with each other.

Lemma 9 (Unboxing set preservation). *If* $\mathrm{C}; \varrho \vdash \rho$, $\mathrm{C}; \varrho \vdash e$, $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$, *and* $(\rho, e) \longmapsto (\rho, e')$ *then* $\mathrm{lbl}(e) \overset{\mathrm{T},\Upsilon}{\simeq} \mathrm{lbl}(e')$.

Proof. All of the cases for which $\mathrm{lbl}(e) = \mathrm{lbl}(e')$ follow immediately. For the remaining cases:

- If $(\rho, x^k) \longmapsto (\rho, v^j)$ where $x^i{:}t = v^j \in \rho$ then by the assumptions we have that $\varrho(x) = i$ (1), $\mathrm{C}(j) \subseteq \mathrm{C}(i)$ (2) and $\mathrm{C}(\varrho(x)) \subseteq \mathrm{C}(k)$ (3), so by transitivity we have $\mathrm{C}(j) \subseteq \mathrm{C}(k)$ (4). By Inhabitance (Lemma 7) we have an $s \in \mathrm{C}(j)$ (5) such that $\mathrm{lbl}(s) = j$ (6), and so by Agreement (Lemma 8) we have $k \overset{\mathrm{T},\Upsilon}{\simeq} j$. Since $\mathrm{lbl}(e) = k$ and $\mathrm{lbl}(e') = j$, the result follows.
- If $(\rho, (\mathtt{unbox}\, \langle v^i{:}t \rangle^j)^k) \longmapsto (\rho, v^i)$ then we must show that $k \overset{\mathrm{T},\Upsilon}{\simeq} i$. By Inhabitance we have $s \in \mathrm{C}(i)$ with $\mathrm{lbl}(s) = i$, so by Agreement, it suffices to show that $\mathrm{C}(i) \subseteq \mathrm{C}(k)$. By the box rule for an acceptable analysis, there is a $s = (\mathtt{box}_t\, m)^j \in \mathrm{C}(j)$ such that $\mathrm{C}(i) \subseteq \mathrm{C}(m)$. Since $s \in \mathrm{C}(j)$, by the rule for unbox, $\mathrm{C}(m) \subseteq \mathrm{C}(k)$, so $\mathrm{C}(i) \subseteq \mathrm{C}(k)$ and we're done.
- If $(\rho, \rho'(v^i)^j) \longmapsto (\rho, v^i)$ then we must show that $j \overset{\mathrm{T},\Upsilon}{\simeq} i$. By Inhabitance, there is an $s \in \mathrm{C}(i)$, and by the acceptable analysis rule for frames we have that $\mathrm{C}(i) \subseteq \mathrm{C}(j)$, so by Agreement we have that $j \overset{\mathrm{T},\Upsilon}{\simeq} i$.

We can show a single-step preservation result: if an unboxing is acceptable for a term and that term can take a reduction step, then the unboxing is still acceptable for the result of the reduction.

Lemma 10 (Single step preservation of consistency). *If* $\mathrm{C}; \varrho \vdash \rho$, $\mathrm{C}; \varrho \vdash e$, $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$, $\mathrm{T}, \Upsilon \vdash \rho$, $\mathrm{T}, \Upsilon \vdash e$, *and* $(\rho, e) \longmapsto (\rho, e')$ *then* $\mathrm{T}, \Upsilon \vdash e'$.

Proof.
- If $e = x^k$, $e' = v^j$, and $x^i{:}t = v^j \in \rho$ then the conclusion is a subderivation of $\mathrm{T}, \Upsilon \vdash \rho$.
- If $e = (\lambda x^i{:}t.e_1)^j$ and $e' = \langle \rho, \lambda x^i{:}t.e_1 \rangle^j$ then this follows immediately, since the premises line up exactly in the introduction and value rules, with the addition of $\mathrm{T}, \Upsilon \vdash \rho$ in the value rule which holds by assumption.
- If $e = (\mathtt{box}_t\, {v_t}^i)^j$ and $e' = \langle {v_t}^i{:}t \rangle^j$ then this follows immediately as in the lambda case.
- If $e = (e_1\ e_2)^j$ and $e' = (e_1'\ e_2)^j$ then we apply the induction hypothesis to $e_1 \longmapsto e_1'$ and we're done.
- If $e = (e_1\ e_2)^j$ and $e' = (e_1\ e_2')^j$ then the argument follows by the symmetric argument to the previous case.

– If $e = (\langle \rho', \lambda x^i{:}t.e'' \rangle^j\ v^k)^l$ and $e' = (\rho', x^i{:}t = v^k)(e'')^l$. By assumption, we have that $\mathrm{T}, \Upsilon \vdash e''$, so it suffices to show that $\mathrm{T}, \Upsilon \vdash (\rho', x^i{:}t = v^k)$. By assumption we have that $\mathrm{T}, \Upsilon \vdash \rho'$, so it suffices to show that $\underline{\mathrm{T}}(i,t) = \underline{\mathrm{T}}(k,t) \wedge \mathrm{T}, \Upsilon \vdash v^k$. We have that $\mathrm{T}, \Upsilon \vdash v^k$ by assumption. To show that $\underline{\mathrm{T}}(i,t) = \underline{\mathrm{T}}(k,t)$ we argue as follows.
By the assumption $\mathrm{C}; \varrho \vdash e''$ we have $(i{:}t \to \mathrm{lbl}(e''))^j \in \mathrm{C}(j)$ (1) and $\mathrm{C}(k) \subseteq \mathrm{C}(i')$ for all $(i'{:}t' \to n')^{j'} \in \mathrm{C}(j)$ (2). So by (1) and (2) we have that $\mathrm{C}(k) \subseteq \mathrm{C}(i)$. By Inhabitance we have that $s \in \mathrm{C}(k)$, so by Agreement we have that $k \overset{\mathrm{T},\Upsilon}{\simeq} i$, and hence we have that $\underline{\mathrm{T}}(i,t) = \underline{\mathrm{T}}(k,t)$ as required.
– If $e = (\mathtt{box}_t\, e_1)^j$, $e' = (\mathtt{box}_t\, e_2)^j$, and $\mathrm{lbl}(e_1) = i$ then:
  • If $j \in \Upsilon$ then let $k = \mathrm{lbl}(e_2)$:
    By induction we have that $\mathrm{T}, \Upsilon \vdash e_2$ so it suffices to show that $\mathrm{T}(j) = \mathrm{T}(k)$. We have by assumption that $\mathrm{C}; \varrho \vdash \rho$, $\mathrm{C}; \varrho \vdash e$, $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$, and by the reduction rule we have that $(\rho, e_1) \longmapsto (\rho, e_2)$ so by Lemma 9 we have that $i \overset{\mathrm{T},\Upsilon}{\simeq} k$, so we have that $\mathrm{T}(i) = \mathrm{T}(k)$, and we have $\mathrm{T}(i) = \mathrm{T}(j)$ by assumption, so we're done.
  • If $j \notin \Upsilon$ then we apply the induction hypothesis to $(\rho, e_1) \longmapsto (\rho, e_2)$ and all of the other premises still hold, so we're done.
– If $e = (\mathtt{unbox}\, e_1)^j$ and $e' = (\mathtt{unbox}\, e_2)^j$, then we apply the induction hypothesis to the derivation of $\mathrm{T}, \Upsilon \vdash e_1$ to get $\mathrm{T}, \Upsilon \vdash e_2$ and we're done.
– If $e = (\mathtt{unbox}\, \langle v^i{:}t \rangle^j)^k$ and $e' = v^i$ then by assumption and inversion (using either of the two box rules) we have $\mathrm{T}, \Upsilon \vdash v^i$ and we're done.
– If $e = \rho'(e_1)^i$ and $e' = \rho'(e_2)^i$ then by assumption $\mathrm{T}, \Upsilon \vdash \rho'$ and $\mathrm{T}, \Upsilon \vdash e_1$, and by the acceptable analysis assumption we have that $\mathrm{C}; \varrho \vdash \rho'$, $\mathrm{C}; \varrho \vdash e_1$. So by the induction hypothesis $\mathrm{T}, \Upsilon \vdash e_2$, and we're done.
– If $e = \rho'(v^i)^j$ and $e' = v^i$ then $\mathrm{T}, \Upsilon \vdash e'$ is a subderivation of $\mathrm{T}, \Upsilon \vdash e$.

Lemma 11 (Many step preservation of acceptability). *If* $\mathrm{C}; \varrho \vdash M$ *and* $\mathrm{C} \vdash M \downharpoonleft\downharpoonright (\mathrm{T}, \Upsilon)$ *and* $M \longmapsto^* M'$ *then* $\mathrm{C} \vdash M' \downharpoonleft\downharpoonright (\mathrm{T}, \Upsilon)$.

Proof. The proof is by induction on the length of the derivation, using Lemma 10.

Finally, we show coherence by observing that if two terms are related by reduction, then their images under the unboxing function induced by a good unboxing are related by the many step reduction relation.

Lemma 12 (Single step reduction coherence). *If* $\mathrm{C}; \varrho \vdash M$ *and* $\mathrm{C} \vdash M \downharpoonleft\downharpoonright (\mathrm{T}, \Upsilon)$ *and* $M \longmapsto M'$ *then* $\lfloor M \rfloor^{\mathrm{T}}_{\Upsilon} \longmapsto^* \lfloor M' \rfloor^{\mathrm{T}}_{\Upsilon}$.

First, a technical lemma about lifting reduction sequences into certain contexts and a lemma that the unboxing of a value has the traceability given by the traceability map.

Lemma 13 (Many step compositionality).

– *If* $(\rho, e_1) \longmapsto^* (\rho, e_1')$ *then* $(\rho, (e_1\ e_2)^i) \longmapsto^* (\rho, (e_1'\ e_2)^i)$

- *If* $(\rho, e_2) \longmapsto^* (\rho, e_2')$ *then* $(\rho, (v^j\ e_2)^i) \longmapsto^* (\rho, (v^j\ e_2')^i)$
- *If* $(\rho, e) \longmapsto^* (\rho, e')$ *then* $(\rho, (\mathtt{box}_t\ e)^i) \longmapsto^* (\rho, (\mathtt{box}_t\ e')^i)$
- *If* $(\rho, e) \longmapsto^* (\rho, e')$ *then* $(\rho, (\mathtt{unbox}\ e)^i) \longmapsto^* (\rho, (\mathtt{unbox}\ e')^i)$
- *If* $(\rho', e) \longmapsto^* (\rho', e')$ *then* $(\rho, \rho'(e)^i) \longmapsto^* (\rho, \rho'(e')^i)$

Proof. The proof is by an easy induction on the length of the reduction sequences.

Lemma 14 (Unboxed values traceability correctness). *If* $\mathrm{T}, \Upsilon \vdash {v_t}^i$ *and* $\lfloor {v_t}^i \rfloor^{\mathrm{T}}_{\Upsilon} = {v'_{t'}}^k$ *then* $t' = \underline{\mathrm{T}}(i, t)$.

Proof. By induction on ${v_t}^i$. The proof proceeds by cases:

- If $v_t = c$ then $t = \mathtt{b}$ and by definition of $\mathrm{T}, \Upsilon \vdash {v_t}^i$, $\mathrm{T}(i) \geq \mathtt{b}$.
- If $v_t = \langle \rho, \lambda x^j{:}t''.e \rangle$ then $t = \mathtt{r}$ and by definition of $\mathrm{T}, \Upsilon \vdash {v_t}^i$, $\mathrm{T}(i) \geq \mathtt{r}$. Also $\lfloor {v_t}^i \rfloor^{\mathrm{T}}_{\Upsilon}$ is a lambda value so $t' = \mathtt{r}$.
- If $v_t = \langle {v''_{t''}}^j{:}t'' \rangle$ then $t = \mathtt{r}$.
  - If $i \in \Upsilon$ then by the definition of $\mathrm{T}, \Upsilon \vdash {v_t}^i$, $\mathrm{T}(i) = \mathrm{T}(j)$. Also $\lfloor {v_t}^i \rfloor^{\mathrm{T}}_{\Upsilon} = \lfloor {v''_{t''}}^j \rfloor^{\mathrm{T}}_{\Upsilon} = {v'_{t'}}^k$. By induction then we have that $t' = \underline{\mathrm{T}}(j, t'')$. Since $\mathrm{T}(j)$ defined and equal to $\mathrm{T}(i)$ then $t' = \underline{\mathrm{T}}(i, t)$ as required.
  - If $i \notin \Upsilon$ then by the definition of $\mathrm{T}, \Upsilon \vdash {v_t}^i$, $\mathrm{T}(i) \geq \mathtt{r}$. Also $\lfloor {v_t}^i \rfloor^{\mathrm{T}}_{\Upsilon}$ is a box value so $t' = \mathtt{r}$.

Proof. The proof is by induction on the derivation of $M \longmapsto M'$, consider the cases for the last rule used to derive it:

- If $(\rho, x^k) \longmapsto (\rho, v^j)$ where $x^i{:}t = v^j \in \rho$ then by definition $\lfloor M \rfloor^{\mathrm{T}}_{\Upsilon} = (\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, x^k)$, $\lfloor M' \rfloor^{\mathrm{T}}_{\Upsilon} = (\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor v^j \rfloor^{\mathrm{T}}_{\Upsilon})$, and $x^i{:}\underline{\mathrm{T}}(i, t) = \lfloor v^j \rfloor^{\mathrm{T}}_{\Upsilon} \in \lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}$. Thus $\lfloor M \rfloor^{\mathrm{T}}_{\Upsilon} \longmapsto \lfloor M' \rfloor^{\mathrm{T}}_{\Upsilon}$ by the same rule.
- If $(\rho, (\lambda x^i{:}t.e_1)^j) \longmapsto (\rho, \langle \rho, \lambda x^i{:}t.e_1 \rangle^j)$, then the unboxings of the e and e' are of the same form, and the same reduction step applies.
- If $(\rho, (\mathtt{box}_t\ {v_t}^i)^j) \longmapsto (\rho, \langle {v_t}^i{:}t \rangle^j)$ then:
  - If $j \notin \Upsilon$ then:

    By the definition of unboxing we have that $\lfloor e \rfloor^{\mathrm{T}}_{\Upsilon} = (\mathtt{box}_{\underline{\mathrm{T}}(i,t)}\ {v'_{t'}}^k)^j$ where ${v'_{t'}}^k = \lfloor {v_t}^i \rfloor^{\mathrm{T}}_{\Upsilon}$. By Lemma 14 we have that $t' = \underline{\mathrm{T}}(i, t)$, and by definition of reduction $(\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, (\mathtt{box}_{\underline{\mathrm{T}}(i,t)}\ {v'_{\underline{\mathrm{T}}(i,t)}}^k)^j) \longmapsto (\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, \langle {v'_{\underline{\mathrm{T}}(i,t)}}^k{:}\underline{\mathrm{T}}(i, t) \rangle^j)$.
  - If $j \in \Upsilon$ then:

    By definition of unboxing

    $\lfloor e \rfloor^{\mathrm{T}}_{\Upsilon} = {v'_{t'}}^k$ where ${v'_{t'}}^k = \lfloor {v_t}^i \rfloor^{\mathrm{T}}_{\Upsilon}$

    By definition of reduction

    $(\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, {v'_{t'}}^k) \longmapsto^* (\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, {v'_{t'}}^k)$
- If $(\rho, (e_1\ e_2)^j) \longmapsto (\rho, (e_1'\ e_2)^j)$ then:

  By definition of $\mathrm{C}; \varrho \vdash M$ and $\mathrm{C} \vdash M \downharpoonleft\downharpoonright (\mathrm{T}, \Upsilon)$ we have that $\mathrm{C}; \varrho \vdash \rho$, $\mathrm{C}; \varrho \vdash e_1$ and $\mathrm{T}, \Upsilon \vdash e_1$. Hence we have that $\mathrm{C}; \varrho \vdash (\rho, e_1)$ and $\mathrm{C} \vdash (\rho, e_1) \downharpoonleft\downharpoonright (\mathrm{T}, \Upsilon)$, and hence by induction we have that $(\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e_1 \rfloor^{\mathrm{T}}_{\Upsilon}) \longmapsto^* (\lfloor \rho \rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e_1' \rfloor^{\mathrm{T}}_{\Upsilon})$.

By Lemma 13
$(\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, (\lfloor e_1\rfloor^{\mathrm{T}}_{\Upsilon}\ \lfloor e_2\rfloor^{\mathrm{T}}_{\Upsilon})^j) \longmapsto^* (\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, (\lfloor e_1'\rfloor^{\mathrm{T}}_{\Upsilon}\ \lfloor e_2\rfloor^{\mathrm{T}}_{\Upsilon})^j)$
By definition of unboxing
$(\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor(e_1\ e_2)^j\rfloor^{\mathrm{T}}_{\Upsilon}) \longmapsto^* (\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor(e_1'\ e_2)^j\rfloor^{\mathrm{T}}_{\Upsilon})$

– If $(\rho, (e_1\ e_2)^j) \longmapsto (\rho, (e_1\ e_2')^j)$ then the argument follows by the symmetric argument to the previous case.

– If $(\rho, (\langle\rho', \lambda x^i{:}t.e''\rangle^j\ v_t{}^k)^l) \longmapsto (\rho, (\rho', x^i{:}t = v_t{}^k)(e'')^l)$ then:
By definition of unboxing we have that
$\lfloor e\rfloor^{\mathrm{T}}_{\Upsilon} = (\langle\lfloor\rho'\rfloor^{\mathrm{T}}_{\Upsilon}, \lambda x^i{:}\underline{\mathrm{T}}(i,t).\lfloor e''\rfloor^{\mathrm{T}}_{\Upsilon}\rangle^j\ \lfloor v_t{}^k\rfloor^{\mathrm{T}}_{\Upsilon})^l$
By definition of unboxing we have that
$\lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon} = (\lfloor\rho'\rfloor^{\mathrm{T}}_{\Upsilon}, x^i{:}\underline{\mathrm{T}}(i,t) = \lfloor v_t{}^k\rfloor^{\mathrm{T}}_{\Upsilon})(\lfloor e''\rfloor^{\mathrm{T}}_{\Upsilon})^l$
By Lemma 14 we have that $\lfloor v_t{}^k\rfloor^{\mathrm{T}}_{\Upsilon} = v'_{\underline{\mathrm{T}}(k,t)}{}^m$, so it suffices to show that $\underline{\mathrm{T}}(k,t) = \underline{\mathrm{T}}(i,t)$. For this it suffices to show that $k \overset{\mathrm{T},\Upsilon}{\simeq} i$. By Inhabitance, we have that $\mathrm{C}(k)$ is inhabited, and by the acceptable analysis rules we have that $\mathrm{C}(k) \subseteq \mathrm{C}(i)$, so by Agreement we have $k \overset{\mathrm{T},\Upsilon}{\simeq} i$ and we're done.

– If $(\rho, (\mathtt{box}_t\, e)^j) \longmapsto (\rho, (\mathtt{box}_t\, e')^j)$ then:
- If $j \in \Upsilon$ then:
  By definition of unboxing
  $\lfloor(\mathtt{box}_t\, e)^j\rfloor^{\mathrm{T}}_{\Upsilon} = \lfloor e\rfloor^{\mathrm{T}}_{\Upsilon}$
  By definition of unboxing
  $\lfloor(\mathtt{box}_t\, e')^j\rfloor^{\mathrm{T}}_{\Upsilon} = \lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon}$
  By induction
  $(\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e\rfloor^{\mathrm{T}}_{\Upsilon}) \longmapsto^* (\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon})$
- If $j \notin \Upsilon$ then let $i = \mathrm{lbl}(e), k = \mathrm{lbl}(e')$:
  By definition of unboxing we have that $\lfloor(\mathtt{box}_t\, e)^j\rfloor^{\mathrm{T}}_{\Upsilon} = (\mathtt{box}_{\underline{\mathrm{T}}(i,t)}\lfloor e\rfloor^{\mathrm{T}}_{\Upsilon})^j$, and $\lfloor(\mathtt{box}_t\, e')^j\rfloor^{\mathrm{T}}_{\Upsilon} = (\mathtt{box}_{\underline{\mathrm{T}}(k,t)}\lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon})^j$. By the induction hypothesis we have that $(\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e\rfloor^{\mathrm{T}}_{\Upsilon}) \longmapsto^* (\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon})$, so by Lemma 13 we have that:
  $(\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, (\mathtt{box}_{\underline{\mathrm{T}}(i,t)}\lfloor e\rfloor^{\mathrm{T}}_{\Upsilon})^j) \longmapsto^* (\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, (\mathtt{box}_{\underline{\mathrm{T}}(i,t)}\lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon})^j)$.
  We must show that:
  $(\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, (\mathtt{box}_{\underline{\mathrm{T}}(i,t)}\lfloor e\rfloor^{\mathrm{T}}_{\Upsilon})^j) \longmapsto^* (\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, (\mathtt{box}_{\underline{\mathrm{T}}(k,t)}\lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon})^j)$
  So it suffices to show that $\underline{\mathrm{T}}(i,t) = \underline{\mathrm{T}}(k,t)$, and hence it suffices to show that $i \overset{\mathrm{T},\Upsilon}{\simeq} k$. This follows by Lemma 9.

– If $(\rho, (\mathtt{unbox}\, e)^j) \longmapsto (\rho, (\mathtt{unbox}\, e')^j)$ then let $i = \mathrm{lbl}(e)$ and $i' = \mathrm{lbl}(e')$. By Lemma 9 we have that $i \overset{\mathrm{T},\Upsilon}{\simeq} i'$, and hence $i \in \Upsilon$ iff $i' \in \Upsilon$.
- If $i, i' \in \Upsilon$ then:
  By definition of unboxing
  $\lfloor(\mathtt{unbox}\, e)^j\rfloor^{\mathrm{T}}_{\Upsilon} = \lfloor e\rfloor^{\mathrm{T}}_{\Upsilon}$
  By definition of unboxing
  $\lfloor(\mathtt{unbox}\, e')^j\rfloor^{\mathrm{T}}_{\Upsilon} = \lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon}$
  By induction
  $(\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e\rfloor^{\mathrm{T}}_{\Upsilon}) \longmapsto^* (\lfloor\rho\rfloor^{\mathrm{T}}_{\Upsilon}, \lfloor e'\rfloor^{\mathrm{T}}_{\Upsilon})$

- If $i, i' \notin \Upsilon$ then:
  By definition of unboxing
  $\lfloor (\mathtt{unbox}\, e)^j \rfloor_\Upsilon^{\mathrm{T}} = (\mathtt{unbox}\, \lfloor e \rfloor_\Upsilon^{\mathrm{T}})^j$
  By definition of unboxing
  $\lfloor (\mathtt{unbox}\, e')^j \rfloor_\Upsilon^{\mathrm{T}} = (\mathtt{unbox}\, \lfloor e' \rfloor_\Upsilon^{\mathrm{T}})^j$
  By induction
  $(\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor e \rfloor_\Upsilon^{\mathrm{T}}) \longmapsto^* (\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor e' \rfloor_\Upsilon^{\mathrm{T}})$
  By Lemma 13
  $(\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor (\mathtt{unbox}\, e)^j \rfloor_\Upsilon^{\mathrm{T}}) \longmapsto^* (\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor (\mathtt{unbox}\, e')^j \rfloor_\Upsilon^{\mathrm{T}})$

– If $(\rho, (\mathtt{unbox}\, \langle v^i{:}t \rangle^j)^k) \longmapsto (\rho, v^i)$ then:
  - If $j \in \Upsilon$ then:
    By definition of unboxing
    $\lfloor (\mathtt{unbox}\, \langle v^i{:}t \rangle^j)^k \rfloor_\Upsilon^{\mathrm{T}} = \lfloor \langle v^i{:}t \rangle^j \rfloor_\Upsilon^{\mathrm{T}} = \lfloor v^i \rfloor_\Upsilon^{\mathrm{T}}$
    So in zero steps
    $(\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor (\mathtt{unbox}\, \langle v^i{:}t \rangle^j)^k \rfloor_\Upsilon^{\mathrm{T}}) \longmapsto^* (\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor v^i \rfloor_\Upsilon^{\mathrm{T}})$
  - If $j \notin \Upsilon$ then:
    By definition of unboxing
    $\lfloor (\mathtt{unbox}\, \langle v^i{:}t \rangle^j)^k \rfloor_\Upsilon^{\mathrm{T}} = (\mathtt{unbox}\, \lfloor \langle v^i{:}t \rangle^j \rfloor_\Upsilon^{\mathrm{T}})^k$
    By definition of unboxing
    $(\mathtt{unbox}\, \lfloor \langle v^i{:}t \rangle^j \rfloor_\Upsilon^{\mathrm{T}})^k = (\mathtt{unbox}\, \langle \lfloor v^i \rfloor_\Upsilon^{\mathrm{T}} {:} \underline{\mathrm{T}}(i,t) \rangle^j)^k$
    By definition of reduction
    $(\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, (\mathtt{unbox}\, \langle \lfloor v^i \rfloor_\Upsilon^{\mathrm{T}} {:} \underline{\mathrm{T}}(i,t) \rangle^j)^k) \longmapsto (\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor v^i \rfloor_\Upsilon^{\mathrm{T}})$

– If $(\rho, \rho'(e_1)^i) \longmapsto (\rho, \rho'(e_2)^i)$ then:
  By definition of unboxing
  $\lfloor \rho'(e_1)^i \rfloor_\Upsilon^{\mathrm{T}} = \lfloor \rho' \rfloor_\Upsilon^{\mathrm{T}} (\lfloor e_1 \rfloor_\Upsilon^{\mathrm{T}})^i$
  By definition of unboxing
  $\lfloor \rho'(e_2)^i \rfloor_\Upsilon^{\mathrm{T}} = \lfloor \rho' \rfloor_\Upsilon^{\mathrm{T}} (\lfloor e_2 \rfloor_\Upsilon^{\mathrm{T}})^i$
  By induction
  $(\lfloor \rho' \rfloor_\Upsilon^{\mathrm{T}}, \lfloor e_1 \rfloor_\Upsilon^{\mathrm{T}}) \longmapsto^* (\lfloor \rho' \rfloor_\Upsilon^{\mathrm{T}}, \lfloor e_2 \rfloor_\Upsilon^{\mathrm{T}})$
  By Lemma 13
  $(\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor \rho' \rfloor_\Upsilon^{\mathrm{T}} (\lfloor e_1 \rfloor_\Upsilon^{\mathrm{T}})^i) \longmapsto^* (\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor \rho' \rfloor_\Upsilon^{\mathrm{T}} (\lfloor e_2 \rfloor_\Upsilon^{\mathrm{T}})^i)$

– If $(\rho, \rho'(v^i)^j) \longmapsto (\rho, v^i)$ then:
  By definition of unboxing
  $\lfloor \rho'(v^i)^j \rfloor_\Upsilon^{\mathrm{T}} = \lfloor \rho' \rfloor_\Upsilon^{\mathrm{T}} (\lfloor v^i \rfloor_\Upsilon^{\mathrm{T}})^j$
  Unboxed value is a value, so by reduction rules
  $(\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor \rho' \rfloor_\Upsilon^{\mathrm{T}} (\lfloor v^i \rfloor_\Upsilon^{\mathrm{T}})^j) \longmapsto (\lfloor \rho \rfloor_\Upsilon^{\mathrm{T}}, \lfloor v^i \rfloor_\Upsilon^{\mathrm{T}})$

Proof of Theorem 2

Proof. – By induction on reduction derivations, using Lemma 12.

1. If $M \longmapsto^* (\rho, v^i)$ in zero steps, then the result follows immediately.
2. If $M \longmapsto^* (\rho, v^i)$ in n steps, then by definition, $M \longmapsto M'$ and $M' \longmapsto^* (\rho, v^i)$ in $n-1$ steps.
   By Lemma 12
   $\downarrow M \downarrow^{\mathrm{T}}_{\Upsilon} \longmapsto \downarrow M' \downarrow^{\mathrm{T}}_{\Upsilon}$
   By Lemma 2
   $\mathrm{C}; \varrho \vdash M'$
   By Lemma 11
   $\mathrm{C} \vdash M' \downarrow\downarrow (\mathrm{T}, \Upsilon)$
   By induction
   $\downarrow M' \downarrow^{\mathrm{T}}_{\Upsilon} \longmapsto^* (\downarrow \rho \downarrow^{\mathrm{T}}_{\Upsilon}, \downarrow v^i \downarrow^{\mathrm{T}}_{\Upsilon})$
   By the defininition of many step reduction
   $\downarrow M \downarrow^{\mathrm{T}}_{\Upsilon} \longmapsto^* (\downarrow \rho \downarrow^{\mathrm{T}}_{\Upsilon}, \downarrow v^i \downarrow^{\mathrm{T}}_{\Upsilon})$

– In the operational semantics, there are six leaf reductions. Two of them take expression forms to value forms, but otherwise leave the term unchanged. One of the them takes unbox of box of a value to that value. One of them takes a frame of a value to that value. Thus if we measure a term by adding its size, number of lambda expressions, and number of box expressions, then this metric strictly decreases for these three leaf reductions. Therefore, in any infinite reduction sequence, there must be an infinite number of steps whose leaf reduction is a variable reduction or an application beta reduction. Then observe in the proof of Lemma 12 that the unboxing of a variable redex or of an application beta redex will always take a step, and that Lemma 13 preserves this. Thus the unboxing will also take an infinite number of steps.

A.4 Proofs for Section 5

Lemma 15 (Initial unboxing cache consistency). *Given any acceptable analysis* (C, ϱ)*, the* $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ *quadruple computed by the algorithm in this section induces an unboxing* (T, Υ) *such that* $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$ *(where* $\mathrm{T} = \mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}$*).*

Proof. By the definition of connected components, $i \in \mathcal{EC}(i)$. By the construction of the initial flow graph, it is also the case that every $s \in \mathrm{C}(i)$ is in the same connected component as i, and hence has the same definition in $\mathcal{T}_{\mathcal{EC}}$ (since the domain of $\mathcal{T}$ is equivalence classes of connected components), and consequently has the same definition in the induced traceability map T. Since the induced unboxing set Υ is empty, we have that $i \overset{\mathrm{T},\Upsilon}{\simeq} \mathrm{lbl}(s)$ for every $s \in \mathrm{C}(i)$ and hence we have that $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$.

Lemma 16 (Traceability map consistency). *Given an acceptable analysis* (C, ϱ)*, the* $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ *quadruple computed by the algorithm in this section induces an unboxing* (T, Υ) *(where* $\mathrm{T} = \mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}$*) such that for every term* e *in the program with* $\mathrm{lbl}(e) = i$*, if* e *is an introduction form* $(\lambda x^j{:}t.e)^i$*,* $(\mathtt{box}_t\, e)^i$*, or* c^i*,* $\mathrm{T}(i) \geq itr(e)$*.*

Proof. By the definition of an acceptable analysis in Figure 4 $\mathrm{C}(i)$ contains a shape s of the same form as e, and hence with the same immediate traceability as e. If s and e are not boxes then $\mathcal{T}(i)$ includes $itr(s)$ and so $\mathcal{T}(i)$ must be at least $itr(s)$. If s and e are boxes, then $\mathcal{EC}(i)$ is in $\mathcal{CN}$, and hence by definition of the induced traceability map T, $\mathrm{T}(i) = \mathcal{T}(\mathcal{EC}(i)) \sqcup \mathtt{r} \geq \mathtt{r}$.

Proof of Lemma 3

Proof. $\mathrm{C} \vdash (\mathrm{T}, \Upsilon)$ follows by Lemma 15. It suffices to show that $\mathrm{T}, \Upsilon \vdash e$. The proof is by induction on terms. Since the induced Υ is empty, it suffices to consider only the rules that apply when the top level label is not in Υ.

Note that for any two labels i and j that are in the same connected component, T must agree on i and j, since they are in the same equivalence class.

- The variable case follows immediately.
- For the lambda case, $\mathrm{T}(j) \geq \mathtt{r}$ follows from Lemma 16, and the consistency of the body follows by induction.
- In the application case, the consistency of the two sub-terms follows by induction.
- For the box introduction case, $\mathcal{EC}(i) \in \mathcal{CN}$, and hence $\mathrm{T}(i) \geq r$ by the definition of $\mathcal{T}^{\mathcal{CN}}$.
- For the unbox case, the consistency of the sub-term follows by induction.
- For the constant case, $\mathrm{T}(i) \geq b$ follows from Lemma 16.

Observation 1 (Connected components). *For any two box introductions in the program* $(\mathtt{box}_t\ e)^i$ *and* $(\mathtt{box}_{t'}\ e')^i$, $\mathrm{lbl}(e)$ *and* $\mathrm{lbl}(e')$ *are in the same connected component in the induced flow graph. To see this note that the construction of the graph includes edges between* $\mathrm{lbl}(e)$ *and* $\mathrm{lbl}(e'')$ *for every* $(\mathtt{box}_t\ j)^i \in \mathrm{C}(i)$, *and similarly for* $(\mathtt{box}_{t'}\ e')^i$. *By inspection of the acceptable analysis rules,* $\mathrm{C}(i)$ *is inhabited, so* $\mathrm{lbl}(e)$ *and* $\mathrm{lbl}(e')$ *must both be in the same connected component as* j, *and hence in the same connected component as each other.*

Lemma 17 (Extended Traceability Map Monotonicity). *Throughout the unboxing phase, for any unboxing step starting with* $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ *and resulting in* $(\mathcal{T}', \Upsilon', \mathcal{CN}', \mathcal{EC}')$*:*

- $\forall i : \mathcal{T}(\mathcal{EC}(i)) \leq \mathcal{T}'(\mathcal{EC}'(i))$.
- $\forall i : \mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}(i) < \top \implies \mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'}(i) < \top$.

Proof. $-\ \forall i : \mathcal{T}(\mathcal{EC}(i)) \leq \mathcal{T}'(\mathcal{EC}'(i))$. Whenever the algorithm combines two equivalence classes, it sets the new extended traceability map for the combined equivalence class to be the least upper bound of the extended traceabilities of the two original equivalence classes. When the algorithm discards an equivalence class from the candidate set, it sets $\mathcal{T}$ to map the equivalence class the least upper bound of its old value and $\mathtt{r}$, and hence can only move the extended traceability up the lattice.

– $\forall i : \mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(i) < \top \implies \mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'}(i) < \top$. In any step, we only change the traceability map for the equivalence classes associated with $\text{lbl}(e)$ and i.
  - In the case that we discard the equivalence class, we have that $\mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(i) = \mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathtt{r}$, since $\mathcal{EC}(i) \in \mathcal{CN}$, and $\mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}}(i) = \mathcal{T}'_{\mathcal{EC}}(i)$ since $\mathcal{EC}(i)$ is not in $\mathcal{CN}'$. By definition, $\mathcal{T}'_{\mathcal{EC}}(i) = \mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathtt{r}$, so we have that $\mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(i) = \mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}}(i)$. Since $\mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(i) < \top$, we have that $\mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}}(i) < \top$
  - In the case that we choose to unbox, we have $\mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(\text{lbl}(e)) < \top$. By definition, we have that $\mathcal{T}'_{\mathcal{EC}'}(i) = \mathcal{T}'_{\mathcal{EC}'}(\text{lbl}(e)) = \mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathcal{T}_{\mathcal{EC}}(\text{lbl}(e))$, so it suffices to show that this is less than $\top$. Note that $\mathcal{EC}'(\text{lbl}(e)) = \mathcal{EC}'(i)$.
    * if $\mathcal{EC}'(\text{lbl}(e)) \notin \mathcal{CN}$ then $\mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'}(i) = \mathcal{T}'_{\mathcal{EC}'}(i) = \mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathcal{T}_{\mathcal{EC}}(\text{lbl}(e)) < \top$ by assumption (since we do not meet the unboxing criteria otherwise).
    * if $\mathcal{EC}'(\text{lbl}(e)) \in \mathcal{CN}$ then:

$$\begin{aligned}
\mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'}(i) &= \mathcal{T}'_{\mathcal{EC}'}(i) \sqcup \mathtt{r} \\
&= (\mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathcal{T}_{\mathcal{EC}}(\text{lbl}(e))) \sqcup \mathtt{r} \\
&= \mathcal{T}_{\mathcal{EC}}(i) \sqcup (\mathcal{T}_{\mathcal{EC}}(\text{lbl}(e)) \sqcup \mathtt{r}) \\
&= \mathcal{T}_{\mathcal{EC}}(i) \sqcup (\mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(\text{lbl}(e)) \\
&< \top \text{ by assumption}
\end{aligned}$$

Lemma 18 (Traceability Map Monotonicity). *Throughout the unboxing phase, for any unboxing step starting with a quadruple $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ and resulting in a quadruple $(\mathcal{T}', \Upsilon', \mathcal{CN}', \mathcal{EC}')$, if $\text{T}(i)$ is defined, then $T'(i)$ is defined, where $T' = \mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'}$.*

Proof. By definition, if $\text{T}(i)$ is defined then $\mathcal{T}_{\mathcal{EC}}^{\mathcal{CN}}(i) < \top$, and hence $\mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'} < \top$ by lemma 17. In all other cases, we have that $\text{T}'(i)$ is defined. If $\mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'}(i) = \bot$ then $\text{T}'(i) = \mathtt{r}$ by definition. If $\mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'}(i) = t$ then $\text{T}'(i) = t$ by definition.

Lemma 19 (Incremental unboxing cache consistency). *After every step of the algorithm starting with $(\mathcal{T}, \Upsilon, \mathcal{CN}, \mathcal{EC})$ resulting in a new $(\mathcal{T}', \Upsilon', \mathcal{CN}', \mathcal{EC}')$ quadruple, we have that $\text{C} \vdash (\text{T}', \Upsilon')$ (where $\text{T}' = \mathcal{T}'^{\mathcal{CN}'}_{\mathcal{EC}'}$)*

Proof. Anything initially in the same equivalence class remains in the same equivalence class, so the argument from Lemma 15 continues to hold for T'. For Υ', note that the algorithm only grows Υ, and every step adds an entire equivalence class to Υ. Consequently, if $s \in \text{C}(i)$ then $\text{lbl}(s)$ and i are in the same equivalence class and hence $\text{lbl}(s) \overset{U}{\simeq} i$.

Proof of Lemma 4

Proof. Note that the properties shown in Lemma 16 continue to hold, since anything initially in the same equivalence class remains in the same equivalence class.

By Lemma 19 we have that $\mathrm{C} \vdash (\mathrm{T}', \Upsilon')$ so it suffices to show that the $\mathrm{T}', \Upsilon', e \vdash$ continues to hold.

If the unboxing step is not taken, (that is, we discard the candidate), then it is straightforward to verify that the result is still a good unboxing since the induced traceability map remains unchanged (${\mathcal{T}'}^{\mathcal{CN}'}_{\mathcal{EC}'}(i) = \mathcal{T}'_{\mathcal{EC}'}(i) = \mathcal{T}_{\mathcal{EC}}(i) \sqcup \mathtt{r} = \mathcal{T}^{\mathcal{CN}}_{\mathcal{EC}}(i)$).

For the case that the unboxing step is taken, we consider the unboxing (T', Υ') induced by the quadruple $(\mathcal{T}', \Upsilon', \mathcal{CN}', \mathcal{EC}')$ resulting from a single step. The cases are as follows (in order of appearance in Figure 6).

- For the variable case, the argument from Lemma 3 continues to hold.
- For the lambda case, the argument from Lemma 3 continues to hold.
- For the application case, the argument from Lemma 3 continues to hold.
- For the box introduction case $(\mathtt{box}_t\, e)^i$, if $i \in \Upsilon$ then for some $(\mathtt{box}_{t'}\, e')^i$ and $(\mathcal{T}^0, \Upsilon^0, \mathcal{CN}^0, \mathcal{EC}^0)$ we had that $\mathcal{EC}^0(i) \in \mathcal{CN}^0$, and that $\mathcal{T}^0_{\mathcal{EC}^0}(i) \sqcup {\mathcal{T}^0}^{\mathcal{CN}^0}_{\mathcal{EC}^0_0}(\mathrm{lbl}(e)) < \top$, and hence we unboxed i as described above. Consequently, in the resulting unboxing $(\mathcal{T}^1, \Upsilon^1, \mathcal{CN}^1, \mathcal{EC}^1)$ we had that $\mathcal{EC}^1(i) = \mathcal{EC}^1(\mathrm{lbl}(e'))$. As observed in Observation 1 above, we have that $\mathrm{lbl}(e)$ and $\mathrm{lbl}(e')$ are in the same connected component, and hence in the same equivalence class, and hence $\mathcal{EC}^1(i) = \mathcal{EC}^1(\mathrm{lbl}(e))$. Since $\mathcal{T}^0_{\mathcal{EC}^0}(i) \sqcup {\mathcal{T}^0}^{\mathcal{CN}^0}_{\mathcal{EC}^0_0}(\mathrm{lbl}(e)) < \top$ we have that ${\mathcal{T}^0}^{\mathcal{CN}^0}_{\mathcal{EC}^0_0}(\mathrm{lbl}(e)) < \top$, and hence we have that $\mathrm{T}(\mathrm{lbl}(e))$ is defined. By Lemma 18 then, we have that $\mathrm{T}^1(\mathrm{lbl}(e)$ is defined, and hence also $\mathrm{T}^1(i)$ as well. The algorithm grows the equivalence classes monotonically, so we have that $\mathcal{EC}'(i) = \mathcal{EC}'(\mathrm{lbl}(e))$, and so we have that $\mathrm{T}'(i) = \mathrm{T}'(\mathrm{lbl}(e))$.
- For the box introduction case, if $i \notin \Upsilon$, the argument from Lemma 3 continues to hold.
- For the unbox case, if $\mathrm{lbl}(e) \in \Upsilon$, then the second quantified premise follows by Lemma 16. For the first quantified premise, if $\mathrm{lbl}(e) \in \Upsilon$ then for some $(\mathtt{box}_{t'}\, e')^j$ in the program and $(\mathcal{T}^0, \Upsilon^0, \mathcal{CN}^0, \mathcal{EC}^0)$ such that $\mathrm{lbl}(e) \in \mathcal{EC}^0(j)$ we had that $\mathcal{EC}^0(j) \in \mathcal{CN}^0$, and that $\mathcal{T}^0_{\mathcal{EC}^0}(i) \sqcup {\mathcal{T}^0}^{\mathcal{CN}^0}_{\mathcal{EC}^0_0}(\mathrm{lbl}(e)) < \top$, and hence we unboxed j as described above. In particular, we have that $\Upsilon^1 = \Upsilon^0 \cup \mathcal{EC}^0(j)$, and since Υ grows monotontically, we have $\mathcal{EC}^0(j) \subseteq \Upsilon$, and hence $\mathcal{EC}^0(\mathrm{lbl}(e)) \subseteq \Upsilon$. By the construction of the induced flow graph, we have edges between $\mathrm{lbl}(s)$ and $\mathrm{lbl}(e)$ for every $s \in \mathrm{C}(\mathrm{lbl}(e))$, hence every shape s in $\mathrm{C}(\mathrm{lbl}(e))$ is in the same connected component as $\mathrm{lbl}(e)$, and hence $\mathrm{lbl}(s) \in \mathcal{EC}^0(\mathrm{lbl}(e))$ and hence $\mathrm{lbl}(s) \in \Upsilon$. This satisfies the second quantified premise.
- For the unbox case, if $\mathrm{lbl}(e) \notin \Upsilon$, the argument from Lemma 3 continues to hold.
- In the constant case, the argument from Lemma 3 continues to hold.

Proof of Theroem 3

Proof. Follows immediately by Lemmas 3 and 4.