

Reference Counting with Frame Limited Reuse

Microsoft Technical Report, MSR-TR-2021-30, Nov 20, 2021 (v1).

Anton Lorenzen
University of Bonn, Germany
anton.lorenzen@uni-bonn.de

Daan Leijen
Microsoft Research, USA
daan@microsoft.com

Abstract

The recently introduced *Perceus* algorithm can automatically insert reference count instructions such that the resulting (cycle-free) program is *garbage free*: objects are freed at the very moment they can no longer be referenced. An important extension of *Perceus* is *reuse analysis*. This optimization pairs objects of known size with fresh allocations of the same size and tries to reuse the object in-place at runtime if it happens to be unique. Current implementations of reuse analysis are fragile with respect to small program transformations, or can cause an arbitrary increase in the peak heap usage. We present a novel *drop-guided* reuse algorithm that is simpler and more robust than previous approaches. Moreover, we give a novel formalization of the linear resource calculus where we can precisely characterize *garbage-free* and *frame-limited* evaluations. On each function call, a frame-limited evaluation may hold on to memory longer if the size is bounded by a constant factor. Using this framework we show that our drop-guided reuse is frame-limited and find that an implementation of our new reuse approach in Koka can provide significant speedups.

1 Introduction

Reference counting [8] is a technique for automatic memory management where each allocated object stores the number of references that point to it. Reinking, Xie, de Moura, and Leijen [33] describe the *Perceus* algorithm for automatically inserting reference count instructions such that the resulting (cycle-free) program is *garbage free*: objects are freed at the very moment they can no longer be referenced. Even though the *Perceus* algorithm itself needs an internal calculus with explicit control flow, the authors apply this work in the context of the Koka language [21] which supports full algebraic effect handlers. These can be used to define various kinds of implicit control flow, including features like exceptions, `async/await`, and backtracking [20, 31, 32, 43].

An important extension of *Perceus* is *reuse analysis*. This optimization pairs objects of known size with fresh allocations of the same size and tries to reuse the object in-place at runtime if it happens to be unique. This was first described in earlier work by Ullrich and de Moura [38] in the context of the Lean theorem prover [27]. Unfortunately, the published algorithms for reuse analysis all have various weaknesses; for example, they are fragile with respect to small program transformations, where inlining or rearranging expressions can cause reuse analysis to fail unexpectedly.

Moreover, while *Perceus* itself is garbage-free, reuse analysis does not have this property. By construction, it holds on to memory that is to be reused later, which can lead to an increased peak memory usage. We find the maximum increase is not just a constant factor but can be much larger, and is generally not *safe for space* [1, 30].

In this work we improve upon this with a novel reuse algorithm and a formal framework for reasoning about heap bounds that can be applied to general transformations (including reuse analysis). In particular:

- We define a new approach to reuse called *drop guided reuse* (Section 2.3). In contrast to earlier techniques we perform the reuse *after* *Perceus* has inserted reference count instructions. This both simplifies the analysis and makes it more robust with respect to small program transformations. We formalize drop-guided reuse generally in the form of declarative derivation rules where we can discuss clearly the various choices an algorithm can make (Section 3.4).
- We illustrate this with a practical example in Section 2.4 where we see significantly better reuse for red-black tree balanced insertion in comparison with the previous algorithms. In combination with the tail-recursion-modulo-cons (TRMC) optimization, we show how for red-black tree insertion we get a highly optimized code path reusing a (unique) tree in-place while using minimal stack space. Our straightforward and purely functional implementation is about 19% faster than the manually optimized in-place mutating red-black tree implementation in the C++ STL library (`std::map`) (Section 4).
- We reformulate the original linear resource calculus λ^1 [33] in a normalized form (λ^{1n}) where we can reason precisely about reuse. The declarative derivation rules for the λ^{1n} calculus are non-deterministic and can derive many programs with reference counting instructions that are all correct but may differ in their memory consumption. In the normalized reformulation, it now is possible to add a single logical side condition (\star) to the `LET` rule that captures various important variants concisely (Section 3.3):
 - If the (\star) condition is unrestricted, the resulting programs are *sound* – that is, the reference counting is correct and the final heap contains no garbage.
 - By restricting (\star) in a particular way, we can show that all derived programs are *garbage-free* where at every allocating evaluation step there is no garbage (and all programs resulting from the *Perceus* algorithm fall in this class).

– Finally, we can weaken the (★) condition of the garbage-free system to also allow derivations that we call *frame-limited*, where every function call uses at most a constant factor c more memory.

- Transformations like reuse and borrowing no longer have the garbage-free property as they hold on to some memory for a bit longer (the cell to reuse, or the data that is borrowed). With the new formalization, we can now show that some of these transformations are still *frame-limited*, and we prove that our new *drop-guided* reuse analysis is frame-limited (Section 3.4). In contrast, we show that some previous reuse algorithms and unrestricted borrow inference [38] are *not* frame-limited transformations (and we argue such transformations should therefore be avoided in practice).
- Building upon drop-guided reuse analysis, we discuss new applications of *FBIP*: Functional But In-Place algorithms [33]. In Section 4.2.1 we show how we can use visitor data types (as a derivative of the original data type) to transform functions to become tail-recursive while reusing the traversed data type in place. We demonstrate this for the red-black tree balanced insertion, and also apply this to the parallel `binarytrees` benchmark, and show that with drop-guided reuse our purely functional implementation is competitive with top performing implementations in other systems.

2 Overview

We start with a brief introduction to Koka [19–21, 43] – a strongly typed functional language with effect handlers which tracks (side) effects in the type of every function. For example, we can define a squaring function as:

```
fun square( x : int ) : total int
  x * x
```

Here we see two types in the result: the effect type `total` and the result type `int`. The `total` type signifies that the function can be modeled semantically as a mathematically *total* function, which always terminates without raising an exception (or having any other observable side effect). Effectful functions get more interesting effect types, like:

```
fun println( s : string ) : console ()
fun divide( x : int, y : int ) : exn int
```

where `println` has a `console` effect and `divide` may raise an exception (`exn`) when dividing by zero. It is beyond the scope of this paper to go into full detail, but a novel feature of Koka is that it supports typed algebraic effect handlers which can define new effects like `async/await`, iterators, or co-routines without needing to extend the language

Koka uses algebraic data types extensively. For example, we can define a polymorphic list of elements of type `a` as:

```
type list(a)
  Cons( head : a, tail : list(a) )
  Nil
```

We can match on a list to define a polymorphic `map` function that applies a function `f` to each element of a list `xs`:

```
fun map( xs : list(a), f : a -> e b ) : e list(b)
  match xs
  Cons(x,xx) -> Cons(f(x), map(xx,f))
  Nil        -> Nil
```

Here we transform the list of generic elements of type `a` to a list of generic elements of type `b`. Since `map` itself has no intrinsic effect, the overall effect of `map` is polymorphic, and equals the effect `e` of the function `f` as it is applied to every element. The `map` function demonstrates many interesting aspects of reference counting and we use it as a running example in the following sections.

2.1 Perceus

In this section we give a short recap of the original Perceus algorithm and reuse analysis. By starting from a language with strong static guarantees, Perceus can insert optimized reference count instructions. Note though it still needs separate mechanisms to address cyclic data and mitigate the impact of thread shared reference counts – we refer to the Perceus paper for a in-depth discussion of these [33].

The main attribute that sets Perceus apart is that it is *garbage-free*: for a cycle-free program, an object is freed as soon as no more references remain. This is illustrated with the following function:

```
fun main()
  val xs = list(1,1000000) // allocate a large list
  val ys = map(xs,inc)     // increment each element
  println(ys)
```

Many reference count systems would drop the references to `xs` and `ys` based on the scope as:

```
fun main()
  val xs = list(1,1000000)
  val ys = map(xs,inc)
  println(ys)
  drop(xs)
  drop(ys)
```

where we use a gray background for generated operations. The `drop(xs)` operation decrements the reference count of an object and, if it drops to zero, recursively drops all children of the object and frees its memory. These “scoped lifetime” reference counts are for example used by a C++ `shared_ptr<T>` (calling the destructor at the end of the scope), Rust’s `Rc<T>` (using the `Drop` trait), and Nim (using a `finally` block to call `destroy`) [44]. It is not required by the semantics, but Swift typically emits code like this as well [15].

Implementing reference counting this way is straightforward and integrates well with exception handling where the drop operations are performed as part of stack unwinding. But from a performance perspective, the technique is not always optimal: in the previous example, the large list `xs` is retained in memory while a new list `ys` is built. Moreover, at the end of the scope, a long cascading chain of drop operations happens for each element in both lists.

2.1.1 Ownership. Perceus takes a more aggressive approach where *ownership* of references is passed down into each function: now `map` is in charge of freeing `xs`, and `ys` is freed by `print`: no `drop` operations are emitted inside `main` as all local variables are *consumed* by other functions, while

the `map` and `print` functions drop the list elements as they go. Let’s take a look at what reference count instructions Perceus generates for the `map` function:

```
fun map(xs : list(a), f : a -> e b) : e list(b)
  match xs
  Cons(x,xx) ->
    dup(x); dup(xx); drop(xs)
    Cons( dup(f)(x), map(xx,f))
  Nil ->
    drop(xs); drop(f)
  Nil
```

In the `Cons` branch, first the head and tail of the list are *dupped*, where a `dup(x)` operation increments the reference count of an object and returns itself. The `drop(xs)` then frees the initial list node. We need to `dup f` as well as it is used twice, while `x` and `xx` are consumed by `f` and `map` respectively.

Transferring ownership, rather than retaining it, means we can free an object immediately when no more references remain. This both increases cache locality and decreases memory usage. For `map`, the memory usage is halved: the list `xs` is deallocated while the new list `ys` is being allocated.

2.1.2 Reuse. Reuse analysis is an optimization that tries to reuse objects in-place. We can pair objects of known size with same sized allocated constructors and try to reuse these in-place at runtime. Reuse analysis rewrites `map` into:

```
fun map(xs : list(a), f : a -> e b) : e list(b)
  match xs
  Cons(x,xx) ->
    dup(x); dup(xx); val r = dropru(xs)
    Cons@r( dup(f)(x), map(xx,f))
```

The *reuse token* `r` becomes the address of the `Cons` cell `xs` if `xs` happens to be unique, and `NULL` otherwise. The `Cons@r` allocation reuses `xs` in-place if `r` is non `NULL`, and allocates a fresh `Cons` cell otherwise. In case we map over a unique list, the list elements are updated in-place. A further rewriting technique called *drop specialization* [33] can further optimize this by inlining the `dropru` operation and simplifying such that no reference count operations are necessary in the case that the list is unique:

```
fun map(xs : list(a), f : a -> e b) : e list(b)
  match xs
  Cons(x,xx) ->
    val r = if unique(xs) then &xs
            else dup(x); dup(xx); decref(xs); NULL
    Cons@r( dup(f)(x), map(xx,f))
  ...
```

This is an important optimization in practice but for the purposes of this paper we leave it out in our examples.

2.2 Problems with Reuse

Both Reinking, Xie et al. [33] and Ullrich and de Moura [38] describe reuse algorithms as a pass *before* the main Perceus algorithm. The reason they chose this approach is two-fold: the `dropru` function can be seen as consuming its argument and thus needs no special treatment from Perceus, and similarly, a reuse token can be deallocated by Perceus if it is not used (for example if the constructor is only allocated in one branch of a nested match-statement but not another).

However, in practice we observed that it can lead to situations where the reuse is not optimal. The algorithm (informally) described in Reinking, Xie et al. [33] only reuses an object if it is no longer live after being matched (we call this algorithm P). However, that requirement is too strong, for example, their algorithm can not reuse `x` for `Just(y)`:

```
match x
  Just(_) -> // x is still live here
  match y
    0 -> x
    - -> Just(y)
```

In practice, the actual implementation in Koka (which we will call algorithm K), does not do the liveness analysis at all, and generates code as:

```
match x
  Just(_) ->
    val r = dropru(dup(x))
    match y
      0 -> drop(y); x
      - -> drop(x); Just@r(y)
```

Here `x` is still live in the scope and which prevents any reuse at runtime due to the inserted `dup` operation. A way to improve on this, is by pushing down reuse operations into branches behind the last use of an object. This is the approach described by Ullrich and de Moura [38] (Fig. 3, algorithm D). Unfortunately, it turns out that this approach can lead to an arbitrary increase in peak memory usage. Consider the following example (B):

```
match xs
  Cons(_,_) ->
    val y = f(xs)
    Cons(y,Nil)
```

Using algorithm D, a reuse is inserted right after the call to `f`, which leads to Perceus (running afterwards) inserting a `dup` on the `xs` parameter:

```
match xs
  Cons(_,_) ->
    val y = f(dup(xs))
    val r = dropru(xs)
    Cons@r(y,Nil)
```

Even though the `Cons` cell of `xs` is now available for reuse, it also holds on to the full `xs` list during evaluation of `f`. This not only means that we may use `sizeof(xs)` more memory than necessary, but also that any reuse by `f` of `xs` is prevented (as `xs` is now certainly not unique).

2.3 Drop Guided Reuse

But it turns out there is a much simpler approach possible: instead of running reuse analysis upfront, we can perform reuse analysis *after* doing Perceus dup-drop insertion.

Indeed, Perceus already does precise liveness analysis and inserts optimal dup/drop operations in the sense that the resulting program is garbage-free. This means that the drop operations signify precisely when a cell may become available for reuse – and this is exactly the point where we should rewrite the `drop` to a `dropru` if there is any potential for reuse. The call to `dropru` can not hold on to entire objects longer than necessary since by the garbage-free property an object must be live until directly before it is dropped.

We can keep track of the currently known sizes of each variable (updated at each branch pattern) and if we encounter a drop we can statically determine if it can pair with a later allocation of the same size. Looking at our earlier example (A) from the previous section, Perceus generates first:

```
match x
  Just(_) ->
    match y
      0 -> drop(y); x
      _ -> drop(x); Just(y)
```

With the new *drop-guided* reuse analysis, we can now rewrite the `drop(x)` into a `dropru(x)` as we know that the size of `x` matches the size of the following `Just` allocation:

```
match x
  Just(_) ->
    match y
      0 -> dropru(y); x
      _ -> val r = dropru(x); Just@r(y)
```

For the other example (B) in the previous section, no drop operations are generated in the first place, and no `dropru` is inserted either – perfect!

This leaves the question though of how to free newly created reuse tokens if these happen to not be used in some branch. However, this is quite straightforward: since reuse tokens are only generated in a specific way, we can show they are only used either never or once, and never captured under a lambda or passed as an argument. Thus we can locally check at each branch of a match expression if a given reuse token `r` is used in this branch and insert a `free(r)` instruction if this is not the case.

2.4 Balanced Trees

As a practical example of the improved precision of drop-guided reuse analysis, we consider balanced insertion in red-black trees [17]. Figure 1 shows the implementation in Koka based on Okasaki’s algorithm [28]. The red-black tree has the invariant that the number of black nodes from the root to any of the leaves are the same, and that a red node is never a parent of red node. Together this ensures that the trees are always balanced. When inserting nodes, the invariants are maintained by rebalancing the nodes when needed. Let’s focus on the second branch in the `ins` function:

```
match t
  Node(B, l, kx, vx, r) ->
    if k < kx then
      if is-red(l) then lbal(ins(l,k,v), kx, vx, r)
      else Node(B, ins(l,k,v), kx, vx, r)
```

With all reuse algorithms, the `Node` allocation in the `else` branch will reuse the outer matched node `t`. It gets more interesting for the `then lbal(...)` branch though. First of all, both `is-red` and `lbal` will be inlined (and simplified); we focus just on the first branch of `lbal` here:

```
match t
  Node(B, l, kx, vx, r) ->
    if k < kx then
      match l
        Node(R,_,_,_,_) ->
          val t2 = ins(l,k,v)
          match t2
            Node(_, Node(R, ...), ...) ->
```

```
type color { R; B }

type tree(a)
  Node( clr:color, l:tree(a), key:int, value:a, r:tree(a) )
  Leaf

fun is-red(t : tree(a)) : bool
  match t
    Node(R,_,_,_,_) -> True
    _ -> False

fun lbal(l : tree(a), k : int, v : a, r : tree(a)) : tree(a)
  match l
    Node(_, Node(R, lx, kx, vx, rx), ky, vy, ry) ->
      Node(R, Node(B,lx,kx,vx,rx), ky, vy, Node(B,ry,k,v,r))
    Node(_, ly, ky, vy, Node(R, lx, kx, vx, rx)) ->
      Node(R, Node(B,ly,ky,vy,lx), kx, vx, Node(B,rx,k,v,r))
    Node(_, lx, kx, vx, rx) ->
      Node(B, Node(R,lx,kx,vx,rx), k, v, r)
    Leaf -> Leaf

fun rbal(l : tree(a), k : int, v : a, r : tree(a)) : tree(a)
  ...

fun ins(t : tree(a), k : int, v : a) : tree(a)
  match t
    Leaf -> Node(R, Leaf, k, v, Leaf)
    Node(B, l, kx, vx, r) ->
      if k < kx then
        if is-red(l) then lbal(ins(l,k,v), kx, vx, r)
        else Node(B, ins(l,k,v), kx, vx, r)
      elif k > kx then
        if is-red(r) then rbal(l, kx, vx, ins(r,k,v))
        else Node(B, l, kx, vx, ins(r,k,v))
      else Node(B, l, k, v, r)
    Node(R, l, kx, vx, r) ->
      if k < kx then Node(R, ins(l,k,v), kx, vx, r)
      elif k > kx then Node(R, l, kx, vx, ins(r,k,v))
      else Node(R, l, k, v, r)
```

Fig. 1. Balanced red-black tree insertion

`Node(R, Node(B, ...), ..., Node(B, ...))`

Due to the inlined `is-red` function, we now get a situation where all previous reuse algorithms fall short. This is exactly like example (B) in Section 2.2, and Koka’s old reuse algorithm K would try to reuse `l` even though it is still used later on, leading to:

```
match l
  Node(R,_,_,_,_) ->
    val ru2 = dropru(dup(l))
    val u = ins(l,k,v)
```

where `ru2` is never available for reuse at runtime. Algorithm D does not fare much better as it holds on to `l` preventing the recursive `ins` from reusing the tree:

```
match l
  Node(R,_,_,_,_) ->
    val t2 = ins(dup(l),k,v)
    val ru2 = dropreuse(l)
```

With drop-guided reuse, none of these problems occur and results in optimal reuse:

```
match t
  Node(B, l, kb, vb, r) ->
    dup(l); dup(kb); dup(vb); dup(r)
    val ru = dropreuse(t)
    if k < kb then
      match l
        Node(R,_,_,_,_) ->
          val t2 = ins(l,k,v)
```

```

match t2
  Node(_, l2 as Node(R, ...), ..., r2) ->
    val ru2 = dropreuse(t2)
    val ru3 = dropreuse(l2)
    Node@ru(R, Node@ru2(B, ...), ..., Node@ru3(B, ...))

```

It turns out that if we consider all branches, we can see that we either match one `Node` and allocate one, or we match three nodes deep and allocate three (when rebalancing). This means that for a unique tree every `Node` is reused in the fast path without doing any allocations!

Essentially, for a unique tree, *the purely functional* algorithm above adapts at runtime to an in-place mutating algorithm re-balancing along the spine without any further allocations. Moreover, if we use the tree *persistently* [29], and the tree is shared or has shared parts, the algorithm adapts to copying exactly the shared spine of the tree (and no more), while still rebalancing in place for any unshared parts.

2.4.1 TRMC: Tail-Recursion-Modulo-Cons. Koka also implements *tail-recursion-modulo-cons* (TRMC) optimization. Usually, with *tail-recursion* any functions that call themselves recursively in a tail position are transformed into a loop instead (using no extra stack space). With TRMC, such function can make its tail call inside any tail-position *expression* consisting of just constructors and non-allocating total expressions. For example, `map` is such function where the recursion in `Cons(f(x), map(f,xx))` is transformed into a loop. This is done by pre-allocating the `Cons` node ahead of the recursive call with a *hole* in the tail field, which is later assigned by the recursive call.

TRMC interacts well with reuse analysis as often the recursive call is inside a constructor that is reused. In the `map` example, this will result in traversing the list in a tight loop while updating each element in-place when the list `xs` happens to be unique.

TRMC is also effective for red-black tree rebalancing. We can see in Figure 1 that there are four TRMC recursive `ins` calls and only in the `rbal/lbal` cases this does not hold. When we study the generated C code the final result is quite sophisticated: there is an outer TRMC loop for the four TRMC `ins` calls, but it is interspersed at runtime with the two actual recursive `ins` calls. Moreover, due to reuse, the code updates unique nodes in place when rebalancing the spine, but adapts to copying for shared subtrees. Overall this is very efficient code that would be difficult to write directly by hand.

2.5 Borrowing

Even though the first example in Section 2.1 argues that arguments should be passed owned to the callee, this is not always optimal – sometimes it is better to pass arguments as *borrowed* instead [38].

Consider converting a list into an unbalanced binary tree:

```

fun make-tree( xs : list(a) ) : tree(a)
  match xs
    Cons(x, xx) -> Node( R, Leaf, 0, x, make-tree(xx) )
    Nil -> Leaf

```

Perceus inserts `dup(x)`; `dup(xx)`; `drop(xs)`; in the `Cons` branch, which causes some overhead, especially since there are no

reuse opportunities. When we annotate a parameter like `xs` to be borrowed (as `^xs`), the the caller keeps ownership of the parameter. As a result, no reference count operations need to be performed at all for the `xs` parameter in our example.

As shown by Ullrich and de Moura [38], this is not always beneficial though: if `xs` happens to be unique, we allocate the tree while `xs` is still live – exactly what we wanted to avoid in Section 2.1. In general, borrowing can increase the memory usage of a program by an arbitrary amount and it is it generally not frame-limited. Ullrich and de Moura [38] describe a borrow inference algorithm that marks `xs` automatically as borrowed. However, given that this is not safe for space, we argue that borrow inference should be further restricted to give a frame-limited guarantee. Therefore, Koka only uses borrowing for built-in primitives (like (big) integer operations), and currently has no borrow inference.

2.6 Frame-Limited Transformations

Perceus is *garbage-free*, and uses minimal memory for a particular program since objects are freed as soon as they are no longer live. We believe that this property is of central importance, as it not only reduces memory usage, but makes it more predictable and easier to reason about. Nevertheless, the property may be too strong, as we saw in the previous sections that many interesting optimizations do not preserve garbage-freeness.

Consider drop guided reuse analysis: At any evaluation step, it can hold on to a single cell per reuse token `r` that was created by `dropru`, but not used at a constructor yet. Since any function can only contain a constant number of `dropru` calls, one might expect the total overhead to be constant as well, which would make drop guided reuse *safe for space* [1, 30] (in the sense that the maximum peak memory increase is bounded by a constant). However, before a reuse token is used there might be a recursive call! Consider the `map` function we first viewed:

```

val r = dropru(xs)
Cons@r( dup(f)(x), map(xx, f) )

```

Here, `r` is live during the recursive call and so reuse analysis can hold on to as many `Cons` cells as either the list is long or the stack allows. In other words, we can only hope to bound the extra memory needed for reuse analysis by a constant factor times the current number of stack frames – we call this *frame-limited*. We formalize this notion in the next section, and show in Section 3.4 that drop guided reuse is a frame-limited transformation.

Even though weaker than being garbage-free or safe for space, we argue that frame-limited transformations are still good in practice. First of all, programmers are already aware of recursion and take steps to avoid unbounded recursion. Secondly, in practice the stack size is usually already bounded – in such case, that makes the frame-limited bound constant and thus actually safe for space.

In contrast, as we showed in Section 2.2, the reuse algorithm D [38] is not frame limited and can lead to an arbitrary

$$\begin{array}{l}
e ::= v \\
\quad | e e \\
\quad | \text{let } x = e \text{ in } e \\
\quad | \text{match } e \{ \overline{p_i \rightarrow e_i} \} \\
v ::= x \\
\quad | \lambda x. e \\
\quad | C \overline{v} \\
p ::= C \overline{x}
\end{array}$$

Semantics:

$$\begin{array}{l}
E ::= \square \mid E e \mid v E \\
\quad | \text{let } x = E \text{ in } e \\
\quad | \text{match } E \{ \overline{p_i \rightarrow e_i} \} \\
\frac{e \rightarrow e'}{E[e] \mapsto E[e']} \text{ [EVAL]} \\
(\text{app}) \quad (\lambda x. e) v \longrightarrow e[x:=v] \\
(\text{let}) \quad \text{let } x = v \text{ in } e \longrightarrow e[x:=v] \\
(\text{match}) \quad \text{match } (C \overline{v}) \{ \overline{p_i \rightarrow e_i} \} \longrightarrow e_i[\overline{x}:=\overline{v}] \\
\text{where } p_i = C \overline{x}
\end{array}$$

Fig. 2. Syntax and semantics of λ^1

increase in memory usage. Also, as shown in the previous section, borrowing in general is not frame-limited either: Once `make-tree` returns we still hold on to `xs`, which can be of arbitrary size. However, for borrowing there is hope: if we only borrow values that are bounded in size by a constant, then this is again frame-limited: In any function at most one function call happens at a time and so we can assign the constant cost of the borrowed values to the stack frame belonging to the surrounding function.

Note that in practice, backend optimizations like tail-recursion may optimize stack frames away. However, we formalize frame-limited in terms of the size of the evaluation context (i.e. the recursion depth) instead of actual stack frames and thus the bound still applies.

3 Formalization

We start out with the syntax and semantics given in [33] (see figure 2). This is essentially just lambda calculus extended with let bindings and pattern matching. We assume that the patterns p_i in a match are all distinct. The semantics for λ^1 is standard using strict evaluation contexts E [42]. The evaluation contexts uniquely determine where to apply an evaluation step using the `EVAL` rule. As such, evaluation contexts neatly abstract from the usual implementation context of a stack and program counter. The small step evaluation rules perform function application (`app`), let-binding (`let`), and pattern matching (`match`).

3.1 Heap Semantics

To reason precisely about reference counting, we need a semantics with an explicit heap. To make evaluation order and sharing explicit, we first translate any expression e into *normalized form* $\lfloor e \rfloor$, as defined in Figure 3, where all arguments become variables instead of values [13].

Figure 4 defines the syntax of the normalized linear resource calculus λ^{1n} where all arguments are now variables. Moreover the syntactic constructs in gray are only generated in derivations of the calculus and are not exposed to

$$\begin{array}{l}
\lfloor x \rfloor = x \\
\lfloor \lambda x. e \rfloor = \lambda x. \lfloor e \rfloor \\
\lfloor e e' \rfloor = \text{let } f = \lfloor e \rfloor \text{ in let } x = \lfloor e' \rfloor \text{ in } f x \\
\lfloor C v_1 \dots v_n \rfloor = \text{let } x_1 = \lfloor v_1 \rfloor \text{ in } \dots \text{let } x_n = \lfloor v_n \rfloor \text{ in } C x_1 \dots x_n \\
\lfloor \text{match } e \{ \overline{p_i \rightarrow e_i} \} \rfloor = \text{let } x = \lfloor e \rfloor \text{ in match } x \{ \overline{p_i \rightarrow \lfloor e_i \rfloor} \} \\
\lfloor \text{let } x = e_1 \text{ in } e_2 \rfloor = \text{let } x = \lfloor e_1 \rfloor \text{ in } \lfloor e_2 \rfloor
\end{array}$$

Fig. 3. Normalization. All f and x are fresh

users. Among those constructs, `dup` and `drop` form the basic instructions of reference counting, while `r←dropru` is used for reuse. Moreover, every lambda $\lambda_{\overline{z}}x.e$ is annotated with its free variables \overline{z} which becomes important during evaluation.

Contexts Δ, Γ are *multisets* containing variable names. We use the compact comma notation for summing (or splitting) multisets. For example, (Γ, x) adds x to Γ , and (Γ_1, Γ_2) appends two multisets Γ_1 and Γ_2 . The set of free variables of an expression e is denoted by $\text{fv}(e)$, and the set of bound variables of a pattern p by $\text{bv}(p)$.

Using our normalized calculus, Figure 5 defines the semantics in terms of a reference counted heap, where sharing of values becomes explicit, and substitution only substitutes variables. Here, each heap entry $x \mapsto^n v$ points to a value v with a reference count of n (with $n \geq 1$). In these semantics, values other than variables are allocated in the heap with rule (`lamr`) and rule (`conr`). The evaluation rules discard entries from the heap when the reference count drops to zero. Any allocated lambda is annotated as $\lambda_{\overline{z}}x.e$ to clarify that these are essentially *closures* holding an environment \overline{z} and a code pointer $\lambda x. e$. Note that it is important that the environment \overline{z} is a multi-set. After the initial translation, \overline{z} will be equivalent to the free variables in the body (see rule `LAM`), but during evaluation substitution may substitute several variables with the same reference. To keep reference counts correct, we need to keep considering each one as a separate entry in the closure environment.

When applying an abstraction, rule (`appr`) needs to satisfy the assumptions made when deriving the abstraction in rule `LAM` (shown in Figure 6). First, the (`appr`) rule inserts `dup` to duplicate the free closure variables \overline{z} , as these are owned in rule `LAM`. It then drops the reference to the closure itself.

A difference between (`appr`) and (`matchr`) is that for application the free variables \overline{z} are dynamic and thus the duplication must be done at runtime. In contrast, a match knows the the bound variables in a pattern statically and we therefore generate the required `dup` and `drop` operations statically during elaboration for each branch (as shown in Figure 6) – this is essential as that enables the further static optimizations of `dup/drop` pairs and reuse analysis.

We discuss the reuse evaluation rules later in Section 3.4.

3.2 Perceus

Figure 6 defines the Perceus logical derivation rules for inserting reference count instructions such that the resulting

$ \begin{array}{l} e ::= v \\ e x \\ \text{let } x = e \text{ in } e \\ \text{match } x \{ \overline{p_i} \mapsto e_i \} \\ \text{dup } x; e \\ \text{drop } x; e \\ r \leftarrow \text{dropru } x; e \end{array} $	$ \begin{array}{l} v ::= x \\ \lambda_{\overline{z}} x. e \\ C \overline{x} \\ p ::= C \overline{x} \\ \Delta, \Gamma ::= \emptyset \mid \Delta \cup x \end{array} $
---	--

$\lambda x. e \doteq \lambda_{\overline{z}} x. e \quad (\overline{z} = \text{fv}(e))$

Fig. 4. The normalized linear resource calculus λ^{ln} .

expression can be soundly evaluated by the heap semantics of Figure 5. The derivation $\Delta \mid \Gamma \vdash e \rightsquigarrow e'$ in Figure 6 reads as follows: given a *borrowed environment* Δ , a *linear environment* Γ , an expression e is translated into an expression e' with explicit reference counting instructions. We call variables in the linear environment *owned*.

The key idea is that each resource (i.e., owned variable) is consumed *exactly* once. That is, a resource needs to be explicitly duplicated (in rule `DUP`) if it is needed more than once; or be explicitly dropped (in rule `DROP`) if it is not needed. The rules are closely related to linear typing.

The rules are very close to the original rules [33] but differ in important details. In particular, by using a normalized form we only split the owned environment Γ in the `LET` rule, and no longer in the `APP` and `CON` rules which are now much simpler. This in turn allows us to parameterize the system by a single side condition (\star) that allows us to concisely capture *garbage-free* and *frame-limited* transformations as shown in Section 3.3.

The `LET` rule splits the owned environment Γ into two separate contexts Γ_1 and Γ_2 for expression e_1 and e_2 respectively. Each expression then consumes its corresponding owned environment. Since Γ_2 is consumed in the e_2 derivation, we know that resources in Γ_2 are surely alive when deriving e_1 , and thus we can *borrow* Γ_2 in the e_1 derivation. The rule is quite similar to the `[LET!]` rule of Wadler’s linear type rules [41, pg.14] where a linear type can be “borrowed” as a regular type during evaluation of a binding.

The `LAM` rule is interesting as it essentially derives the body of the lambda independently. The premise $\Gamma = \text{fv}(\lambda x. e)$ requires that exactly the free variables in the lambda are owned – this corresponds to the notion that a lambda is allocated as a closure at runtime that holds all free variables of the lambda (and thus the lambda expression *consumes* the free variables). The body of a lambda is evaluated only when applied, so it is derived under an empty borrowed environment only owning the argument and the free variables (in the closure). The translated lambda is also annotated with Γ , as $\lambda_{\Gamma} x. e$, so we know precisely the resources the lambda should own when evaluated in a heap semantics. We often omit the annotation when it is irrelevant.

Another difference from earlier work is that the `MATCH` rule now statically generates `dup` instructions for the pattern bindings which is essential to actually perform further `dup/drop` optimizations on the final derived expression. Finally, we also added the `DROPRU` rule to reason precisely about reuse analysis described later.

Properties. The logical derivation rules precisely elaborate expressions with reference count operations such that they can be correctly evaluated by the target heap semantics, as stated in the following theorem:

Theorem 1. (*Reference-counted heap semantics is sound*)

If we have $\emptyset \mid \emptyset \vdash e \rightsquigarrow e'$ and $e \mapsto^* v$, then we also have $\emptyset \mid e' \mapsto^*_r H \mid x$ with $[H]x = v$.

We use a novel approach to prove this theorem in separate steps: first we show soundness in a heap semantics that ignores reference count instructions (Appendix C.2), then use a separate resource calculus to show reference counts are correct (Appendix C.3), and finally combine these results for the final proof (Appendix C.4).

Second, we prove that the reference counting semantics never *hold on* to unused variables. We first define the notion of *reachability*.

Definition 1. (*Reachability*)

We say a variable x is reachable in terms of a heap H and an expression e , denoted as $\text{reach}(x, H \mid e)$, if (1) $x \in \text{fv}(e)$; or (2) for some y , we have $\text{reach}(y, H \mid e) \wedge y \mapsto^n v \in H \wedge \text{reach}(x, H \mid v)$.

With reachability, we can show (see Appendix C.5):

Theorem 2. (*Reference counting leaves no garbage*)

Given $\emptyset \mid \emptyset \vdash e \rightsquigarrow e'$, and $\emptyset \mid e' \mapsto^*_r H \mid x$, then for every intermediate state $H_i \mid e_i$, we have for all $y \in \text{dom}(H_i)$, $\text{reach}(y, H_i \mid e_i)$.

Note that λ^{ln} does not model *mutable references*. A natural extension of the system is to include mutable references and thus cycles. In that case, we could generalize Theorem 2, where the conclusion would be that for all resources in the heap, it is either reachable from the expression, or it is part of a cycle.

These theorems establish the correctness of the reference-counted heap semantics. However, correctness does not imply that evaluation is *garbage-free*. Eventually all live data is discarded but an evaluation may well hold on to live data too long by delaying drop operations. As an example, consider $y \mapsto^1 () \mid (\lambda x. x) (\text{drop } y; ())$, where y is reachable but dropped too late: it is only dropped after the lambda gets allocated. In contrast, a garbage-free algorithm would produce $y \mapsto^1 () \mid \text{drop } y; (\lambda x. x) ()$.

3.3 Parameterizing Perceus Derivations

Reinking, Xie et al. [33] give declarative derivation rules for reference counting but then provide a separate *algorithm* that is then proven to be garbage-free. This is not ideal as it does not provide any particular insight why the algorithm is

$H : x \mapsto (\mathbb{N}^+, v)$ $E ::= \square \mid E x \mid \text{let } x = E \text{ in } e$	$\frac{H \mid e \longrightarrow_r H' \mid e'}{H \mid E[e] \mapsto_r H' \mid E[e']} \text{ [EVAL]}$
$(\text{lam}_r) \quad H \mid \lambda_{\bar{z}} x. e \quad \longrightarrow_r \quad H, f \mapsto^1 \lambda_{\bar{z}} x. e \quad \mid f \quad \text{fresh } f$ $(\text{con}_r) \quad H \mid C x_1 \dots x_n \quad \longrightarrow_r \quad H, z \mapsto^1 C x_1 \dots x_n \quad \mid z \quad \text{fresh } z$ $(\text{app}_r) \quad H \mid f y \quad \longrightarrow_r \quad H \mid \text{dup } \bar{z}; \text{ drop } f; e[x:=y] \quad (f \mapsto^n \lambda_{\bar{z}} x. e) \in H$ $(\text{match}_r) \quad H \mid \text{match } y \{ \overline{p_i} \rightarrow e_i \} \quad \longrightarrow_r \quad H \mid e_i[\bar{x}:=\bar{z}] \quad \text{with } p_i = C \bar{x} \text{ and } (y \mapsto^n C \bar{z}) \in H$ $(\text{let}_r) \quad H \mid \text{let } x = z \text{ in } e \quad \longrightarrow_r \quad H \mid e[x:=z]$	
$(\text{dup}_r) \quad H, x \mapsto^n v \quad \mid \text{dup } x; e \quad \longrightarrow_r \quad H, x \mapsto^{n+1} v \quad \mid e$ $(\text{drop}_r) \quad H, x \mapsto^{n+1} v \quad \mid \text{drop } x; e \quad \longrightarrow_r \quad H, x \mapsto^n v \quad \mid e \quad \text{if } n \geq 1$ $(\text{dlam}_r) \quad H, x \mapsto^1 \lambda_{\bar{x}} y. e' \quad \mid \text{drop } x; e \quad \longrightarrow_r \quad H \mid \text{drop } \bar{x}; e$ $(\text{dcon}_r) \quad H, x \mapsto^1 C \bar{x} \quad \mid \text{drop } x; e \quad \longrightarrow_r \quad H \mid \text{drop } \bar{x}; e$	
Extension with reuse:	
$(\text{drop}_{ru}) \quad H, x \mapsto^{n+1} v \quad \mid r \leftarrow \text{dropru } x; e \quad \longrightarrow_r \quad H, x \mapsto^n v, z \mapsto^1 () \mid e[r:=z] \quad \text{fresh } z, \text{ if } n \geq 1$ $(\text{dlam}_{ru}) \quad H, x \mapsto^1 \lambda_{\bar{x}} y. e' \quad \mid r \leftarrow \text{dropru } x; e \quad \longrightarrow_r \quad H, z \mapsto^1 () \mid \text{drop } \bar{x}; e[r:=z] \quad \text{fresh } z$ $(\text{dcon}_{ru}) \quad H, x \mapsto^1 C \bar{x} \quad \mid r \leftarrow \text{dropru } x; e \quad \longrightarrow_r \quad H, \bar{z} \mapsto^1 (), z \mapsto^1 C \bar{z} \mid \text{drop } \bar{x}; e[r:=z] \quad \text{fresh } z, \bar{z}$	

Fig. 5. Reference-counted heap semantics for λ^{in} .

$\Delta \mid \Gamma \vdash e \rightsquigarrow e' \quad (\uparrow \text{ is input, } \downarrow \text{ is output})$ $\uparrow \quad \uparrow \quad \uparrow \quad \downarrow$	$\Delta \text{ and } \Gamma \text{ are multisets of the borrowed and owned variables in scope}$
$\frac{}{\Delta \mid x \vdash x \rightsquigarrow x} \text{ [VAR]} \quad \frac{}{\Delta \mid \bar{x} \vdash C \bar{x} \rightsquigarrow C \bar{x}} \text{ [CON]} \quad \frac{\Delta \mid \Gamma, x \vdash e \rightsquigarrow e' \quad x \in \Delta, \Gamma}{\Delta \mid \Gamma \vdash e \rightsquigarrow \text{dup } x; e'} \text{ [DUP]}$	
$\frac{\Delta, x \mid \Gamma \vdash e \rightsquigarrow e'}{\Delta \mid \Gamma, x \vdash e x \rightsquigarrow e' x} \text{ [APP]} \quad \frac{\Delta \mid \Gamma \vdash e \rightsquigarrow e'}{\Delta \mid \Gamma, x \vdash e \rightsquigarrow \text{drop } x; e'} \text{ [DROP]} \quad \frac{\Delta \mid \Gamma, r \vdash e \rightsquigarrow e' \quad \text{fresh } r}{\Delta \mid \Gamma, x \vdash e \rightsquigarrow r \leftarrow \text{dropru } x; e'} \text{ [DROPRU]}$	
$\frac{\emptyset \mid \Gamma, x \vdash e \rightsquigarrow e' \quad \Gamma = \text{fv}(\lambda x. e)}{\Delta \mid \Gamma \vdash \lambda x. e \rightsquigarrow \lambda_{\Gamma} x. e'} \text{ [LAM]} \quad \frac{\Delta, \Gamma_2 \mid \Gamma_1 \vdash e_1 \rightsquigarrow e'_1 \quad \Delta \mid \Gamma_2, x \vdash e_2 \rightsquigarrow e'_2 \quad (\star)}{\Delta \mid \Gamma_1, \Gamma_2 \vdash \text{let } x = e_1 \text{ in } e_2 \rightsquigarrow \text{let } x = e'_1 \text{ in } e'_2} \text{ [LET]}$	
$\frac{\Delta \mid \Gamma, \bar{z}_i \vdash e_i \rightsquigarrow e'_i \quad p_i = C_i \bar{z}_i \quad x \in \Delta, \Gamma}{\Delta \mid \Gamma \vdash \text{match } x \{ \overline{p_i} \mapsto e_i \} \rightsquigarrow \text{match } x \{ \overline{p_i} \mapsto \text{dup}(\bar{z}_i); e'_i \}} \text{ [MATCH]}$	

Fig. 6. Perceus logical derivation rules. The rules are parameterized by the (\star) condition on LET. We write \vdash_{GF} for garbage-free derivations, and use_{FL} for derivations that are frame-limited.

garbage-free, and if other approaches exist as well.

It turns out we can actually capture the garbage-free property *declaratively* as a single side condition on the LET in our new derivation rules. Moreover, by weakening the condition, this can also be used to capture frame-limited derivations. We can thus instantiate the rules by giving a specific (\star) condition:

- *General* derivations \vdash . When we define (\star) to be true, the evaluation of any derived expression is *sound* (Theorem 1).
- *Garbage-free* derivations \vdash_{GF} . When we define the (\star) condition as $\Gamma_2 \subseteq \text{fv}(e_2)$, then the evaluation of any derived expression is *garbage-free* (Definition 2 and Theorem 3). At the LET rule we have the freedom to split Γ into Γ_1 and Γ_2 in any way. For garbage-free derivations though we try to minimize borrowing in the e_1 derivation (of Γ_2) and

thus the condition captures intuitively that we should use the smallest Γ_2 possible, and not include variables that are not needed for the e_2 derivation.

- *Frame-limited* derivations \vdash_{FL} . When we define (\star) as $\Gamma_2 = \Gamma', \Gamma''$ where $\Gamma' \subseteq \text{fv}(e_2)$ and $\text{sizeof}(\Gamma'') \leq c$ for some constant c , then the evaluation for any derived expression is *frame-limited* (Definition 3 and Theorem 4). We define $\text{sizeof}(\Gamma'')$ as the sum of the sizes of each element: $\sum_{y \in \Gamma''} \text{sizeof}(y)$. This weakens the garbage-free condition to allow borrowing of any y where the runtime size of y is known to be limited by a constant. This is just enough for transformations like reuse where we borrow a reuse token until it can be reused.

Garbage-Free Evaluation. We define *garbage-free* evaluation formally as:

$\begin{array}{c} \text{S} \mid \text{R} \Vdash e \rightsquigarrow e' \quad (\uparrow \text{ is input, } \downarrow \text{ is output}) \\ \uparrow \quad \uparrow \quad \uparrow \quad \downarrow \end{array}$		
$\text{S maps variables to their heap size (if known), R maps reuse variables } r \text{ to their available heap size, } \text{dom}(\text{R}) \cap \text{fv}(e)$		
$\frac{}{\text{S} \mid \emptyset \Vdash x \rightsquigarrow x} \text{ [RVAR]}$	$\frac{\text{S} \mid \text{R} \Vdash e \rightsquigarrow e'}{\text{S} \mid \text{R}, r : n \Vdash e \rightsquigarrow \text{drop } r; e'} \text{ [RDROPR]}$	
$\frac{}{\text{S} \mid \emptyset \Vdash C \bar{x} \rightsquigarrow C \bar{x}} \text{ [RCON]}$	$\frac{x : n \in \text{S} \quad \text{S} \mid \text{R}, r : n \Vdash e \rightsquigarrow e' \quad \text{fresh } r}{\text{S} \mid \text{R} \Vdash \text{drop } x; e \rightsquigarrow r \leftarrow \text{dropru } x; e'} \text{ [RDROP-REUSE]}$	
$\frac{\emptyset \mid \emptyset \Vdash e \rightsquigarrow e'}{\text{S} \mid \emptyset \Vdash \lambda_{\Gamma} x. e \rightsquigarrow \lambda_{\Gamma} x. e'} \text{ [RLAM]}$	$\frac{\text{S} \mid \text{R}_1 \Vdash e_1 \rightsquigarrow e'_1 \quad \text{S} \mid \text{R}_2 \Vdash e_2 \rightsquigarrow e'_2}{\text{S} \mid \text{R}_1, \text{R}_2 \Vdash \text{let } x = e_1 \text{ in } e_2 \rightsquigarrow \text{let } x = e'_1 \text{ in } e'_2} \text{ [RLET]}$	
$\frac{\text{S} \mid \text{R} \Vdash e \rightsquigarrow e'}{\text{S} \mid \text{R} \Vdash e x \rightsquigarrow e' x} \text{ [RAPP]}$	$\frac{\text{S}, x : n \mid \text{R} \Vdash e_i \rightsquigarrow e'_i \quad p_i = C x_1 \dots x_n}{\text{S} \mid \text{R} \Vdash \text{match } x \{ \overline{p_i} \mapsto e_i \} \rightsquigarrow \text{match } x \{ \overline{p_i} \mapsto e'_i \} } \text{ [RMATCH]}$	
$\frac{\text{S} \mid \text{R} \Vdash e \rightsquigarrow e'}{\text{S} \mid \text{R} \Vdash \text{dup } x; e \rightsquigarrow \text{dup } x; e'}$	$\frac{\text{S} \mid \text{R} \Vdash e \rightsquigarrow e'}{\text{S} \mid \text{R} \Vdash \text{drop } x; e \rightsquigarrow \text{drop } x; e'}$	$\frac{\text{S} \mid \text{R}, r \Vdash e \rightsquigarrow e'}{\text{S} \mid \text{R} \Vdash r \leftarrow \text{dropru } x; e \rightsquigarrow r \leftarrow \text{dropru } x; e}$

Fig. 7. Declarative rules for reuse analysis

Definition 2. (*Garbage free evaluation*)

An evaluation $\emptyset \mid e' \mapsto_r^* \text{H} \mid x$ is called *garbage-free* iff for every intermediate state $\text{H}_i \mid \text{E}[v]$ in the evaluation, we have that for all $y \in \text{dom}(\text{H}_i)$, $\text{reach}(y, \text{H}_i \mid \text{E}[v])$.

where we use the notation $[e]$ to erase all drop and dup in the expression e . This is a refinement of the definition given in [33], which considered any non dup/drop steps, while we weaken this and consider only value steps v . By using $\text{E}[v]$ we consider exactly those points where we are at an allocation step ($C \bar{x}$ or $\lambda_{\bar{x}} x. e$), and this is exactly the point where we want to ensure that there is no garbage. In particular, match and let are heap invariant, and applications just expand to dups, drops, and substitution. This also gives more freedom to garbage-free algorithms as it becomes possible for example to push a drop down into the branches of a match which was not possible before. Using our new definition, we can then prove (Appendix C.6):

Theorem 3. (*\vdash_{GF} derivations are garbage-free*)

If $\emptyset \mid \emptyset \vdash_{\text{GF}} e \rightsquigarrow e'$ and $\emptyset \mid e' \mapsto_r^* \text{H} \mid x$, then the evaluation is garbage-free.

Even with the garbage-free side-condition, there are still many choice points in the derivations which gives us freedom to consider various algorithms. Generally though, when implementing Perceus one wants to dup as late as possible (push up dup into the leaves of the derivation), and do drops as early as possible. The original Perceus algorithm [33] does this, and also trivially satisfies our garbage-free condition as it has the invariant $\Gamma \subseteq \text{fv}(e)$ for any derivation step.

Frame-Limited Evaluation. Similar to garbage-free evaluations, we can define frame-limited evaluations:

Definition 3. (*Frame-limited evaluation*)

An evaluation $\emptyset \mid e' \mapsto_r^* \text{H} \mid x$ is called *frame-limited* iff for every intermediate state $\text{H}_i \mid \text{E}[v]$ in the evaluation, we have that H_i equals H_1, H_2 such that for all $y \in \text{dom}(\text{H}_1)$, $\text{reach}(y, \text{H}_1 \mid \text{E}[v])$ and $|\text{H}_2| \leq c \cdot |\text{E}|$ for some constant c .

This expresses that at every allocation step, we may now hold on to extra heap space H_2 , but the size of H_2 is limited by a constant amount times the size of the evaluation context (i.e. the stack). We can now prove (Appendix C.7):

Theorem 4. (*\vdash_{FL} derivations are frame-limited*)

If $\emptyset \mid \emptyset \vdash_{\text{FL}} e \rightsquigarrow e'$ and we have $\emptyset \mid e' \mapsto_r^* \text{H} \mid x$, then the evaluation is frame-limited with respect to the constant c chosen in the (\star) condition.

Borrowing a variable for a function call can also be frame-limited: Whenever an owned variable is passed as borrowed, we need to move it from the linear environment to the borrowed environment. But this can only happen in the APP and LET rules. With the APP rule, we can safely borrow since we use the variable later and with the LET rule this is frame-limited whenever the size of the borrowed variables is bounded by a constant.

3.4 Drop-Guided Reuse

Figure 7 shows the declarative derivation rules for drop-guided reuse analysis. A rule $\text{S} \mid \text{R} \Vdash e \rightsquigarrow e'$ states we can derive e' from e given a mapping S from variables to their heap cell size (if known), and a mapping R from reuse tokens r to their available heap size. Again, we use a multi-set for S and R; we need to ensure we only use a reuse token r once, and similar to the owned environment Γ the leaf derivations require R to be empty (as in RVAR, RCON, and RLAM).

We state drop-guided reuse in terms of derivation rules instead of a specific algorithm in order to clearly expose the

choice points. At any drop x ; e expression we can choose to either leave it as is, or use `RDROP-REUSE` to try to reuse it at runtime. We can only use `RDROP-REUSE` though if the size of x is known statically – in our rules this only happens by matching on a particular constructor (in `RMATCH`) but in practice we may also use type information for example. Furthermore, in an implementation we also would only apply `RDROP-REUSE` if there is an actual opportunity in e for reuse to occur – that is, we look first if e contains an occurrence of $C\ x_1 \dots x_n$.

Another choice point is in the `RLET` rule where we can freely split R into R_1 and R_2 , where we may either reuse early or late if reuse is possible in both e_1 and e_2 . Also, rule `RDROPR` is not syntax directed and thus we have a choice in what reuse token to use if there are multiple reuse tokens of the right size.

To simplify the formalization and proofs, we do not introduce a new form of constructor reuse that we used before (as `C@R`) but instead assume the runtime evaluation recognizes a pair drop r ; $C\ \bar{x}$ as a reuse opportunity, that is:

$$\begin{aligned} H, r \mapsto^1 C' \perp_1 \dots \perp_n \mid \text{drop } r; C\ x_1 \dots x_n &\longrightarrow_r \\ H, x \mapsto^1 C\ x_1 \dots x_n \mid x & \end{aligned}$$

Also, to simplify the formalization, we do not add \perp either and use allocated unit constructors $()$ instead. The rules for dropru evaluation are given in Figure 5. We can now show that drop-guided reuse is sound and frame-limited (Appendix C.8):

Theorem 5. (*Reuse is frame-limited*)

If $\Delta \mid \Gamma \vdash_{\text{GF}} e \rightsquigarrow e'$, and reuse derives $S \mid R \Vdash e' \rightsquigarrow e''$, then $\Delta \mid \Gamma, R \vdash_{\text{FL}} e \rightsquigarrow e''$.

That is, if we have a garbage free derivation followed by drop-guided reuse, we could have derived that same expression directly as well using a frame-limited derivation. This is how Koka implements this as well: first a Perceus algorithm (that satisfies garbage-free derivations) followed by a drop-guided reuse algorithm (that satisfies reuse derivations).

4 Benchmarks

We measured the performance drop-guided reuse in Koka, versus the previous algorithm, and also against state-of-the-art memory reclamation implementations in various other languages. Since we compare across languages we need to interpret the results with care – the results depend not only on memory reclamation but also on the different optimizations performed by each compiler and how well we can translate each benchmark to that particular language. These results are therefore mostly useful to get a sense of the absolute performance of Koka, and as evidence that our compilation techniques (Perceus, drop-guided reuse, TRMC) are viable and can be competitive.

We selected mature comparison systems that use a range of memory reclamation techniques and are considered best-in-class. The systems we compare are:

- Koka v2.3.3 with drop-guided reuse, compiling the generated C code with gcc 9.3.0 using a customized version of the mimalloc allocator [22]. We also run Koka “-fno-opttrmc” where TRMC is disabled to measure the impact of those optimizations. We also measure “Koka old” which is a branch of Koka (v2.3.3-old) with the old-style reuse algorithm without borrowing.
- Multi-core OCaml 4.12, This has a concurrent generational collector with a minor and major heap. The minor heap uses a copying collector, while a tracing collector is used for the major heap [11, 26, Chap.22, 35]. The Koka benchmarks correspond essentially one-to-one to the OCaml versions.
- Haskell, GHC 8.6.5. A highly optimizing compiler with a multi generational garbage collector. The benchmark sources again correspond very closely, but since Haskell has lazy semantics, we used strictness annotations in the data structures to speed up the benchmarks, as well as to ensure that the same amount of work is done.
- Swift 5.5. The only other language in this comparison where the compiler uses reference counting [7, 39]. The benchmarks are directly translated to Swift in a functional style without using direct mutation. However, we translated tail-recursive definitions to explicit loops.
- Java 17 LTS, Uses the HotSpot JVM and the G1 concurrent, low-latency, generational garbage collector. The benchmarks are directly translated from Swift.
- C++, gcc 9.3.0 using the standard libc allocator. A highly optimizing compiler with manual memory management. We use C++ as our performance baseline: For the `rbtree` benchmark we used the standard STL `std::map` implementation that uses a highly optimized in-place updating version of red-black trees [14]. The `binarytrees` benchmark use a monotonic buffer resource for memory management, while the other benchmarks (`nqueens`, `deriv`, and `cfold`) do not reclaim memory at all.

4.1 Benchmarking Balanced Trees

We use the same red-black tree benchmarks as used in the Perceus paper [33]. There are two versions:

- `rbtree`: inserts 4200000 elements in tree and afterwards folds over the tree counting the `True` elements.
- `rbtree-ck`: like `rbtree` but also keeps a list of every 10th tree that is generated, effectively sharing many subtrees. This implies many subtrees have a non-unique reference count which causes many more slow paths to be taken in the Koka code. This is not done for C++ as `std::map` does not support persistence.

The results (together with other benchmarks) are shown in Figure 8 (on an AMD5950X, see Figure 9 in the Appendix for the results on an M1). Even though Koka has few optimizations (besides Perceus) in comparison with these mature systems, it performs surprisingly well. In fact, it outperforms the C++ implementation by almost 20% – how is that even possible? We conjecture this is mainly due to two factors: 1)

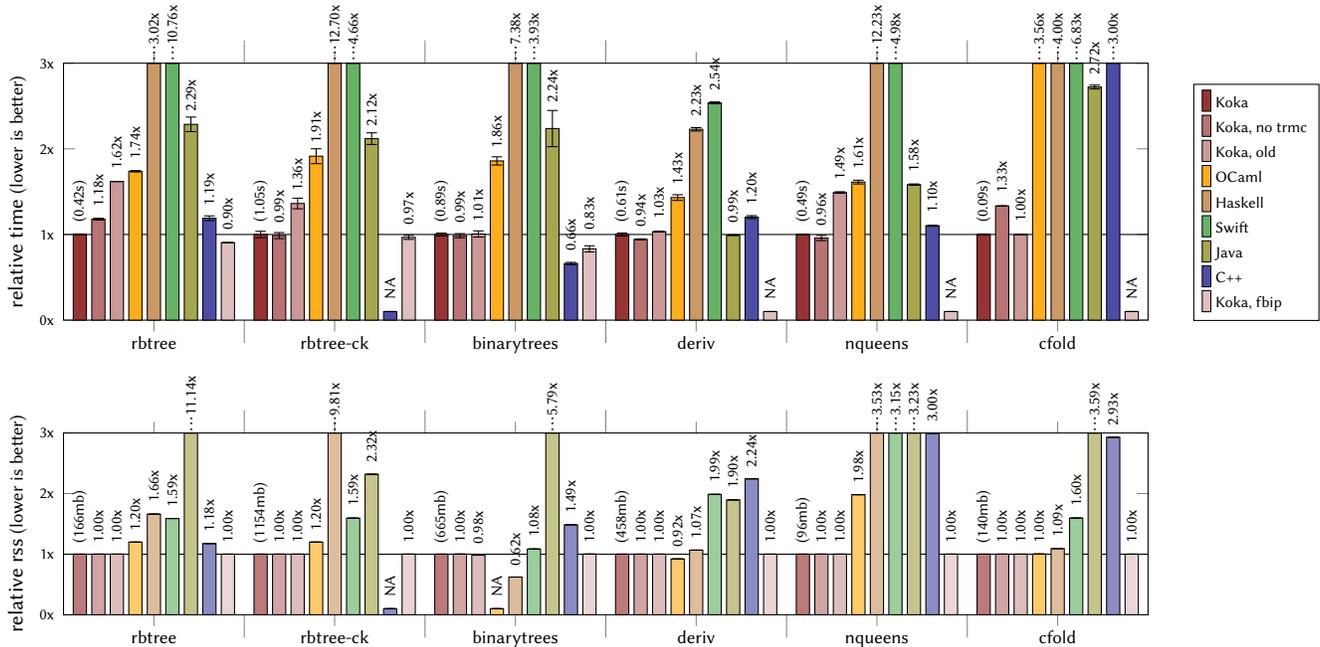


Fig. 8. All benchmarks with relative execution time and peak working set with respect to Koka. Using a 32-core x64 AMD5950X 3.4Ghz with 32GiB 3600Mhz memory (32KiB L1, 512KiB L2), Ubuntu 20.04.

the TRMC optimization leads to using less stack space than the C++ implementation (and perhaps less register-spills), and 2) the allocator that Koka uses can use 8-byte alignment while C++ requires 16-byte minimal alignment which leads to less allocated memory. If we compare the non-TRMC optimized version it performs only slightly better than C++ which seems to support our thesis.

Even though this is a typical functional style algorithm both OCaml and Haskell are much slower (1.7× and 3×). The Swift version is about 10× slower even though it uses reference counting as well – this may be a combination of no reuse (due to default borrowing) and a more complex implementation of the reference counts.

The `rbtrees-ck` version with many shared subtrees increases the relative speed of Koka even further. This is also somewhat surprising as it is clearly much worse for reuse (with many shared subtrees), but of course it similarly causes more pressure on garbage collectors as well as more objects get promoted to older generations. The Multi-core OCaml GC performs especially well here staying within 20% of Koka’s *garbage-free* memory usage.

If we compare Koka with drop-guided reuse against Koka “old”, we can see the new algorithm does about 1.6× better on this benchmark due to the improved reuse for the `is-red` test as explained in Section 2.4. Finally, what is the fastest “Koka fbip” version? We will discuss this in Section 4.2.1.

4.2 Binary Trees

In this section we look at the `binarytrees` benchmark from the Computer Language Benchmark Game [36] which is an

adaptation of Hans Boehm’s GCbench benchmark [5] (which in turn was adapted from a benchmark by John Ellis and Pete Kovac). This is an interesting benchmark as it uses *concurrent* allocation of many binary trees and calculating their checksums. Moreover, we can compare against the best performing implementations that were created by experts in each of our comparison languages.

The top 17 implementations are all languages with manual allocation (C++, C, Rust, and Free Pascal) with the top entry being C++ (#7) followed very closely by Rust and two other C++ entries (#5 and #4). For our purposes, we use C++ benchmark #5 since the performance of #7 was uneven across systems¹. The C++ entries all use very efficient allocation by using a `monotonic_buffer_resource` for bump-pointer allocation where all nodes are freed at once at every iteration.

Figure 8 shows the benchmark results of `binarytrees` (the Koka implementation can be found in Appendix A.1). Besides C++, Koka outperforms all other languages here even though it uses a simple thread pool implementation without work-stealing. The C++ implementation is still quite a bit faster though (0.66×). Since the benchmark does concurrent allocations reference counting generally needs to be atomic. However, due to careful language design, Koka can avoid most of the overhead by checking upfront if a reference count needs to be atomic or not [33].

¹Performance of benchmark #7 is dependent on the version of the Intel TBB library, and more than twice as slow as #5 on arm64.

4.2.1 Folding Binary Trees using FBIP. Most of the work in the `binarytrees` benchmark is in creating the trees and calculating their size using the `check` function:

```
type tree
  Node( left: tree, right: tree )
  Tip

fun check( t : tree ) : int
  match t
  | Node(l,r) -> check(l) + check(r) + 1
  | Tip       -> 0
```

This consists of two non-tail-recursive calls to `check`. We can improve upon this using FBIP. In particular, for any datatype we can derive a *visitor* datatype. For our purposes we define:

```
type visit
  NodeR( right: tree, rest: visit )
  Done
```

We can use the new visitor datatype to write a `check` function that uses no extra stack space as all calls are tail-recursive:

```
fun check-fbip( t : tree, v : visit, acc : int ) : int
  match t
  | Node(l,r) -> check-fbip( l, NodeR(r,v), acc + 1 ) // (A)
  | Tip       -> match v
  | NodeR(r,v') -> check-fbip( r, v', acc ) // (B)
  | Done        -> acc // (C)
```

At every `Node` we directly go down the left branch but remember that we still need to visit the right node by extending our `visit` datatype (A). When we reach a `Tip` we go through our visitor to now visit the saved right nodes (B) until we are done (C).

This may look more expensive, but when `t` happens to be unique at runtime, the `NodeR` allocations in (1) will reuse the `Node` that is matched – effectively updating these nodes in place to create a list of nodes that still need to be visited. At runtime this becomes tight loop that directly reuses the memory of the tree it is checking. In Figure 8 this is the *Koka fbip* variant and with this implementation of `check` Koka becomes 17% faster and within 25% of the performance of the C++ implementation (0.8× versus Koka fbip).

We can apply the same FBIP technique on our previous red-black tree balanced insertion: as we saw, due to TRMC most calls to `ins` are tail recursive but not the ones that needed rebalancing. We can define a visitor for red-black trees as the derivative of the `tree` datatype:

```
type zipper(a)
  NodeR(cclr:color, l:tree(a), key:int, value:a, zip:zipper(a))
  NodeL(cclr:color, zip:zipper(a), key:int, value:a, r:tree(a))
  Done
```

Using this datatype we can now first traverse down a tree in tail-recursive way to the insertion point, building up a `zipper` that is then used to reconstruct the tree going back up. The traversal back up is also tail recursive and rebalances the tree as needed [28] – the implementation can be found in Appendix A.2. This is the *Koka fbip* variant in Figure 8 which is about 10% faster than the regular version and now around 30% faster than the C++ STL version.

4.3 Other Benchmarks

The `nqueens`, `cfold`, and `deriv` benchmarks are the same as used in the Perceus paper [33]. We included them here for completeness but for these benchmarks the new drop guided reuse gives very similar results to the “old” algorithm.

The `nqueens` benchmark computes a list of all solutions to the N-queens problem of size 21. The large speedup here compared to Koka “old” is actually due to borrowing in the safe function. The `cfold` benchmark performs constant folding in program, and the `deriv` benchmark computes a symbolic derivative of a large expression. Again, in these benchmarks the difference with the previous reuse algorithm is minimal.

5 Related Work

Our work is closely based earlier work by Reinking, Xie et al. [33] and Ullrich and de Moura [38] (in the context of the Lean theorem prover [27]). In this work we improve upon the both of the two earlier reuse algorithms, and show that drop-guided reuse is strictly better as it is frame-limited and can find more opportunities for reuse. Our λ^n calculus, heap semantics, and derivation rules differ in important details from [33]: the normalization simplifies the derivation rules and allows us to concisely express the garbage-free and frame-limited side conditions (and no longer as a particular property of a specific algorithm). It is also now immediate that the previous published Perceus algorithm is garbage-free. The new frame-limited condition allows us to reason about various transformations that are no longer garbage-free but can still be bounded.

The notion of safe-for-space was introduced by Appel [1] and further studied by Paraskevopoulou and Appel [30]. Similar to our reuse transformation and frame-limited notion, the latter work introduces a general framework to show that the closure conversion transformation with flat environments is safe-for-space, while linked closure conversion is not. Other examples where a program transformation was proven to respect a resource bound include Crary and Weirich [9] and Minamide [25].

Using explicit reference count instructions in order to optimize them via static analysis is described as early as Barth [3]. Mutating unique references in place has traditionally focused on array updates [18], as in functional array languages like Sisal [24] and SaC [16, 34]. Férey and Shankar [12] provide functional array primitives that use in-place mutation if the array has a unique reference which is also present in the Koka implementation. We believe this would work especially well in combination with reuse-analysis for BTree-like structures using trees of small functional arrays.

The λ^1 and λ^{ln} calculus are closely based on linear logic. Turner and Wadler [37] give a heap-based operational interpretation which does not need reference counts as linearity is tracked by the type system. In contrast, Chirimar et al. [6] give an interpretation of linear logic in terms of reference counting, but in their system, values with a linear type are not guaranteed to have a unique reference at runtime.

Generally, a system with linear types [41], like linear Haskell [4], or the uniqueness typing of Clean [2, 40], can offer *static* guarantees that the corresponding objects are unique at runtime, so that destructive updates can always be performed safely. However, this usually also requires writing multiple versions of a function for each case (unique-versus shared argument). By contrast, reuse analysis relies on dynamic runtime information, and thus reuse can be performed generally. This is also what enables FBIP to use a single function that can be used for both unique or shared objects (since the uniqueness property is *not* part of the type). These two mechanisms could be combined: if our system is extended with unique types, then reuse analysis could statically eliminate corresponding uniqueness checks.

The Swift language is widely used in iOS development and uses reference counting with an explicit representation in its intermediate language. There is no reuse analysis but, as remarked by Ullrich and de Moura [38], this may not be so important for Swift as typical programs mutate objects in-place. There is no cycle collection for Swift, but despite the widespread usage of mutation this seems to be not a large problem in practice. Since it can be easy to create accidental cycles through the `self` pointer in callbacks, Swift has good support for *weak* references to break such cycles in a declarative manner. Ungar et al. [39] optimize atomic reference counts by tagging objects that can be potentially thread-shared. Later work by Choi et al. [7], uses *biased* reference counting to avoid many atomic updates. They remove the need for tagging by extending the object header with the ID of the owning thread together with two reference counts: a shared one that needs atomic updates, and the biased one for the owning thread.

The CPython implementation also uses reference counting, and uses ownership-based reference counts for parameters but still only drops the reference count of local variables when exiting the frame. Another recent language that uses reference counting is Nim. The reference counting method is scope-based and uses non-atomic operations (and objects cannot be shared across threads without extra precautions). Nim can be configured to use ORC reference counting which extends the basic ARC collector with a cycle collection [44]. Nim has the `acyclic` annotation to identify data types that are (co)-inductive, as well as the `(unsafe) cursor` annotation for variables that should not be reference counted.

In our work we focus on *precise* and *garbage free* reference counting which enables static optimization of reference count instructions. On the other extreme, Deutsch and Bobrow [10] consider *deferred* reference counting – any reference count operations on stack-based local variables are *deferred* and only the reference counts of fields in the heap are maintained. Much like a tracing collector, the stack roots are periodically scanned and deferred reference counting operations are performed. Levanoni and Petrank [23] extend this work and present a high performance reference counting collector for Java that uses the *sliding view* algorithm to

avoid many intermediate reference counting operations and needs no synchronization on the write barrier.

6 Conclusion and Future Work

In this work we show the effectiveness of drop-guided reuse, and give a precise characterization of garbage-free and frame-limited evaluations. However, this works well partly because in a functional style language like Koka it is uncommon to create cycles (which can only be created through mutable references). We would like to see if we can combine some of the static analysis with cycle collection. Also, we are interested in language support to guarantee that reuse is happening at compile time.

References

- [1] Andrew W. Appel. *Compiling with Continuations*. Cambridge University Press. 1991. doi:10.1017/CBO9780511609619.
- [2] Erik Barendsen, and Sjaak Smetsers. “Uniqueness Typing for Functional Languages with Graph Rewriting Semantics.” *Mathematical Structures in Computer Science* 6 (6). Cambridge University Press: 579–612. 1996. doi:10.1017/S0960129500070109.
- [3] Jeffrey M. Barth. *Shifting Garbage Collection Overhead to Compile Time*. UCB/ERL M524. EECS Department, University of California, Berkeley. Jun. 1975. <http://www2.eecs.berkeley.edu/Pubs/TechRpts/1975/29109.html>.
- [4] Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. “Linear Haskell: Practical Linearity in a Higher-Order Polymorphic Language.” *Proc. ACM Program. Lang.* 2 (POPL). Dec. 2017. doi:10.1145/3158093.
- [5] Hans Boehm. “GC Bench.” 2000. https://hboehm.info/gc/gc_bench.
- [6] Jawahar Chirimar, Carl A. Gunter, and Jon G. Riecke. “Reference Counting as a Computational Interpretation of Linear Logic.” *Journal of Functional Programming* 6: 6–2. 1996.
- [7] Jiho Choi, Thomas Shull, and Josep Torrellas. “Biased Reference Counting: Minimizing Atomic Operations in Garbage Collection.” In *Proceedings of the 27th International Conference on Parallel Architectures and Compilation Techniques*. PACT ’18. Limassol, Cyprus. 2018. doi:10.1145/3243176.3243195.
- [8] George E Collins. “A Method for Overlapping and Erasure of Lists.” *Communications of the ACM* 3 (12). ACM New York, NY, USA: 655–657. 1960.
- [9] Karl Cray, and Stephnie Weirich. “Resource Bound Certification.” In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 184–198. POPL ’00. Boston, MA, USA. 2000. doi:10.1145/325694.325716.
- [10] L. Peter Deutsch, and Daniel G. Bobrow. “An Efficient, Incremental, Automatic Garbage Collector.” *Communications of the ACM* 19 (9): 522–526. Sep. 1976. doi:10.1145/360336.360345.
- [11] Damien Doligez, and Xavier Leroy. “A Concurrent, Generational Garbage Collector for a Multithreaded Implementation of ML.” In *Proceedings of the 20th ACM Symposium on Principles of Programming Languages (POPL)*, 113–123. ACM press. Jan. 1993.
- [12] Gaspard Férey, and Natarajan Shankar. “Code Generation Using a Formal Model of Reference Counting.” In *NASA Formal Methods*, edited by Sanjai Rayadurgam and Oksana Tkachuk, 150–165. Springer International Publishing. 2016.
- [13] Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. “The Essence of Compiling with Continuations.” In *Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*, 237–247. PLDI ’93. Albuquerque, New Mexico, USA. 1993. doi:10.1145/155090.155113.
- [14] Free Software Foundation, Silicon Graphics, and Hewlett-Packard Company. “Internal Red-Black Tree Implementation for “Stl::map.”” 1994. <https://github.com/gcc-mirror/gcc/tree/master/libstdc++v3/src/c++98/tree.cc>.
- [15] Matt Gallagher. “Reference Counted Releases in Swift.” Dec. 2016. <https://www.cocoawithlove.com/blog/resources-releases-reentrancy.html>. Blog post.
- [16] Clemens Grelck, and Kai Trojahnner. “Implicit Memory Management for SAC.” In *6th International Workshop on Implementation and Application of Functional Languages (IFL ’04)*. Lübeck, Germany. Sep. 2004.
- [17] Leo J Guibas, and Robert Sedgewick. “A Dichromatic Framework for Balanced Trees.” In *19th Annual Symposium on Foundations of Computer Science (sfcs 1978)*, 8–21. IEEE. 1978.
- [18] Paul Hudak, and Adrienne Bloss. “The Aggregate Update Problem in Functional Programming Systems.” In *Proceedings of the 12th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, 300–314. POPL ’85. ACM, New Orleans, Louisiana, USA. 1985. doi:10.1145/318593.318660.
- [19] Daan Leijen. “Koka: Programming with Row Polymorphic Effect Types.” In *MSFP’14, 5th Workshop on Mathematically Structured Functional Programming*. 2014. doi:10.4204/EPTCS.153.8.
- [20] Daan Leijen. “Type Directed Compilation of Row-Typed Algebraic Effects.” In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL ’17)*, 486–499. Paris, France. Jan. 2017. doi:10.1145/3009837.3009872.
- [21] Daan Leijen. “The Koka Language.” 2021. <https://koka-lang.github.io>.
- [22] Daan Leijen, Zorn Ben, and Leo de Moura. “Mimalloc: Free List Sharding in Action.” *Programming Languages and Systems*, LNCS, 11893. Springer International Publishing. 2019. doi:10.1007/978-3-030-34175-6_13. APLAS’19.
- [23] Yossi Levroni, and Erez Petrank. “An On-the-Fly Reference-Counting Garbage Collector for Java.” *ACM Trans. Program. Lang. Syst.* 28 (1): 1–69. Jan. 2006. doi:10.1145/1111596.1111597.
- [24] J. McGraw, S. Skedzielewski, S. Allan, D. Grit, R. Oldehoeft, J. Glauert, I. Dobes, and P. Hohensee. *SISAL: Streams and Iteration in a Single-Assignment Language. Language Reference Manual, Version 1. 1*. LLL/M-146, ON: DE83016576. Lawrence Livermore National Lab., CA, USA. Jul. 1983.
- [25] Yasuhiko Minamide. “Space-Profiling Semantics of the Call-by-Value Lambda Calculus and the CPS Transformation.” *Electronic Notes in Theoretical Computer Science* 26: 105–120. 1999. doi:10.1016/S1571-0661(05)80286-5. HOOTS ’99, Higher Order Operational Techniques in Semantics.
- [26] Yaron Minsky, Anil Madhavapeddy, and Jason Hickey. *Real World OCaml: Functional Programming for the Masses*. 2012. <https://dev.realworldocaml.org>.
- [27] Leonardo de Moura, and Sebastian Ullrich. “The Lean 4 Theorem Prover and Programming Language.” In *Automated Deduction – CADE 28*, edited by André Platzer and Geoff Sutcliffe, 625–635. 2021.
- [28] Chris Okasaki. “Red-Black Trees in a Functional Setting.” *Journal of Functional Programming* 9 (4). Cambridge University Press: 471–477. 1999. doi:10.1017/S0956796899003494.
- [29] Chris Okasaki. *Purely Functional Data Structures*. Columbia University, New York. Jun. 1999.
- [30] Zoe Paraskevopoulou, and Andrew W Appel. “Closure Conversion Is Safe for Space.” *Proceedings of the ACM on Programming Languages* 3 (ICFP). ACM New York, NY, USA: 1–29. 2019.
- [31] Gordon D. Plotkin, and John Power. “Algebraic Operations and Generic Effects.” *Applied Categorical Structures* 11 (1): 69–94. 2003. doi:10.1023/A:1023064908962.
- [32] Gordon D. Plotkin, and Matija Pretnar. “Handlers of Algebraic Effects.” In *18th European Symposium on Programming Languages and Systems*, 80–94. ESOP’09. York, UK. Mar. 2009. doi:10.1007/978-3-642-00590-9_7.
- [33] Reinking, Xie, de Moura, and Leijen. “Perceus: Garbage Free Reference Counting with Reuse.” In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 96–111. PLDI 2021. ACK, New York, NY, USA. 2021. doi:10.1145/3453483.3454032.
- [34] Sven-Bodo Scholz. “Single Assignment C: Efficient Support for High-Level Array Operations in a Functional Setting.” *Journal of Functional Programming* 13 (6). Cambridge University Press, USA: 1005–1059. Nov. 2003. doi:10.1017/S0956796802004458.
- [35] KC Sivaramakrishnan, Stephen Dolan, Leo White, Sadiq Jaffer, Tom Kelly, Anmol Sahoo, Sudha Parimala, Atul Dhiman, and Anil Madhavapeddy. “Retrofitting Parallelism onto OCaml.” *Proc. ACM Program. Lang.* 4 (ICFP). ACM. Aug. 2020. doi:10.1145/3408995.
- [36] “The Computer Language Benchmark Game – Binarytrees.” Nov. 2021. <https://benchmarksgame-team.pages.debian.net/benchmarksgame/performance/binarytrees.html>.
- [37] David N. Turner, and Phillip Wadler. “Operational Interpretations of Linear Logic,” number 227: 231–248. 1999.

- [38] Sebastian Ullrich, and Leonardo de Moura. “Counting Immutable Beans – Reference Counting Optimized for Purely Functional Programming.” In *Proceedings of the 31st Symposium on Implementation and Application of Functional Languages (IFL’19)*. Singapore. Sep. 2019.
- [39] David Ungar, David Grove, and Hubertus Franke. “Dynamic Atomicity: Optimizing Swift Memory Management.” In *Proceedings of the 13th ACM SIGPLAN International Symposium on Dynamic Languages*, 15–26. DLS 2017. Vancouver, BC, Canada. 2017. doi:[10.1145/3133841.3133843](https://doi.org/10.1145/3133841.3133843).
- [40] Edsko de Vries, Rinus Plasmeijer, and David M. Abrahamson. “Uniqueness Typing Simplified.” In *Implementation and Application of Functional Languages (IFL’08)*, edited by Olaf Chitil, Zoltán Horváth, and Viktória Zsóka, 201–218. Springer. 2008.
- [41] Phillip Wadler. “Linear Types Can Change the World!” In *Programming Concepts and Methods*. 1990.
- [42] Andrew K. Wright, and Matthias Felleisen. “A Syntactic Approach to Type Soundness.” *Inf. Comput.* 115 (1): 38–94. Nov. 1994. doi:[10.1006/inco.1994.1093](https://doi.org/10.1006/inco.1994.1093).
- [43] Ningning Xie, and Daan Leijen. “Generalized Evidence Passing for Effect Handlers – Efficient Compilation of Effect Handlers to C.” In *Proceedings of the 26th ACM SIGPLAN International Conference on Functional Programming (ICFP’2021)*. ICFP ’21. Virtual. Aug. 2021.
- [44] Danil Yarantsev. “ORC - Nim’s Cycle Collector.” Oct. 2020. <https://nim-lang.org/blog/2020/10/15/introduction-to-arc-orc-in-nim.html>.

Appendix

A Sources

A.1 Binary Trees in Koka

This is the implementation of the `binarytrees` benchmark of the Computer Language Benchmark Game [36]:

```
import std/os/env
import std/os/task

type tree
  Node( left : tree, right : tree )
  Tip

fun make( depth : int ) : div tree
  if depth >= 1
  then Node( make(depth - 1), make(depth - 1) )
  else Node( Tip, Tip )

fun check( t : tree ) : int
  match t
  Node(l,r) -> l.check + r.check + 1
  Tip      -> 0

fun sum-count( count : int, depth : int ) : div int
  fold-int(count+1,0) fn(i,csum)
  csum + make(depth).check

fun gen-depth( min-depth : int, max-depth : int ) : pure list((int,int,promise(int)))
  list(min-depth, max-depth, 2) fn(d) // list from min to max with stride 2
  val count = 2^(max-depth + min-depth - d)
  (count, d, task{ sum-count(count, d) }) // one task per depth

fun show( msg : string, depth : int, check : int ) : console ()
  println(msg ++ " of depth " ++ depth.show ++ "\t check: " ++ check.show)

public fun main()
  val n = get-args().head.default("").parse-int.default(21)
  val min-depth = 4
  val max-depth = max(min-depth + 2, n)

  // allocate and free the stretch tree
  val stretch-depth = max-depth.inc
  show( "stretch tree", stretch-depth, make(stretch-depth).check )

  // allocate long lived tree
  val long = make(max-depth)

  // allocate and free many trees in parallel
  val trees = gen-depth( min-depth, max-depth )
  trees.foreach fn((count,depth,csum))
  show( count.show ++ "\t trees", depth, csum.await )

  // and check if the long lived tree is still good
  show( "long lived tree", max-depth, long.check )
```

A.2 Red-Black Trees with FBIP

This is the implementation of red-black tree rebalancing using the FBIP approach described in Section 4.2.1. Note how `move-up`, `balance-red`, `ins` are all nicely tail-recursive:

```
type color
  Red
  Black

type tree
  Node(color : color, lchild : tree, key : int, value : bool, rchild : tree)
  Leaf

type zipper
  NodeR(color : color, lchild : tree, key : int, value : bool, zip : zipper)
  NodeL(color : color, zip : zipper, key : int, value : bool, rchild : tree)
  Done

fun move-up(z : zipper, t : tree)
  match z
  NodeR(c, l, k, v, z1) -> z1.move-up(Node(c, l, k, v, t))
  NodeL(c, z1, k, v, r) -> z1.move-up(Node(c, t, k, v, r))
```

```

Done -> t

fun balance-red( z : zipper, l : tree, k : int, v : bool, r : tree ) : tree
match z
  NodeR(Black, l1, k1, v1, z1) -> z1.move-up( Node( Black, l1, k1, v1, Node(Red,l,k,v,r) ) )
  NodeL(Black, z1, k1, v1, r1) -> z1.move-up( Node( Black, Node(Red,l,k,v,r), k1, v1, r1 ) )
  NodeR(Red, l1, k1, v1, z1) -> match z1
    NodeR(_c2, l2, k2, v2, z2) -> z2.balance-red( Node(Black, l2, k2, v2, l1), k1, v1, Node(Black, l, k, v, r) )
    NodeL(_c2, z2, k2, v2, r2) -> z2.balance-red( Node(Black, l1, k1, v1, l), k, v, Node(Black, r, k2, v2, r2) )
    Done -> Node(Black, l1, k1, v1, Node(Red,l,k,v,r))
  NodeL(Red, z1, k1, v1, r1) -> match z1
    NodeR(_c2, l2, k2, v2, z2) -> z2.balance-red( Node(Black, l2, k2, v2, l), k, v, Node(Black, r, k1, v1, r1) )
    NodeL(_c2, z2, k2, v2, r2) -> z2.balance-red( Node(Black, l, k, v, r), k1, v1, Node(Black, r1, k2, v2, r2) )
    Done -> Node(Black, Node(Red,l,k,v,r), k1, v1, r1)
  Done -> Node(Black,l,k,v,r)

fun ins(t : tree, k : int, v : bool, z : zipper) : tree
match t
  Node(c, l, kx, vx, r)
  -> if k < kx then ins(l, k, v, NodeL(c, z, kx, vx, r))
    elif k > kx then ins(r, k, v, NodeR(c, l, kx, vx, z))
    else z.move-up(Node(c, l, kx, vx, r))
  Leaf -> z.balance-red(Leaf, k, v, Leaf)

fun insert(t : tree, k : int, v : bool) : tree
ins(t, k, v, Done)

```

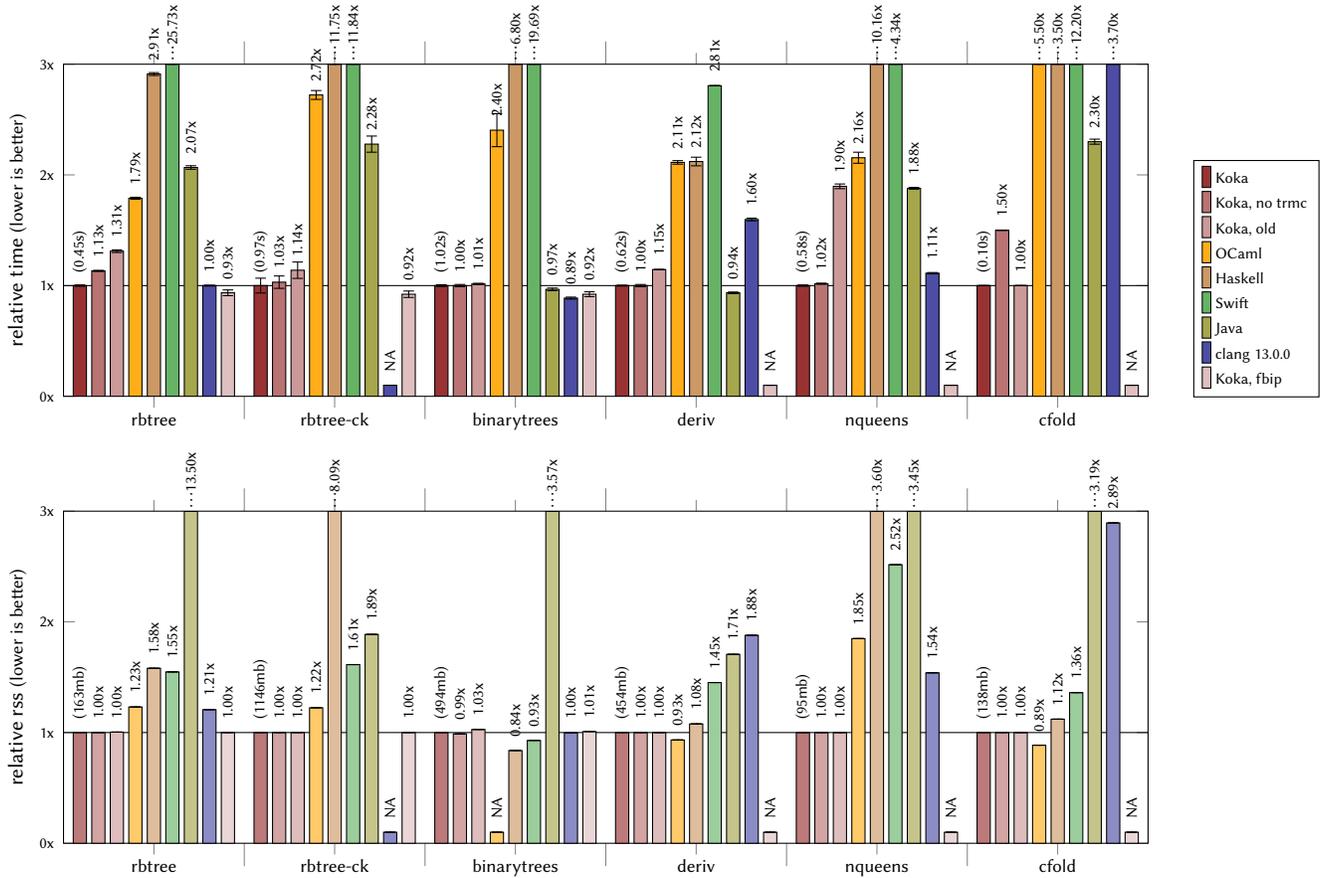


Fig. 9. Relative execution time and peak working set with respect to Koka. Using an 8-core arm64 M1 3.2Ghz (192+128KiB L1, 12MiB shared L2) with 16GiB 4266Mhz memory, macOS 11.6

B Benchmarks on an arm64 M1

Figure 9 gives the benchmark results on an arm64 M1. The relative execution times are roughly similar to that of the AMD5950X except for the `rbtree` and `binarytrees` benchmarks. In the `rbtree` benchmark, Koka now has similar performance as C++. This may be partly the use of the `clang` (vs `gcc`) which uses a slightly different implementation of red-black tree rebalancing in `std::map`². Another factor may be the larger shared L2 cache on the M1. Also, TRMC seems relatively less effective which may be due to slightly faster call frames on the M1.

In the `binarytrees` benchmark, Koka is now just 11% slower as C++ #5 and the `fbip` version is almost as fast. This may be mostly due to the concurrency being limited to 8 cores. All benchmark implementations other than the C++ one are limited in concurrency to about $\sim 10\times$ due to the `depth` steps in the benchmark. The C++ implementation also uses parallelism within the iterations per depth-step so on the 32-core AMD5950X this may account partly for the improved performance.

C Examples of worst-case heap usage

C.1 Borrowing is not frame-limited

In general, borrowing is not frame-limited. For example, the function `go` below does not use its argument `xs`, but marks it as borrowed so it needs to stay live. It then calls itself recursively with another `big-list()` for `xs`.

```
fun big-list()
  list(1, 100)

fun go( n, ^xs )
  val y = big-list()
  if(n >= 1)
    then go(n - 1, y)
  else []
```

²The LLVM `std::map` implementation can be found at https://github.com/llvm/llvm-project/blob/main/libcxx/include/__tree

```
fun main(n)
  go(n, big-list())
```

If we would ignore the borrow annotation on `xs`, we could evaluate this program with only one `big-list()` live at a time. With the borrow annotation, the program keeps up to `n` many `big-list()`s alive at a time. In a certain sense this program illustrates a worst-case: Any value we borrow must have been constructed before and so the heap usage can not increase due to a single borrowing function call by more than the maximum heap usage of the non-borrowing program. In other words, if H_{max} denotes the maximum heap size of the non-borrowing program, the heap usage of a program with borrowing is bounded at any evaluation step by H_{max} times the number of stack frames.

C.2 Algorithm D is not frame-limited

We can modify the program from the last section to show that the reuse Algorithm D [38] can use the same amount of memory as borrowing:

```
fun big-list()
  list(1, 100)

fun go( n, xs )
  val y = big-list()
  if(n >= 1)
    then match y
      Cons(.,_) ->
        go(n - 1, y)
        // drop(y) is inserted here
      Cons(1, [])
      Nil -> [] // won't happen
    else []

fun main(n)
  go(n, big-list())
```

Unlike previously, `xs` is not borrowed and so the first `big-list()` generated in `main` will be discarded. But in the subsequent iterations, `y` will be held on to at the callsite so that it can be reused for the `Cons(1, [])` allocation.

C.3 Lean's implementation of Algorithm D is not frame-limited

In practice, Lean [27] performs several transformations that make it difficult to see that the problematic program outlined above for Algorithm D also applies to their implementation. In particular:

- Lean lifts fully applied constructors like `Cons(1, [])` to a definition, so that they only have to be allocated once. We therefore ask the user for input and compute the tail `[]` based on this input. This ensures that no inlining can lead the constructor to be fully applied.
- An unused argument like `xs` for `go` is always be marked as borrowed. We print it to the console so that it is marked as owned.
- A pure function call of which we don't use the result like `go(n - 1, y)` above would be eliminated. Thanks to the modifications of the last two bullets, the function is not pure and will not get eliminated.
- The result of `big_list` is not recomputed every time we call it if we only pass `()`. We take the length of the user-provided string to compute the length of `big_list` so that the recomputation has to happen.

Keeping the above in mind, we can write the example as:

```
open List

set_option trace.compiler.ir.result true

def big_list (n : Nat) : List Nat :=
  List.range n

def print_head (xs : List Nat) :=
  match xs with
  | (cons h _) => IO.println (toString h)
  | nil => IO.println "Nil"

def get_tail : IO (List Nat × Nat) := do
  let i <- IO.getStdin
  let s <- i.getLine
  match String.toList s with
  | (cons _ _) => (cons 1 [], String.length s)
  | nil => (nil, 1)

partial def go (n : Nat) (xs : List Nat) : IO (List Nat) := do
  print_head xs
  let (t, n) <- get_tail
```

```

let y := big_list n;
if n >= 1 then
  match y with
  | (cons _ _) => do
    let _ <- go (n - 1) y
    return cons 1 t
  | nil => return t
else return t

def main : IO Unit := do
  let _ <- go 10 (big_list 10)
  IO.println "Done"

```

As of this writing (Lean version 4.0.0, commit 6475e3d5ccaf, Release) Lean emits the following IR:

```

...
let x_12 : obj := List.range x_11; // inlined big_list
...
inc x_12; // hold on to the list
let x_16 : obj := go x_15 x_12 x_9; // recursive call with (n - 1), the list and the IO state
...
let x_18 : obj := proj[1] x_12; // drop the children of the first list cell
dec x_18;
let x_19 : obj := proj[0] x_12;
dec x_19;
...
set x_12[1] := x_10; // use the list cell for the 'cons 1 t' constructor
set x_12[0] := x_13;
...

```

$$\begin{array}{l}
\text{H} : x \mapsto (\mathbb{N}^+, v) \\
\text{E} ::= \square \mid \text{E } e \mid x \text{E} \mid C x_1 \dots x_{i-1} \text{E } v_{i+1} \dots v_n \\
\quad \mid \text{let } x = \text{E in } e \mid \text{match } \text{E} \{ \overline{p_i \rightarrow e_i} \}
\end{array}
\qquad
\frac{\text{H} \mid e \longrightarrow_h \text{H}' \mid e'}{\text{H} \mid \text{E}[e] \mapsto_h \text{H}' \mid \text{E}[e']} \text{ [EVAL]}$$

Evaluation steps:

$$\begin{array}{llll}
(\text{lam}_h) & \text{H} \mid \lambda_{\bar{z}} x. e & \longrightarrow_h & \text{H}, f \mapsto^1 \lambda_{\bar{z}} x. e \mid f \quad \text{fresh } f \\
(\text{con}_h) & \text{H} \mid C x_1 \dots x_n & \longrightarrow_h & \text{H}, z \mapsto^1 C x_1 \dots x_n \mid z \quad \text{fresh } z \\
(\text{app}_h) & \text{H} \mid f y & \longrightarrow_h & \text{H} \mid \text{dup } \bar{z}; \text{drop } x; e[x:=y] \quad (f \mapsto^n \lambda_{\bar{z}} x. e) \in \text{H} \\
(\text{match}_h) & \text{H} \mid \text{match } y \{ \overline{p_i \rightarrow e_i} \} & \longrightarrow_h & \text{H} \mid e_i[\bar{x}:=\bar{z}] \quad \text{with } p_i = C \bar{x} \text{ and } (y \mapsto^n C \bar{z}) \in \text{H} \\
(\text{let}_h) & \text{H} \mid \text{let } x = z \text{ in } e & \longrightarrow_h & \text{H} \mid e[x:=z] \\
(\text{dup}_h) & \text{H} \mid \text{dup } x; e & \longrightarrow_h & \text{H} \mid e \\
(\text{drop}_h) & \text{H} \mid \text{drop } x; e & \longrightarrow_h & \text{H} \mid e \\
(\text{dropru}_h) & \text{H} \mid r \leftarrow \text{dropru } x; e & \longrightarrow_h & \text{H}, z \mapsto^1 () \mid e[r:=z] \quad \text{fresh } z
\end{array}$$

Fig. 10. Non-reference-counted heap semantics for λ^{1n} .

D Proofs

D.1 Extending strict evaluation semantics

If we add $\text{dup } e$ and $\text{drop } e$ in the syntax, as well as add to the standard semantics in Figure 2 the following rules:

$$\begin{array}{lll}
(\text{dup}) & \text{dup } e'; e & \longrightarrow e \\
(\text{drop}) & \text{drop } e'; e & \longrightarrow e \\
(\text{dropru}) & r \leftarrow \text{dropru } e'; e & \longrightarrow e
\end{array}$$

we immediately see that translations do not change evaluation:

Theorem 6. (Translation is sound)

If $e \mapsto^* v$ with $\emptyset \mid \emptyset \vdash e \rightsquigarrow e'$, then also $e' \mapsto^* v$.

Proof. (Of Theorem 6) Follows directly from Lemma 1 and the three reduction rules (dup), (drop) and (dropru). Note that the expression e can not depend on a reuse token r , since they are fresh. Therefore r will only appear as the result of a dropru and as the argument to a dup , drop or dropru . \square

Lemma 1. (Translation only inserts dup/drop)

If $\Delta \mid \Gamma \vdash e \rightsquigarrow e'$ then $e = [e']$.

Proof. (Of Lemma 1) By straightforward case analysis of each derivation. \square

D.2 Soundness of non-reference-counted heap semantics

To show that our logical rules are sound and leave no garbage, we extend the formalization of [33]. Lemma 9 and 10 of their proof rely on the heap evaluation context being not too different from the standard evaluation context. Unfortunately, this is not the case in our formalization: Our heap evaluation context only works for normalized terms and evaluation may make normalized terms only partially normalized. We also can not just extend the evaluation context for the heap semantics, because the match rule relies on the scrutinee being a variable. We will therefore first show that a heap semantics that ignores dups, drops and drop-reuses (Figure 10) will evaluate an expression to the same value as the standard strict evaluation.

Lemma 2.

Let $\text{H} \mid e \mapsto_h \text{H}' \mid e'$ be a heap derivation and z a fresh variable. Then $\text{H}, z \mapsto^n v \mid e \mapsto_h \text{H}', z \mapsto^n v \mid e'$.

Proof. By straightforward case analysis on the evaluation rules.

Lemma 3.

Let H be a heap and v a value. Then $\text{H} \mid v \mapsto_h^* \text{H}' \mid x$ with $[\text{H}']x = [\text{H}]v$.

Proof. By induction on v .

case x .

$$\begin{array}{ll}
\text{H} \mid x \mapsto_h^* \text{H} \mid x & (1), \text{ given} \\
[\text{H}]x = [\text{H}]x & (2), \text{ trivial}
\end{array}$$

case $[\lambda x. e]$.

$$\begin{aligned} H \mid \lambda x. e \mapsto^*_h H, f \mapsto^1 \lambda x. e \mid f & \quad (1), \text{ by } (lam_h) \\ [H, f \mapsto^1 \lambda x. e]f = [H](\lambda x. e) & \quad (2), \text{ by } (1) \end{aligned}$$

case $[C \bar{v}]$.

$$\begin{aligned} H \mid C \bar{v} \mapsto^*_h H, \bar{x} \mapsto^1 \bar{v} \mid C \bar{x} & \quad (1), \text{ by the evaluation context} \\ [H]\bar{v} = [H, \bar{x} \mapsto^1 \bar{v}]\bar{x} & \quad (2), \text{ by induction} \\ H, \bar{x} \mapsto^1 \bar{v} \mid C \bar{x} \mapsto^*_h H, \bar{x} \mapsto^1 \bar{v}, x \mapsto^1 C \bar{x} \mid x & \quad (3), \text{ by } (con_h) \\ [H, \bar{x} \mapsto^1 \bar{v}, x \mapsto^1 C \bar{x}]x = [H](C \bar{v}) & \quad (4), \text{ by } (2) \end{aligned}$$

Lemma 4.

Let H be a heap and e an expression. If $[H]e \mapsto^* v$, then $H \mid e \mapsto^*_h H' \mid x$ with $[H']x = v$.

Proof. By induction on the judgment $[H]e \mapsto^* v$, where \mapsto^* is the reflexive transitive closure of \mapsto .

case $[[H]v \mapsto^* [H]v]$ (reflexive case).

$$\begin{aligned} [H]v \mapsto^* [H]v & \quad (1), \text{ given} \\ H \mid v \mapsto^*_h H' \mid x & \quad (2), \text{ by lemma 3} \\ [H']x = [H]v & \quad (3), \text{ by lemma 3} \end{aligned}$$

case $[[H]e \mapsto [H]e' \text{ and } [H]e' \mapsto^* v]$ (transitive case).

We proceed with induction on the judgment $[H]e \mapsto [H]e'$.

case $[E = \square \text{ and } e = (\lambda x. \underline{e}) v]$.

$$\begin{aligned} e' = \underline{e}[x:=v] & \quad (1), \text{ given} \\ H \mid (\lambda x. \underline{e}) v \mapsto_h H, f \mapsto^1 \lambda x. \underline{e} \mid f v & \quad (2), \text{ by } (lam_h) \\ H, f \mapsto^1 \lambda x. \underline{e} \mid f v \mapsto_h H', f \mapsto^1 \lambda x. \underline{e} \mid f z & \quad (3), \text{ by lemma 3} \\ H', f \mapsto^1 \lambda x. \underline{e} \mid f z \mapsto^*_h H', f \mapsto^1 \lambda x. \underline{e} \mid \underline{e}[x:=z] & \quad (4), \text{ by } (app_h), (dup_h), (drop_h) \\ [H]v = [H']z & \quad (5), \text{ by lemma 3} \\ [H]e' = [H, z \mapsto^1 v]\underline{e}[x:=z] = [H']\underline{e}[x:=z] & \quad (6), \text{ by } (1) \text{ and } (5) \\ [H]e' = [H', f \mapsto^1 \lambda x. \underline{e}]\underline{e}[x:=z] & \quad (7), \text{ by monotonicity} \end{aligned}$$

case $[E = \square \text{ and } e = \text{let } x = v \text{ in } \underline{e}]$.

$$\begin{aligned} e' = \underline{e}[x:=v] & \quad (1), \text{ given} \\ H \mid \text{let } x = v \text{ in } \underline{e} \mapsto_h H' \mid \text{let } x = z \text{ in } \underline{e} & \quad (2), \text{ by lemma 3} \\ H' \mid \text{let } x = z \text{ in } \underline{e} \mapsto_h H' \mid \underline{e}[x:=z] & \quad (3), \text{ by } (let_h) \\ [H]v = [H']z & \quad (4), \text{ by lemma 3} \\ [H]e' = [H, z \mapsto^1 v]\underline{e}[x:=z] = [H']\underline{e}[x:=z] & \quad (5), \text{ by } (1) \text{ and } (4) \end{aligned}$$

case $[E = \square \text{ and } e = \text{match } (C \bar{v}) \{ \overline{p_i \rightarrow e_i} \}]$.

$$\begin{aligned} e' = e_i[\bar{x}:=\bar{v}] & \quad (1), \text{ given where } p_i = C \bar{x} \\ H \mid \text{match } (C \bar{v}) \{ \overline{p_i \rightarrow e_i} \} \mapsto_h H' \mid \text{match } z \{ \overline{p_i \rightarrow e_i} \} & \quad (2), \text{ by lemma 3} \\ [H](C \bar{v}) = [H']z & \quad (3), \text{ by lemma 3} \\ (z \mapsto^n C \bar{z}) \in H' & \quad (4), \text{ by } (3) \\ H' \mid \text{match } z \{ \overline{p_i \rightarrow e_i} \} \mapsto_h H' \mid e_i[\bar{x}:=\bar{z}] & \quad (5), \text{ by } (match_h) \text{ and } (4) \\ [H]e' = [H, \bar{z} \mapsto^1 \bar{v}]e_i[\bar{x}:=\bar{z}] = [H']e_i[\bar{x}:=\bar{z}] & \quad (6), \text{ by } (1) \text{ and } (3) \end{aligned}$$

case $[E = \square \text{ and } e = \text{dup } x; e_2]$.

$$\begin{aligned} e' = e_2 & \quad (1), \text{ given} \\ [H]e' \mapsto^* v & \quad (2), \text{ given} \\ H \mid e' \mapsto^*_h H' \mid x & \quad (3), \text{ by induction} \\ [H']x = v & \quad (4), \text{ by induction} \end{aligned}$$

case $[E = \square \text{ and } e = \text{drop } x; e_2]$.

$$\begin{aligned} e' = e_2 & \quad (1), \text{ given} \\ [H]e' \mapsto^* v & \quad (2), \text{ given} \\ H \mid e' \mapsto^*_h H' \mid x & \quad (3), \text{ by induction} \\ [H']x = v & \quad (4), \text{ by induction} \end{aligned}$$

case $[E = \square \text{ and } e = r \leftarrow \text{dropru } x; e_2]$.

$e' = e_2$ (1), given
 $[H]e' \mapsto^* v$ (2), given
 $H \mid e' \mapsto^* H' \mid x$ (3), by induction
 $[H']x = v$ (4), by induction
 $H, z \mapsto^1 () \mid e' \mapsto^* H', z \mapsto^1 () \mid x$ (5), by (3) as z fresh
 $[H', z \mapsto^1 ()]x = v$ (6), by (4) and lemma 2

case $[E = E_1 e'' \text{ and } e = \underline{e} e'']$.

$e' = w e''$ (1), given
 $[H]\underline{e} \mapsto^* w$ (2), given
 $H \mid \underline{e} \mapsto^*_h H' \mid x$ (3), by induction
 $[H]w = [H']x$ (4), by induction
 $[H]e' = [H'](x e'')$ (5), by (4)
 $[H]e' \mapsto^* v$ (6), given
 $[H'](x e'') \mapsto^* v$ (7), by (5)
 $H' \mid (x e'') \mapsto^*_h H'' \mid y$ (8), by induction
 $[H]v = [H'']y$ (9), by induction
 $H \mid (\underline{e} e'') \mapsto^*_h H'' \mid y$ (10), by (3) and (8)
 $H \mid e \mapsto^*_h H'' \mid y$ (11), by (10)

case $[E = v E_1 \text{ and } e = v e'']$.

$e' = v w$ (1), given
 $H \mid v \mapsto^*_h H' \mid x$ (2), by lemma 3
 $[H']x = [H]v$ (3), by lemma 3
 $[H']w = [H]w$ (4), since new names are fresh
 $H' \mid e'' \mapsto^*_h H'' \mid y$ (5), by induction
 $[H'']y = [H]w$ (6), by induction and (4)
 $[H](v w) \mapsto^* u$ (7), given
 $[H''](x y) \mapsto^* u$ (8), by (7),(3),(6)
 $H'' \mid x y \mapsto^*_h H_3 \mid z$ (9), by induction
 $[H_3]z = u$ (10), by induction
 $H \mid v e'' \mapsto^*_h H_3 \mid z$ (11), by (2),(5),(9)

case $[E = \text{let } x = E_1 \text{ in } e_3 \text{ and } e = \text{let } x = e_2 \text{ in } e_3]$.

$[H]e_2 \mapsto^* v$ (1), given
 $e' = \text{let } x = v \text{ in } e_3$ (2), given
 $H \mid e_2 \mapsto^* H' \mid y$ (3), by induction
 $[H']y = v$ (4), by induction
 $[H]e' \mapsto^* w$ (5), by induction
 $H' \mid (\text{let } x = y \text{ in } e_3) \mapsto^* H'' \mid z$ (6), by induction
 $[H'']z = w$ (7), by induction
 $H \mid (\text{let } x = e_2 \text{ in } e_3) \mapsto^* H'' \mid z$ (8), by (3),(4)

case $[E = \text{match } E_1 \{\overline{p_i \rightarrow e_i}\} \text{ and } e = \text{match } e_2 \{\overline{p_i \rightarrow e_i}\}]$.

$[H]e_2 \mapsto^* v$ (1), given
 $e' = \text{match } v \{\overline{p_i \rightarrow e_i}\}$ (2), given
 $H \mid e_2 \mapsto^* H' \mid y$ (3), by induction
 $[H']y = v$ (4), by induction
 $[H]e' \mapsto^* w$ (5), by induction
 $H' \mid (\text{match } v \{\overline{p_i \rightarrow e_i}\}) \mapsto^* H'' \mid z$ (6), by induction
 $[H'']z = w$ (7), by induction
 $H \mid (\text{match } e_2 \{\overline{p_i \rightarrow e_i}\}) \mapsto^* H'' \mid z$ (8), by (3),(4)

D.3 A Heap Reference Counting Calculus

$$\begin{array}{c}
\frac{}{H, x \mapsto^{n+1} v \vdash x \dashv H, x \mapsto^n v} \text{ [DRVAR]} \\
\\
\frac{H \vdash \bar{y} \dashv H_1}{H, x \mapsto^1 \lambda_{\bar{y}} z. e \vdash x \dashv H_1} \text{ [DRVARLAM]} \\
\\
\frac{H \vdash \bar{y} \dashv H_1}{H, x \mapsto^1 C \bar{y} \vdash x \dashv H_1} \text{ [DRVARCON]} \\
\\
\frac{H, x \mapsto^{n+1} v \vdash e \dashv H_1}{H, x \mapsto^n v \vdash \text{dup } x; e \dashv H_1} \text{ [DRDUP]} \\
\\
\frac{H, x \mapsto^n v \vdash e \dashv H_1}{H, x \mapsto^{n+1} v \vdash \text{drop } x; e \dashv H_1} \text{ [DRDROP]} \\
\\
\frac{H \vdash \text{drop } \bar{y}; e \dashv H_1}{H, x \mapsto^1 C \bar{y} \vdash \text{drop } x; e \dashv H_1} \text{ [DRDROPCON]} \\
\\
\frac{H \vdash \text{drop } \bar{y}; e \dashv H_1}{H, x \mapsto^1 \lambda_{\bar{y}} z. e \vdash \text{drop } x; e \dashv H_1} \text{ [DRDROPLAM]} \\
\\
\frac{H, x \mapsto^n v, r \mapsto^1 () \vdash e \dashv H_1}{H, x \mapsto^{n+1} v \vdash r \leftarrow \text{dropru } x; e \dashv H_1} \text{ [DRDROP RU]} \\
\\
\frac{H, r \mapsto^1 () \vdash \text{drop } \bar{y}; e \dashv H_1}{H, x \mapsto^1 C \bar{y} \vdash r \leftarrow \text{dropru } x; e \dashv H_1} \text{ [DRDROP CONRU]} \\
\\
\frac{H, r \mapsto^1 () \vdash \text{drop } \bar{y}; e \dashv H_1}{H, x \mapsto^1 \lambda_{\bar{y}} z. e \vdash r \leftarrow \text{dropru } x; e \dashv H_1} \text{ [DRDROPLAMRU]} \\
\\
\frac{H \vdash x_1 \dashv H_1 \dots H_{n-1} \vdash x_n \dashv H_n}{H \vdash C x_1 \dots x_n \dashv H_n} \text{ [DRCON]} \\
\\
\frac{H \vdash \bar{y} \dashv H_1 \quad \bar{y} \mapsto^1 (), x \mapsto^1 () \vdash e \dashv \emptyset}{H \vdash \lambda_{\bar{y}} x. e \dashv H_1} \text{ [DRLAM]} \\
\\
\frac{H \vdash e \dashv H_1 \quad H_1 \vdash x \dashv H_2}{H \vdash e x \dashv H_2} \text{ [DRAPP]} \\
\\
\frac{H \vdash e_1 \dashv H_1 \quad H_1, x \mapsto^1 () \vdash e_2 \dashv H_2 \quad x \notin H, H_2 \quad (\star)}{H \vdash \text{val } x = e_1; e_2 \dashv H_2} \text{ [DRBIND]} \\
\\
\frac{H, [\bar{z}_i] \vdash e_i \dashv H' \quad \bar{z}_i \notin H, H' \quad x \in H}{H \vdash \text{match } x \{ p_i \mapsto \text{dup}(\bar{z}_i); e_i \} \dashv H'} \text{ [DRMATCH]}
\end{array}$$

For the next sections we will assume that $(\star) = \text{true}$. Note however, that this is only relevant in lemma 12 and lemma 14. All other lemmas work for any choice of (\star) (even false) as they either don't need the DRBIND rule or already assume (\star) as a premise.

Lemma 5. (*Heap Reference Counting Free variables*)

If $H_1 \vdash e \dashv H_2$, then $\text{fv}(e) \in H_1$, and $\text{fv}(H_2) \in H_1$ with same domains.

Proof. (*Of Lemma 5*) By a straightforward induction on the rules. \square

Definition 4. (*Extension*)

H is extended with x , denoted as $H \# x$, where $\#$ works as follows:

- (1) if $H = H', x \mapsto^n v$, then $H \# x = H', x \mapsto^{n+1} v$;
- (2) if $x \notin H$, then $H \# x = (H, x \mapsto^1 v) \# \text{fv}(v)$.

We omit the domain of x in $H \# x$ for simplicity. The domain should always be available by inspecting the heap (in (1)) or via explicit passing (in (2)).

We only focus on situations where there is no cycles in the dependency of x (but we are fine with existing cycles in H), so that the extension terminates. That implies $(H, x \mapsto^1 v) \# \text{fv}(v) = H \# \text{fv}(v), x \mapsto^1 v$ in (2).

Lemma 6. (*Drop is dual to extension*)

If $H_1 \vdash \text{drop } x; () \dashv H_2$, then $H_1 = H_2 \# x$. Similarly, if $H_1 \vdash x \dashv H_2$, then $H_1 = H_2 \# x$.

Proof. (*Of Lemma 6*) By induction on the judgment.

case

$$H, x \mapsto^{n+1} v \vdash \text{drop } x; () \dashv H, x \mapsto^n v \quad \text{DRDROP, DRCON}$$

case

$$\begin{array}{ll} H, x \mapsto^1 \lambda_{\bar{y}} z. e \vdash \text{drop } x; () \dashv H_1 & \text{given} \\ H \vdash \text{drop } \bar{y}; () \dashv H_1 & \text{DRDROPLAM} \\ H = H_1 \# \bar{y} & \text{I.H.} \\ H, x \mapsto^1 \lambda_{\bar{y}} z. e = H_1 \# x & \text{by definition} \end{array}$$

case

$$\begin{array}{ll} H, x \mapsto^1 C \bar{y} \vdash \text{drop } x; () \dashv H_1 & \text{given} \\ H \vdash \text{drop } \bar{y}; () \dashv H_1 & \text{DRDROPCON} \\ H = H_1 \# \bar{y} & \text{I.H.} \\ H, x \mapsto^1 C \bar{y} = H_1 \# x & \text{by definition} \end{array}$$

□

Lemma 7. (*Extension is dual to drop*)

$H \# x \vdash \text{drop } x; () \dashv H$. Similarly, $H \# x \vdash x \dashv H$.

Proof. (*Of Lemma 7*) By induction on $\#x$.

case

$$\begin{array}{ll} H = H', x \mapsto^n \bar{y} & \text{if} \\ H \# x = H', x \mapsto^{n+1} \bar{y} & \text{by definition} \\ H', x \mapsto^{n+1} \bar{y} \vdash \text{drop } x; () \dashv H & \text{DRDROP} \end{array}$$

case

$$\begin{array}{ll} x \notin H & \text{if} \\ \bar{y} = \text{fv}(v) & \text{let} \\ H \# x = H \# \bar{y}, x \mapsto^1 v & \text{by definition} \\ H \# \bar{y} \vdash \bar{y} \dashv H & \text{I.H.} \\ H \# \bar{y}, x \mapsto^1 v \vdash x \dashv H & \text{DRVARLAM OR DRVARCON} \end{array}$$

□

Lemma 8. (*Extension Commutativity*)

$H \# x \# y = H \# y \# x$.

Proof. (*Of Lemma 8*) By induction on $\#x$ and $\#y$, then we do case analysis. **case** $x \in H$. Then $H = H', x \mapsto^n v$.

By definition, $H \# x \# y = (H', x \mapsto^{n+1} v) \# y$. Since the way $\#$ works only depends on whether x exists but not the exact number of its occurrence, we can decrease the number of x , do $\#y$ and then add x back. That is, $(H', x \mapsto^{n+1} v) \# y = (H', x \mapsto^n v) \# y \# x = H \# y \# x$.

case $y \in H$ is similar as the previous case.

case $x, y \notin H$. Then $H \# x = H \# \bar{x}, x \mapsto^1 v$ where $\bar{x} = \text{fv}(v)$.

subcase Assume $\#y$ won't cause $\#x$, then $\#y$ doesn't care about the existence of x .

So $(H \# \bar{x}, x \mapsto^1 v) \# y = H \# \bar{x} \# y, x \mapsto^1 v$
 $= H \# y \# \bar{x}, x \mapsto^1 v$ by I.H.

$= H \# y \# x$ by definition.

subcase Or otherwise $\#y$ will cause $\#x$. Since there is no cycle in the dependency, that means $\#x$ won't cause $\#y$. Then we can prove it as in the previous case. \square

Lemma 9. (*Dropru is dual to extension*)

If $H_1 \vdash r \leftarrow \text{dropru } x; () \dashv H_2$, then $H_1 \# r = H_2 \# x$.

Proof. (*Of Lemma 9*) **case**

$H, x \mapsto^{n+1} v \vdash r \leftarrow \text{dropru } x; () \dashv H, x \mapsto^n v, r \mapsto^1 ()$ DRDROPRU
 $H, x \mapsto^{n+1} v \# r = H, x \mapsto^n v, r \mapsto^1 () \# x$ Lemma 8
case

$H, x \mapsto^1 \lambda_{\bar{y}} z. e \vdash r \leftarrow \text{dropru } x; () \dashv H_1, r \mapsto^1 ()$ given
 $H \vdash \text{drop } \bar{y}; () \dashv H_1$ DRDROPLAMRU
 $H = H_1 \# \bar{y}$ Lemma 6
 $H, x \mapsto^1 \lambda_{\bar{y}} z. e = H_1 \# x$ by definition
 $H, x \mapsto^1 \lambda_{\bar{y}} z. e \# r = H_1, r \mapsto^1 () \# x$ Lemma 8
case

$H, x \mapsto^1 C \bar{y} \vdash r \leftarrow \text{dropru } x; () \dashv H_1, r \mapsto^1 ()$ given
 $H \vdash \text{drop } \bar{y}; () \dashv H_1$ DRDROPCONRU
 $H \# r = H_1 \# \bar{y}$ Lemma 6
 $H, x \mapsto^1 C \bar{y} = H_1 \# x$ by definition
 $H, x \mapsto^1 C \bar{y} \# r = H_1, r \mapsto^1 () \# x$ Lemma 8
 \square

Lemma 10. (*Extension is dual to dropru*)

$H \# x \vdash r \leftarrow \text{dropru } x; () \dashv H \# r$.

Proof. (*Of Lemma 10*) By induction on $\#x$.

case

$H = H', x \mapsto^n \bar{y}$ if
 $H \# x = H', x \mapsto^{n+1} \bar{y}$ by definition
 $H', x \mapsto^{n+1} \bar{y} \vdash r \leftarrow \text{dropru } x; () \dashv H, r \mapsto^1 ()$ DRDROPRU
case

$x \notin H$ if
 $\bar{y} = \text{fv}(v)$ let
 $H \# x = H \# \bar{y}, x \mapsto^1 v$ by definition
 $H \# \bar{y} \vdash \bar{y} \dashv H$ I.H.
 $H \# \bar{y}, x \mapsto^1 v \vdash r \leftarrow \text{dropru } x; () \dashv H, r \mapsto^1 ()$ DRDROPLAMRU OR DRDROPCONRU
 \square

D.3.1 Relating to linear resource calculus.

Definition 5. (*Context to Dependency Heap*)

Given a context Γ , $\llbracket \Gamma \rrbracket$ defines a dependency heap, with all x becoming $x \mapsto^n ()$ if x appears n times in Γ .

Lemma 11.

$\llbracket \Gamma, x \rrbracket \vdash x \dashv \llbracket \Gamma \rrbracket$.

Similarly, if $\llbracket \Gamma \rrbracket \vdash e \dashv H$, then $\llbracket \Gamma, x \rrbracket \vdash \text{drop } x; e \dashv H$.

Similarly, if $\llbracket \Gamma, r \rrbracket \vdash e \dashv H$, then $\llbracket \Gamma, x \rrbracket \vdash r \leftarrow \text{dropru } x; e \dashv H$.

Proof. The goal holds by rule `DRVAR` (`DRDROP`, `DRDROPRU`) when $x \in \Gamma$ or by rule `DRVARCON` (`DRDROPCON`, `DRDROPCONRU`) if $x \notin \Gamma$.
□

Lemma 12. (*linear resource calculus relates to reference counting*)

If $\Delta \mid \Gamma \vdash e \rightsquigarrow e'$, then $\llbracket \Delta, \Gamma \rrbracket \vdash e' \dashv \llbracket \Delta \rrbracket$.

Proof. (Of Lemma 12) By induction on the elaboration.

case

$\Delta \mid x \vdash x \rightsquigarrow x$ given
 $\llbracket \Delta, x \rrbracket \vdash x \dashv \llbracket \Delta \rrbracket$ Lemma 11
case

$\Delta \mid \Gamma \vdash e \rightsquigarrow \text{dup } x; e'$ given
 $\Delta \mid \Gamma, x \vdash e \rightsquigarrow e'$ given
 $x \in \Delta, \Gamma$ given
 $\llbracket \Delta, \Gamma, x \rrbracket \vdash e' \dashv \llbracket \Delta \rrbracket$ I.H.
 $\llbracket \Delta, \Gamma \rrbracket \vdash \text{dup } x; e' \dashv \llbracket \Delta \rrbracket$ DRDUP
case

$\Delta \mid \Gamma, x \vdash e \rightsquigarrow \text{drop } x; e'$ given
 $\Delta \mid \Gamma \vdash e \rightsquigarrow e'$ given
 $\llbracket \Delta, \Gamma \rrbracket \vdash e' \dashv \llbracket \Delta \rrbracket$ I.H.
 $\llbracket \Delta, \Gamma, x \rrbracket \vdash \text{drop } x; e' \dashv \llbracket \Delta \rrbracket$ Lemma 11
case

$\Delta \mid \Gamma, x \vdash e \rightsquigarrow r \leftarrow \text{dropru } x; e'$ given
 $\Delta \mid \Gamma, r \vdash e \rightsquigarrow e'$ given
 $\llbracket \Delta, \Gamma, r \rrbracket \vdash e' \dashv \llbracket \Delta \rrbracket$ I.H.
 $\llbracket \Delta, \Gamma, x \rrbracket \vdash r \leftarrow \text{dropru } x; e' \dashv \llbracket \Delta \rrbracket$ Lemma 11
case

$\Delta \mid \bar{y} \vdash \lambda x. e \rightsquigarrow \lambda_{\bar{y}} x. e'$ given
 $\emptyset \mid \bar{y}, x \vdash e \rightsquigarrow e'$ given
 $\bar{y} = \text{fv}(\lambda x. e)$ given
 $\llbracket \Delta, \bar{y} \rrbracket \vdash \bar{y} \dashv \llbracket \Delta \rrbracket$ Lemma 11
 $\llbracket \bar{y}, x \rrbracket \vdash e' \dashv \emptyset$ I.H.
 $\llbracket \Delta, \bar{y} \rrbracket \vdash \lambda_{\bar{y}} x. e \dashv \llbracket \Delta \rrbracket$ DRLAM
case

$\Delta \mid \Gamma \vdash e_1 x \rightsquigarrow e'_1 x$ given
 $\Delta, x \mid \Gamma \vdash e_1 \rightsquigarrow e'_1$ given
 $\llbracket \Delta, x, \Gamma \rrbracket \vdash e'_1 \dashv \llbracket \Delta, x \rrbracket$ I.H.
 $\llbracket \Delta, x \rrbracket \vdash x \dashv \llbracket \Delta \rrbracket$ Lemma 11
 $\llbracket \Delta, \Gamma \rrbracket \vdash e'_1 x \dashv \llbracket \Delta \rrbracket$ DRAPP
case

$\Delta \mid \Gamma_1, \Gamma_2 \vdash \text{val } x = e_1; e_2 \rightsquigarrow \text{val } x = e'_1; e'_2$ given
 $\Delta, \Gamma_2 \mid \Gamma_1 \vdash e_1 \rightsquigarrow e'_1$ given
 $\Delta \mid \Gamma_2, x \vdash e_2 \rightsquigarrow e'_2$ given
 $x \notin \Delta, \Gamma_1, \Gamma_2$ given
 $\llbracket \Delta, \Gamma_1, \Gamma_2 \rrbracket \vdash e'_1 \dashv \llbracket \Delta, \Gamma_2 \rrbracket$ I.H.
 $\llbracket \Delta, \Gamma_2, x \rrbracket \vdash e'_2 \dashv \llbracket \Delta \rrbracket$ I.H.
 $x \notin \llbracket \Delta \rrbracket$ follows
 $\llbracket \Delta, \Gamma_1, \Gamma_2 \rrbracket \vdash \text{val } x = e'_1; e'_2 \dashv \llbracket \Delta \rrbracket$ DRBIND

case

$$\begin{array}{l}
\Delta \mid \Gamma \vdash \text{match } x \{ \overline{p_i \mapsto e_i} \} \rightsquigarrow \text{match } x \{ \overline{p_i \mapsto \text{dup}(\overline{z_i}); e'_i} \} \quad \text{given} \\
x \in \Delta, \Gamma \quad \text{given} \\
\Delta \mid \Gamma, \overline{z_i} \vdash e_i \rightsquigarrow e'_i \quad \text{given} \\
[[\Delta, \Gamma, \overline{z_i}] \vdash e'_i \dashv \llbracket \Delta \rrbracket] \quad \text{I.H.} \\
\hline
[[\Delta, \Gamma] \vdash \text{match } x \{ \overline{p_i \mapsto \text{dup}(\overline{z_i}); e'_i} \} \dashv \llbracket \Delta \rrbracket] \quad \text{DRMATCH} \\
\text{case}
\end{array}$$

$$\begin{array}{l}
\Delta \mid \overline{x} \vdash C \overline{x} \rightsquigarrow C \overline{x} \quad \text{given} \\
[[\Delta, \overline{x}] \vdash \overline{x} \dashv \llbracket \Delta \rrbracket] \quad \text{Lemma 11} \\
[[\Delta, \overline{x}] \vdash C \overline{x} \dashv \llbracket \Delta \rrbracket] \quad \text{DRCON} \\
\quad \square
\end{array}$$

D.3.2 Weakening.

Lemma 13. (Weakening)

If $H_1 \vdash e \dashv H_2$, then $H_1 \# x \vdash e \dashv H_2 \# x$.

Proof. (Of Lemma 13) By induction on the judgment.

case

$$\begin{array}{l}
H_0 \vdash C \overline{x} \dashv H_n \quad \text{given} \\
H_{i-1} \vdash x_i \dashv H_i \quad \text{DRCON} \\
H_{i-1} \# x \vdash x_i \dashv H_i \# x \quad \text{I.H.} \\
H_0 \# x \vdash C \overline{x} \dashv H_n \# x \quad \text{DRCON} \\
\text{case}
\end{array}$$

$$\begin{array}{l}
H \vdash y \dashv H_2 \quad \text{given} \\
H = H_2 \# y \quad \text{Lemma 6} \\
H \# x = H_2 \# y \# x \\
= H_2 \# x \# y \quad \text{Lemma 8} \\
H_2 \# x \# y \vdash y \dashv H_2 \# x \quad \text{Lemma 7}
\end{array}$$

case

$$\begin{array}{l}
H, y \mapsto^n \overline{y} \vdash \text{dup } y; e \dashv H_1 \quad \text{given} \\
H, y \mapsto^{n+1} \overline{y} \vdash e \dashv H_1 \quad \text{DRDUP} \\
(H, y \mapsto^n \overline{y}) \# y \vdash e \dashv H_1 \quad \text{definition of } \# \\
(H, y \mapsto^n \overline{y}) \# y \# x \vdash e \dashv H_1 \# x \quad \text{I.H.} \\
(H, y \mapsto^n \overline{y}) \# y \# x \\
= (H, y \mapsto^n \overline{y}) \# x \# y \quad \text{Lemma 8} \\
(H, y \mapsto^n \overline{y}) \# x \# y \vdash e \dashv H_1 \# x \quad \text{By substitution} \\
(H, y \mapsto^n \overline{y}) \# x \vdash \text{dup } y; e \dashv H_1 \# x \quad \text{DRDUP} \\
\text{case}
\end{array}$$

$$\begin{array}{l}
H \vdash \text{drop } y; e \dashv H_2 \quad \text{given} \\
H \vdash \text{drop } y; () \dashv H_3 \quad \text{follows} \\
H_3 \vdash e \dashv H_2 \quad \text{above} \\
H = H_3 \# y \quad \text{Lemma 6} \\
H \# x = H_3 \# y \# x \\
= H_3 \# x \# y \quad \text{Lemma 8} \\
H_3 \# x \# y \vdash \text{drop } y; () \dashv H_3 \# x \quad \text{Lemma 7} \\
H_3 \# x \vdash e \dashv H_2 \# x \quad \text{I.H.} \\
H_3 \# x \# y \vdash \text{drop } y; e \dashv H_2 \# x \quad \text{Follows} \\
H \# x \vdash \text{drop } y; e \dashv H_2 \# x \quad \text{By substitution}
\end{array}$$

case

$H \vdash r \leftarrow \text{dropru } y; e \dashv H_2$	given
$H \vdash \text{drop } y; () \dashv H_3$	follows
$H_3 \# r \vdash e \dashv H_2$	above
$H = H_3 \# y$	Lemma 6
$H \# x = H_3 \# y \# x$	
$= H_3 \# x \# y$	Lemma 8
$H_3 \# x \# y \vdash r \leftarrow \text{dropru } y; () \dashv H_3 \# x \# r$	Lemma 10
$H_3 \# r \# x \vdash e \dashv H_2 \# x$	I.H.
$H_3 \# x \# r = H_3 \# r \# x$	Lemma 8
$H_3 \# x \# y \vdash r \leftarrow \text{dropru } y; e \dashv H_2 \# x$	Follows
$H \# x \vdash r \leftarrow \text{dropru } y; e \dashv H_2 \# x$	By substitution

case

$H \vdash \lambda_{\bar{y}} z. e \dashv H_1$	given
$H \vdash \bar{y} \dashv H_1$	given
$\bar{y} \mapsto^1 (), z \mapsto^1 () \vdash e \dashv \emptyset$	given
$H \# x \vdash \bar{y} \dashv H_1 \# x$	I.H.
$H \# x \vdash \lambda_{\bar{y}} z. e \dashv H_1 \# x$	DRLAM

case

$H \vdash e_1 e_2 \dashv H_2$	given
$H \vdash e_1 \dashv H_1$	given
$H_1 \vdash e_2 \dashv H_2$	given
$H \# x \vdash e_1 \dashv H_1 \# x$	I.H.
$H_1 \# x \vdash e_2 \dashv H_2 \# x$	I.H.
$H \# x \vdash e_1 e_2 \dashv H_2 \# x$	DRAPP

case

$H \vdash \text{val } z = e_1 ; e_2 \dashv H_2$	given
$H \vdash e_1 \dashv H_1$	given
$H_1, z \mapsto^1 () \vdash e_2 \dashv H_2$	given
$z \notin H$	given
$H \# x \vdash e_1 \dashv H_1 \# x$	I.H.
$(H_1, z \mapsto^1 ()) \# x \vdash e_2 \dashv H_2 \# x$	I.H.
$z \notin H_1$	Lemma 5
$H_1, z \mapsto^1 () = H_1 \# z$	
$(H_1, z \mapsto^1 ()) \# x = H_1 \# z \# x$	
$= H_1 \# x \# z$	Lemma 8
$= H_1 \# x, z \mapsto^1 ()$	
$H_1 \# x, z \mapsto^1 () \vdash e_2 \dashv H_2 \# x$	by substitution
$H \# x \vdash e_1 e_2 \dashv H_2 \# x$	DRBIND

case

$H \vdash \text{match } z \{ \overline{p_i \mapsto \text{dup}(\overline{z_i}); e_i} \} \dashv H'$	given
$z \in H, z \in H \# x$	given
$H, H_i \vdash e_i \dashv H'$	given
$H_i = \llbracket \overline{z_i} \rrbracket$	given
$(H, H_i) \# x \vdash e_i \dashv H' \# x$	I.H.
$\overline{z_i}$ fresh	assume
$H, \llbracket H_i \rrbracket = H \# \overline{z_i}$	
$(H, H_i) \# x = H \# \overline{z_i} \# x$	
$= H \# x \# \overline{z_i}$	Lemma 8
$= H \# x, H_i$	
$H \# x, H_i \vdash e_i \dashv H' \# x$	by substitution
$H \# x \vdash \text{match } z \{ \overline{p_i \mapsto \text{dup}(\overline{z_i}); e_i} \} \dashv H' \# x$	DRMATCH
□	

D.3.3 Substitution.

Lemma 14. (*Substitution preserves dependencies*)

If $y, z \notin H_1, H_2$ and $H_1, \llbracket y \rrbracket \vdash e \dashv H_2$, then $H_1 \# z \vdash e[y:=z] \dashv H_2$.

Proof. (*Of Lemma 14*) By induction on the judgment. The cases DRVAR, DRVARLAM, DRVARCON, DRDROP, DRDROPCON, DRDROPLAM, DRDROPRU, DRDROPCONRU, and DRDROPLAMRU follow directly from lemma 6 (resp. 9) if $y = x$. If $y \neq x$, then the derivation encountered a var, drop or drop-reuse case with $y = x$. The claim then follows by the inductive hypothesis.

case

$H_0 \vdash C \overline{x} - H_n$	given
$H_{i-1} \vdash x_i \dashv H_i$	DRCON
$y \notin H_n$	given
$\exists j. y \in H_j, y \notin H_{j+1}$	follows
$(H_i - \llbracket y \rrbracket) \# z \vdash x_i[y:=z] \dashv (H_{i+1} - \llbracket y \rrbracket) \# z$	for $i < j$, by Lemma 13
$(H_j - \llbracket y \rrbracket) \# z \vdash y[y:=z] \dashv H_{j+1}$	I.H.
$H_i \vdash x_i[y:=z] \dashv H_{i+1}$	for $i > j$, above
$(H_0 - \llbracket y \rrbracket) \# z \vdash (C \overline{x})[y:=z] \dashv H_n$	DRCON

case

$H, \llbracket y \rrbracket \vdash \text{dup } x; e \dashv H_1$	given
$(H \# x), \llbracket y \rrbracket \vdash e \dashv H_1$	by lemma 8
$(H \# x) \# z \vdash e[y:=z] \dashv H_1$	I.H.
$H \# z \vdash \text{dup } x[y:=z]; e[y:=z] \dashv H_1$	DRDUP
$H \# z \vdash (\text{dup } x; e)[y:=z] \dashv H_1$	follows

case

$H, \llbracket y \rrbracket \vdash \lambda_{\overline{y}} x. e \dashv H_1$	given
$H, \llbracket y \rrbracket \vdash \overline{y} \dashv H_1$	given
$\overline{y} \mapsto^1 (), x \mapsto^1 () \vdash e \dashv \emptyset$	given
$H \# z \vdash \overline{y}[y:=z] \dashv H_1$	I.H.
$H \# z \vdash (\lambda_{\overline{y}} x. e)[y:=z] \dashv H_1 \# z$	DRLAM

case

$H, \llbracket y \rrbracket \vdash e_1 x \dashv H_2$	given
$H, \llbracket y \rrbracket \vdash e_1 \dashv H_1$	given
$y \notin H_1$	assume
$H \# z \vdash e_1[y:=z] \dashv H_1$	I.H.
$H_1 \vdash x \dashv H_2$	given
$H_1 \vdash x[y:=z] \dashv H_2$	since $x \neq y$
$H \# z \vdash (e_1 x)[y:=z] \dashv H_2$	DRAPP
$y \in H_1$	else
$H \# z \vdash e_1[y:=z] \dashv (H_1 - \llbracket y \rrbracket) \# z$	by substitution and Lemma 13
$(H_1 - \llbracket y \rrbracket) \# z \vdash x[y:=z] \dashv H_2$	I.H.
$H \# z \vdash (e_1 x)[y:=z] \dashv H_2$	DRAPP
case	

$H, \llbracket y \rrbracket \vdash \text{val } x = e_1; e_2 \dashv H_2$	given
$H, \llbracket y \rrbracket \vdash e_1 \dashv H_1$	given
$y \notin H_1$	assume
$H \# z \vdash e_1[y:=z] \dashv H_1$	I.H.
$H_1, \llbracket x \rrbracket \vdash e_2 \dashv H_2$	given
$H_1, \llbracket x \rrbracket \vdash e_2[y:=z] \dashv H_2$	since $y \notin \text{fv}(e_2)$ by lemma 5
$H \# z \vdash (\text{val } x = e_1; e_2)[y:=z] \dashv H_2$	DRBIND
$y \in H_1$	else
$H \# z \vdash e_1[y:=z] \dashv (H_1 - \llbracket y \rrbracket) \# z$	by substitution and Lemma 13
$(H_1 - \llbracket y \rrbracket, \llbracket x \rrbracket) \# z \vdash e_2[y:=z] \dashv H_2$	I.H.
$H \# z \vdash (\text{val } x = e_1; e_2)[y:=z] \dashv H_2$	DRBIND
case	

$H \vdash \text{match } x \{ \overline{p_i \mapsto \text{dup}(\overline{z_i}); e_i} \} \dashv H'$	given
$x \in H, x[y:=z] \in (H - \llbracket y \rrbracket) \# z$	given
$H, H_i \vdash e_i \dashv H'$	given
$H_i = \llbracket \overline{z_i} \rrbracket$	given
$((H, H_i) - \llbracket y \rrbracket) \# z \vdash e_i[y:=z] \dashv H'$	I.H.
$\overline{z_i}$ fresh	assume
$H, \llbracket H_i \rrbracket = H \# \overline{z_i}$	
$((H, H_i) - \llbracket y \rrbracket) \# z = (H - \llbracket y \rrbracket) \# \overline{z_i} \# z$	
$= (H - \llbracket y \rrbracket) \# z \# \overline{z_i}$	Lemma 8
$= (H - \llbracket y \rrbracket) \# z, H_i$	
$(H - \llbracket y \rrbracket) \# z, H_i \vdash e_i \dashv H'$	by substitution
$(H - \llbracket y \rrbracket) \# z \vdash (\text{match } x \{ \overline{p_i \mapsto \text{dup}(\overline{z_i}); e_i} \})[y:=z] \dashv H'$	DRMATCH
□	

D.4 Soundness of Reference Counting Semantics

Definition 6. (Well-formed Abstractions)

If e ok, then all $(\lambda^{\overline{y}} x. e_1)$ in e satisfies $\llbracket \overline{y}, x \rrbracket \vdash e_1 \dashv \emptyset$.

Definition 7. (Well-formed Heap)

If H ok, then (1) if $x \mapsto^n v \in H$, then $\text{fv}(v) \in H$, and v ok; (2) there are no dependency cycles in H .

Lemma 15. (No Garbage (Small step))

Given H_1 ok and e_1 ok if $H_1 \vdash e_1 \dashv H'$, and $H_1 \mid e_1 \longrightarrow_r H_2 \mid e_2$, then H_2 ok, e_2 ok, and $H_2 \vdash e_2 \dashv H'$.

Proof. (Of Lemma 15) When we place a new variable $z \mapsto^1 v$ in the heap (e.g., (lam)), z is fresh so v cannot refer to z (even indirectly). So there is no dependency cycle. Also, in those cases, since $H_1 \vdash v \dashv H'$, by Lemma 5, we know $\text{fv}(v) \in H_1$. Moreover we have v ok as a precondition.

Heap reduction retains abstractions, with the only change being substitution. If $\llbracket \overline{y}, x \rrbracket \vdash e \dashv \emptyset$, then $\llbracket \overline{y}[y:=z], x \rrbracket \vdash e[y:=z] \dashv \emptyset$ by substitution.

Now we prove $H_2 \vdash e_2 \dashv H'$ by induction on the judgment.

case (app_r) $H \mid f z \longrightarrow_r H \mid \text{dup } \bar{y}; \text{drop } f; e[x:=z] \quad (f \mapsto^n \lambda_{\bar{y}}x. e) \in H$

$H \vdash f z \dashv H_1$	given
$H = H_1 \# f \# z$	Lemma 6
$\bar{y} \in H_1 \# f \# z$	$f \mapsto \lambda_{\bar{y}}x. e$
$H \vdash \text{dup } \bar{y}; () \dashv H \# \bar{y}$	by definition
$H \# \bar{y} = H_1 \# f \# z \# \bar{y}$	
$= H_1 \# \bar{y} \# z \# f$	Lemma 8
$H_1 \# \bar{y} \# z \# f \vdash \text{drop } f; () \dashv H_1 \# \bar{y} \# z$	Lemma 7
$\llbracket \bar{y}, x \rrbracket \vdash e \dashv \emptyset$	$\lambda_{\bar{y}}x. e$ ok
$\llbracket \bar{y}, z \rrbracket \vdash e[x:=z] \dashv \emptyset$	by substitution
$H_1 \# \bar{y} \# z \vdash e[x:=z] \dashv H_1$	Lemma 13 and Lemma 14

case ($match_r$) $H \mid \text{match } x \{ p_i \rightarrow \text{dup } (\bar{x}); e_i \} \longrightarrow_r H \mid \text{dup } (\bar{x}); e_i[\bar{x}:=\bar{y}]$ with $p_i = C \bar{x}$ and $(x \mapsto^n C \bar{y}) \in H$

$H \vdash \text{match } x \{ C \bar{x} \rightarrow \text{dup } \bar{x}; e_i \} \dashv H'$	given
$H, \llbracket \bar{x} \rrbracket \vdash e_i \dashv H'$	given
$\bar{x} \notin H_i$	given
$(x \mapsto^n C \bar{y}) \in H$	given
$\bar{y} \in H$	H ok
$H \vdash \text{dup } \bar{y}; () \dashv H \# \bar{y}$	by definition
$H \# \bar{y} \vdash e_i[\bar{x}:=\bar{y}] \dashv H'$	Lemma 14
$H \vdash \text{dup } (\bar{x}); e_i[\bar{x}:=\bar{y}] \dashv H'$	DUP

case (let_r) $H \mid \text{let } x = z \text{ in } e \longrightarrow_r H \mid e[x:=z]$

$H \vdash \text{let } x = z \text{ in } e \dashv H'$	given
$H \vdash z \dashv H_1$	given
$H_1 = H \# z$	Lemma 6
$H_1, x \mapsto^1 () \vdash e \dashv H'$	given
$H_1 \# z \vdash e[x:=z] \dashv H'$	Lemma 14
$H \vdash e[x:=z] \dashv H'$	

case (lam_r) $H \mid (\lambda_{\bar{y}} x. e) \longrightarrow_h H, f \mapsto^1 \lambda_{\bar{y}} x. e \mid f \quad \text{fresh } f$

$H \vdash \lambda_{\bar{y}} x. e \dashv H_1$	given
$H \vdash \bar{y} \dashv H_1$	given
$H, f \mapsto^1 \lambda_{\bar{y}} x. e \vdash f \dashv H_1$	DRVVAR

case (con_r) $H \mid C x_1 \dots x_n \longrightarrow_r H, z \mapsto^1 C x_1 \dots x_n \mid z$ fresh z

$H \vdash C x_1 \dots x_n \dashv H_1$	given
$H, z \mapsto^1 C x_1 \dots x_n \vdash z \dashv H$	DRCON

case (dup_r) $H, x \mapsto^n v \mid \text{dup } x; e \longrightarrow_r H, x \mapsto^{n+1} v \mid e$

$H, x \mapsto^n v \vdash \text{dup } x; e \dashv H_1$	given
$H, x \mapsto^{n+1} v \vdash e \dashv H_1$	DRDUP

case ($drop_r$) $H, x \mapsto^{n+1} v \mid \text{drop } x; e \longrightarrow_r H, x \mapsto^n v \mid e \quad \text{if } n \geq 1$

$H, x \mapsto^{n+1} v \vdash \text{drop } x; e \dashv H_1$	given
$H, x \mapsto^n v \vdash e \dashv H_1$	DRDROP

case ($diam_r$) $H, x \mapsto^1 \lambda_{\bar{y}}z.e \mid \text{drop } x; e \longrightarrow_r H \mid \text{drop } \bar{y}; e$

$H, x \mapsto^1 \lambda_{\bar{y}}z.e \vdash \text{drop } x; e \dashv H_1$	given
$H \vdash \text{drop } \bar{y}; e \dashv H_1$	DRDROPLAM

case ($dcon_r$) $H, x \mapsto^1 C \bar{y} \mid \text{drop } x; e \longrightarrow_r H \mid \text{drop } \bar{y}; e$

$H, x \mapsto^1 C \bar{y} \vdash \text{drop } x; e \dashv H_1$	given
$H \vdash \text{drop } \bar{y}; e \dashv H_1$	DRDROPCON

case ($drop_{ru}$) $H, x \mapsto^{n+1} v \mid r \leftarrow \text{drop}_r x; e \longrightarrow_r H, x \mapsto^n v, z \mapsto^1 () \mid e[r:=z] \quad \text{fresh } z$

$H, x \mapsto^{n+1} v \vdash r \leftarrow \text{drop}_r x; e \dashv H_1$	given
$H, x \mapsto^n v, r \mapsto^1 () \vdash e \dashv H_1$	DRDROP RU
$H, x \mapsto^n v, z \mapsto^1 () \vdash e[r:=z] \dashv H_1$	by substitution

$\text{case } (d\text{lam}_{ru}) \text{ H}, x \mapsto^1 \lambda_{\bar{x}} y. e' \mid r \leftarrow \text{dropru } x; e \longrightarrow_r \text{H}, z \mapsto^1 () \mid \text{drop } \bar{x}; e[r:=z] \quad \text{fresh } z$
 $\text{H}, x \mapsto^1 \lambda_{\bar{y}} z. e \vdash r \leftarrow \text{dropru } x; e \vdash \text{H}_1 \quad \text{given}$
 $\text{H}, r \mapsto^1 () \vdash \text{drop } \bar{y}; e \vdash \text{H}_1 \quad \text{DRDROPLAMRU}$
 $\text{H}, z \mapsto^1 () \vdash \text{drop } \bar{y}; e[r:=z] \vdash \text{H}_1 \quad \text{by substitution}$
 $\text{case } (d\text{con}_{ru}) \text{ H}, x \mapsto^1 C \bar{x} \mid r \leftarrow \text{dropru } x; e \longrightarrow_r \text{H}, \bar{z} \mapsto^1 (), z \mapsto^1 C \bar{z} \mid \text{drop } \bar{x}; e[r:=z] \quad \text{fresh } z, \bar{z}$
 $\text{H}, x \mapsto^1 C \bar{y} \vdash r \leftarrow \text{dropru } x; e \vdash \text{H}_1 \quad \text{given}$
 $\text{H}, r \mapsto^1 () \vdash \text{drop } \bar{y}; e \vdash \text{H}_1 \quad \text{DRDROPCONRU}$
 $\text{H}, \bar{z} \mapsto^1 (), z \mapsto^1 C \bar{z} \vdash \text{drop } \bar{y}; e[r:=z] \vdash \text{H}_1 \quad \text{by Lemma 14}$
 \square

The ok part reasoning of Lemma 15 can be easily generalized to big step. So from now on we will implicitly assume every expression and heap we discuss is ok.

We introduced the heap evaluation context in Figure 10 as:

$E ::= \square \mid E e \mid x E$
 $\mid \text{let } x = E \text{ in } e$
 $\mid \text{match } E \{ \bar{p}_i \mapsto \bar{e}_i \}$
 $\mid C x_1 \dots x_{i-1} E v_{i+1} \dots v_n$

But that was only necessary for terms, that hadn't been normalized (as we encountered them in lemma 4). But in the next proofs we will only apply the heap semantics to normalized terms: An expression e' such that $\Delta \mid \Gamma \vdash e \rightsquigarrow e'$ needs to be normalized and the heap semantics transform normalized terms to normalized terms since only variables and not values are substituted in. For the remainder of this section we will therefore work with the evaluation context:

$E ::= \square \mid E x \mid \text{val } x = E; e$

Lemma 16. (No Garbage (big step))

If $\text{H}_1 \vdash E[e_1] \vdash \text{H}'$, and $\text{H}_1 \mid E[e_1] \mapsto_r \text{H}_2 \mid E[e_2]$, then $\text{H}_2 \vdash E[e_2] \vdash \text{H}'$.

Proof. (Proof for Lemma 16) By induction on E.

case $E = \square$. Follows by Lemma 15.

case $E = E_1 x$.

$\text{H}_1 \vdash E_1[e_1] x \vdash \text{H}_2 \quad \text{given}$
 $\text{H}_1 \vdash E_1[e_1] \vdash \text{H}_3 \quad \text{DRAPP}$
 $\text{H}_3 \vdash x \vdash \text{H}_2 \quad \text{given}$
 $\text{H}_1 \vdash E_1[e_2] \vdash \text{H}_3 \quad \text{I.H.}$
 $\text{H}_1 \vdash E_1[e_2] x \vdash \text{H}_2 \quad \text{DRAPP}$

case $E = \text{val } x = E_1; e$.

$\text{H}_1 \vdash \text{val } x = E_1[e_1]; e \vdash \text{H}_2 \quad \text{given}$
 $\text{H}_1 \vdash E_1[e_1] \vdash \text{H}_3 \quad \text{DRBIND}$
 $\text{H}_3, x \mapsto^1 () \vdash e \vdash \text{H}_2 \quad \text{given}$
 $x \notin \text{H}_2 \quad \text{given}$
 $\text{H}_1 \vdash E_1[e_2] \vdash \text{H}_3 \quad \text{I.H.}$
 $\text{H}_1 \vdash \text{val } x = E_1[e_2]; e \vdash \text{H}_2 \quad \text{DRBIND}$

\square

The next lemma can't be stated with $[\text{H}_1]e_1 = [\text{H}'_1]e_1$, because reuse tokens are handled differently in the heap calculi and thus we might have $\text{H}_1 = r \mapsto^1 ()$, $\text{H}'_1 = r \mapsto^1 C \bar{v}$, $e_1 = r$. But we know that reuse tokens can only appear as the argument to dup/drop/dropru (as they are chosen fresh in the DROPRU rule), and so we can remove them beforehand ($\lceil e_1 \rceil$). Because of the (lam_h) rule, we denote by $\lceil e_1 \rceil$ here a procedure that deletes all dups, drops and droprus in e_1 but doesn't descend into lambdas (and doesn't delete dups, drops, droprus there).

Lemma 17. (Reference counting semantics is sound (small step))

If $\text{H}_1 \mid e_1 \longrightarrow_h \text{H}_2 \mid e_2$, and $\text{H}'_1 \vdash e_1 \vdash \text{H}_3$, and $[\text{H}_1]\lceil e_1 \rceil = [\text{H}'_1]\lceil e_1 \rceil$, then there exists H'_2 such that $\text{H}'_1 \mid e_1 \longrightarrow^*_r \text{H}'_2 \mid e_2$ and $[\text{H}_2]\lceil e_2 \rceil = [\text{H}'_2]\lceil e_2 \rceil$.

Proof. (Of Lemma 17) By case analysis on the heap judgment. The rules (lam_h) , (con_h) , (app_h) , (match_h) , (let_h) and (lam_r) , (con_r) , (app_r) , (match_r) , (let_r) are identical. Note in particular, that in the (lam_h) and (lam_r) rules we store the same lambda, since $\lceil e_1 \rceil$ doesn't descend into lambdas. We thus only have to show the claim for the (dup_h) , (drop_h) and (dropru_h) rules.

case (dup_h) $H \mid dup\ x; e \longrightarrow_h H \mid e$

$H'_1 \vdash dup\ x; e \vdash H_3$ given
 $x \in H'_1$ follows
 $H'_1 \mid dup\ x; e \longrightarrow_r H'_1 \# x \mid e$ (dup) and $\#$
 $H'_2 = H'_1 \# x$ follows
 $[H_1][e] = [H'_1][e]$ given
 $[H_2][e] = [H_1][e] = [H'_1][e] = [H'_2][e]$ follows

case ($drop_h$) $H \mid drop\ x; e \longrightarrow_h H \mid e$

$H'_1 \vdash drop\ x; e \vdash H_3$ given
 $H'_1 \vdash drop\ x; () \vdash H'_2$ follows
 $H'_2 \vdash e \vdash H_3$ above
 $fv(e) \in H'_2$ Lemma 5
 $H'_1 = H'_2 \# x$ Lemma 6
 $H'_1 \mid drop\ x; e \longrightarrow_r H'_2 \mid e$ ($drop_r$), ($d\text{lam}_r$), ($d\text{con}_r$)
 $[H_2][e] = [H_1][e] = [H'_1][e] = [H'_2][e]$ follows

case ($dropru_h$) $H \mid r \leftarrow dropru\ x; e \longrightarrow_h H, z \mapsto^1 () \mid e[r:=z]$

$H'_1 \vdash r \leftarrow dropru\ x; e \vdash H_3$ given
 $H'_1 \vdash r \leftarrow dropru\ x; () \vdash H'_2$ follows
 $H'_2 \vdash e \vdash H_3$ above
 $fv(e) \in H'_2$ Lemma 5
 $H'_1 \# r = H'_2 \# x$ Lemma 9
 $H'_2 = H'_2[r:=z]$ let
 $H'_1 \mid r \leftarrow dropru\ x; e \longrightarrow_r H'_2 \mid e[r:=z]$ ($drop_{ru}$), ($d\text{lam}_{ru}$), ($d\text{con}_{ru}$)
 $[H_2][e] = [H_1, z \mapsto^1 ()][e] = [H'_1, z \mapsto^1 ()][e]$ since z is fresh
 $[H'_1, z \mapsto^1 ()][e] = [H'_2][e]$ by LAM and DROPRU

□

Lemma 18. (*Reference counting semantics is sound (big step)*)

If $H_1 \mid E[e_1] \longrightarrow_h H_2 \mid E[e_2]$, and $H'_1 \vdash E[e_1] \vdash H_3$, and $[H_1][E[e_1]] = [H'_1][E[e_1]]$, then there exists H'_2 such that $H'_1 \mid E[e_1] \longrightarrow^*_r H'_2$ and $[H_2][E[e_2]] = [H'_2][E[e_2]]$.

Proof. (*Of Lemma 18*) By induction on E .

case $E = \square$. Follows by Lemma 17.

case $E = E_1\ x$.

$H_1 \mid E[e_1] \longrightarrow_h H_2 \mid E[e_2]$ given
 $H_1 \mid E_1[e_1]\ x \longrightarrow_h H_2 \mid E_1[e_2]\ x$ given
 $H_1 \mid E_1[e_1] \longrightarrow_h H_2 \mid E_1[e_2]$ follows
 $H'_1 \vdash E[e_1] \vdash H_3$ given
 $H'_1 \vdash e_1 \vdash H_4$ follows
 $[H_1][E[e_1]] = [H'_1][E[e_1]]$ given
 $[H_1][E_1[e_1]] = [H'_1][E_1[e_1]]$ given
 $H'_1 \mid E_1[e_1] \longrightarrow^*_r H'_2 \mid E_1[e_2]$ I.H.
 $H'_1 \mid E_1[e_1]\ x \longrightarrow^*_r H'_2 \mid E_1[e_2]\ x$ follows
 $H'_1 \mid E[e_1] \longrightarrow^*_r H'_2 \mid E[e_2]$ follows
 $[H_2][E_1[e_1]] = [H'_2][E_1[e_1]]$ above
 $[H_2][E[e_2]] = [H'_2][E[e_2]]$ follows

case $E = \text{val}\ x = E_1; e$.

$H_1 \mid E[e_1] \longrightarrow_h H_2 \mid E[e_2]$	given
$H_1 \mid \text{let } x = E_1[e_1] \text{ in } e \longrightarrow_h H_2 \mid \text{let } x = E_1[e_2] \text{ in } e$	given
$H_1 \mid E_1[e_1] \longrightarrow_h H_2 \mid E_1[e_2]$	follows
$H'_1 \vdash E[e_1] \dashv H_3$	given
$H'_1 \vdash e_1 \dashv H_4$	follows
$[H_1][E[e_1]] = [H'_1][E[e_1]]$	given
$[H_1][E_1[e_1]] = [H'_1][E_1[e_1]]$	given
$H'_1 \mid E_1[e_1] \longrightarrow^*_r H'_2 \mid E_1[e_2]$	I.H.
$H'_1 \mid \text{let } x = E_1[e_1] \text{ in } e \longrightarrow^*_r H'_2 \mid \text{let } x = E_1[e_2] \text{ in } e$	follows
$H'_1 \mid E[e_1] \longrightarrow^*_r H'_2 \mid E[e_2]$	follows
$[H_2][E_1[e_1]] = [H'_2][E_1[e_1]]$	above
$[H_2][E[e_2]] = [H'_2][E[e_2]]$	follows
\square	

Lemma 19. (Reference counting semantics is sound (big step, star))

If $H_1 \mid e_1 \longmapsto^*_h H_2 \mid e_2$, and $H'_1 \vdash e_1 \dashv H_3$, and $[H_1][e_1] = [H'_1][e_1]$, then there exists H'_2 such that $H'_1 \mid e_1 \longmapsto^*_r H'_2 \mid e_2$ and $[H_2][e_2] = [H'_2][e_2]$.

Proof. (Of Lemma 19) We can apply lemma 18 repeatedly: After each application, we know by lemma 16 that $H'_2 \vdash e_2 \dashv H_3$.

Proof. (Of Theorem 1)

$\emptyset \mid \emptyset \vdash e \rightsquigarrow e'$	given
$e \longmapsto^*_r v$	given
$e' \longmapsto^*_r v$	Theorem 6
$\emptyset \mid e' \longmapsto^*_h H_2 \mid x$	Theorem 4
$[H_2]x = v$	above
\emptyset ok	
e' ok	from LAM
$\emptyset \vdash e' \dashv \emptyset$	Lemma 12
$\emptyset \mid e' \longmapsto^*_r H_3 \mid x$	Lemma 19
$[H_3]x = [H_2]x = v$	above
\square	

D.5 No Garbage

Theorem 7. (No garbage)

Given $\emptyset \mid \emptyset \vdash e \rightsquigarrow e_1$, and $\emptyset \mid e_1 \longmapsto^*_r H_i \mid e_i$, then $H_i \vdash e_i \dashv \emptyset$.

Proof. (Of Theorem 7)

\emptyset ok	by construction
e_1 ok	by LAM
$\emptyset \vdash e_1 \dashv \emptyset$	Lemma 12
$H_i \vdash e_i \dashv \emptyset$	Lemma 16
\square	

Lemma 20. (Reachability)

If $H_1 \vdash e \dashv H_2$, then there for all $y \in \text{dom}(H_1) - \text{dom}(H_2)$, $\text{reach}(y, H_1 \mid e)$.

For ease of reference, we denote it as $\text{reach}(H_1 - H_2, H_1 \mid e)$

Proof. (Of Lemma 20) By induction on the judgment.

case

$H_0 \vdash C x_1 \dots x_n \dashv H_n$ given
 $H_0 \vdash x_1 \dashv H_1 \dots H_{n-1} \vdash x_n \dashv H_n$ DRCON
 $\text{reach}(H_{i-1} - H_i, H_{i-1} \mid x_i)$ I.H.
 $\text{reach}(H_{i-1} - H_i, H_0 \mid x_i)$ Lemma 5
 $\text{reach}(H_n - H_0, H_0 \mid C x_1 \dots x_n)$ Follows
case

$H, x \mapsto^{n+1} v \vdash x \dashv H, x \mapsto^n v$ given
 $\text{dom}(H, x \mapsto^{n+1} v) - \text{dom}(H, x \mapsto^n v) = \emptyset$
case

$H, x \mapsto^1 \lambda_{\bar{y}}z. e \vdash x \dashv H_1$ given
 $H \vdash \bar{y} \dashv H_1$ DRVARLAM
 $\text{reach}(H - H_1, H \mid \bar{y})$ I.H.
 $\text{reach}(H - H_1, H \mid \lambda_{\bar{y}}z. e)$ by definition
 $\text{reach}((H, x \mapsto^1 \lambda_{\bar{y}}z. e) - H_1, (H, x \mapsto^1 \lambda_{\bar{y}}z. e) \mid x)$ follows
case

$H, x \mapsto^1 C \bar{y} \vdash x \dashv H_1$ given
 $H \vdash \bar{y} \dashv H_1$ DRVARCON
 $\text{reach}(H - H_1, H \mid \bar{y})$ I.H.
 $\text{reach}(H - H_1, H \mid C \bar{y})$ by definition
 $\text{reach}((H, x \mapsto^1 C \bar{y}) - H_1, (H, x \mapsto^1 C \bar{y}) \mid x)$ follows
case

$H, x \mapsto^n v \vdash \text{dup } x; e \dashv H_1$ given
 $H, x \mapsto^{n+1} v \vdash e \dashv H_1$ DRDUP
 $\text{reach}((H, x \mapsto^{n+1} v) - H_1, (H, x \mapsto^{n+1} v) \mid e)$ I.H.
 $\text{reach}((H, x \mapsto^n v) - H_1, (H, x \mapsto^n v) \mid \text{dup } x; e)$ follows
case

$H, x \mapsto^{n+1} v \vdash \text{drop } x; e \dashv H_1$ given
 $H, x \mapsto^n v \vdash e \dashv H_1$ DRDROP
 $\text{reach}((H, x \mapsto^n v) - H_1, (H, x \mapsto^n v) \mid e)$ I.H.
 $\text{reach}((H, x \mapsto^{n+1} v) - H_1, (H, x \mapsto^{n+1} v) \mid \text{drop } x; e)$ follows
case

$H, x \mapsto^1 C \bar{y} \vdash \text{drop } x; e \dashv H_1$ given
 $H \vdash \text{drop } \bar{y}; e \dashv H_1$ DRDROPCON
 $\text{reach}(H - H_1, H \mid \text{drop } \bar{y}; e)$ I.H.
 $\text{reach}((H, x \mapsto^1 C \bar{y}) - H_1, (H, x \mapsto^1 C \bar{y}) \mid \text{drop } x; e)$ follows
case

$H, x \mapsto^1 \lambda_{\bar{y}}z. e \vdash \text{drop } x; e \dashv H_1$ given
 $H \vdash \text{drop } \bar{y}; e \dashv H_1$ DRDROPLAM
 $\text{reach}(H - H_1, H \mid \text{drop } \bar{y}; e)$ I.H.
 $\text{reach}((H, x \mapsto^1 \lambda_{\bar{y}}z. e) - H_1, (H, x \mapsto^1 \lambda_{\bar{y}}z. e) \mid \text{drop } x; e)$ follows
case

$H, x \mapsto^{n+1} v \vdash r \leftarrow \text{dropru } x; e \dashv H_1$ given
 $H, x \mapsto^n v, r \mapsto^1 () \vdash e \dashv H_1$ DRDROPRU
 $\text{reach}((H, x \mapsto^n v, r \mapsto^1 ()) - H_1, (H, x \mapsto^n v, r \mapsto^1 ())) \mid e$ I.H.
 $\text{reach}((H, x \mapsto^n v) - H_1, (H, x \mapsto^n v, r \mapsto^1 ())) \mid e$ follows
 $\text{reach}((H, x \mapsto^{n+1} v) - H_1, (H, x \mapsto^{n+1} v) \mid r \leftarrow \text{dropru } x; e)$ follows
case

$H, x \mapsto^1 C \bar{y} \vdash r \leftarrow \text{dropru } x; e \vdash H_1$	given
$H, r \mapsto^1 () \vdash \text{drop } \bar{y}; e \vdash H_1$	DRDROPCONRU
$\text{reach}((H, r \mapsto^1 ()) - H_1, H \mid \text{drop } \bar{y}; e)$	I.H.
$\text{reach}(H - H_1, H \mid \text{drop } \bar{y}; e)$	follows
$\text{reach}((H, x \mapsto^1 C \bar{y}) - H_1, (H, x \mapsto^1 C \bar{y}) \mid r \leftarrow \text{dropru } x; e)$	follows
case	

$H, x \mapsto^1 \lambda_{\bar{y}} z. e \vdash r \leftarrow \text{dropru } x; e \vdash H_1$	given
$H, r \mapsto^1 () \vdash \text{drop } \bar{y}; e \vdash H_1$	DRDROPLAMRU
$\text{reach}((H, r \mapsto^1 ()) - H_1, H \mid \text{drop } \bar{y}; e)$	I.H.
$\text{reach}(H - H_1, H \mid \text{drop } \bar{y}; e)$	follows
$\text{reach}((H, x \mapsto^1 \lambda_{\bar{y}} z. e) - H_1, (H, x \mapsto^1 \lambda_{\bar{y}} z. e) \mid r \leftarrow \text{dropru } x; e)$	follows
case	

$H \vdash \lambda_{\bar{y}} x. e \vdash H_1$	given
$H \vdash \bar{y} \vdash H_1$	DRLAM
$\text{reach}(H - H_1, H \mid \bar{y})$	I.H.
$\text{reach}(H - H_1, H \mid \lambda_{\bar{y}} x. e)$	follows
case	

$H \vdash e_1 e_2 \vdash H_2$	given
$H \vdash e_1 \vdash H_1$	DRAPP
$H_1 \vdash e_2 \vdash H_2$	DRAPP
$\text{reach}(H - H_1, H \mid e_1)$	I.H.
$\text{reach}(H_1 - H_2, H_1 \mid e_2)$	I.H.
$\text{reach}(H_1 - H_2, H \mid e_2)$	Lemma 5
$\text{reach}(H - H_2, H \mid e_1 e_2)$	follows
case	

$H \vdash \text{val } x = e_1; e_2 \vdash H_2$	given
$H \vdash e_1 \vdash H_1$	DRBIND
$H_1, x \mapsto^1 () \vdash e_2 \vdash H_2$	DRBIND
$x \notin H, H_2$	DRBIND
$\text{reach}(H - H_1, H \mid e_1)$	I.H.
$\text{reach}((H_1, x \mapsto^1 ()) - H_2, (H_1, x \mapsto^1 ()) \mid e_2)$	I.H.
$\text{reach}((H_1, x \mapsto^1 ()) - H_2, (H, x \mapsto^1 ()) \mid e_2)$	Lemma 5
$\text{dom}(H_1) \subseteq \text{dom}(H_1, x \mapsto^1 ())$	
$\text{reach}(H_1 - H_2, (H, x \mapsto^1 ()) \mid e_2)$	follows
$x \notin H$	known
$x \notin \text{dom}(H) - \text{dom}(H_2)$	follows
$\text{reach}(H - H_2, H \mid \text{val } x = e_1; e_2)$	follows
case	

$H \vdash \text{match } x \{ \overline{p_i \mapsto \text{dup } \bar{z}_i}; e_i \} \vdash H'$	given
$H, [\bar{z}_i] \vdash e_i \vdash H'$	DRMATCH
$\bar{z}_i \notin H, H'$	DRMATCH
$\text{reach}((H, [\bar{z}_i]) - H', (H, [\bar{z}_i]) \mid e_i)$	I.H.
$\text{dom}(H) \subseteq \text{dom}(H, [\bar{z}_i])$	
$\text{reach}(H - H', (H, [\bar{z}_i]) \mid e_i)$	follows
$\bar{z}_i \notin H$	known
$\bar{z}_i \notin \text{dom}(H) - \text{dom}(H')$	follows
$\text{reach}(H - H', H \mid \text{match } x \{ \overline{p_i \mapsto \text{dup } \bar{z}_i}; e_i \})$	follows
□	

Proof. (Of Theorem 2)

$\emptyset \mid \emptyset \vdash e \rightsquigarrow e'$ given
 $H_i \vdash e_i \dashv \emptyset$ Theorem 7
 $\text{reach}(H_i - \emptyset, H_i \mid e_i)$ Lemma 20
 \square

D.6 Precision

In this section we will define $(\star) = \text{reach}(H_1 - H_2, (H_1, x \mapsto^1 ()) \mid \lceil e_2 \rceil)$.

Lemma 21. (linear resource calculus relates to reference counting (Garbage free version))

If $\Delta \mid \Gamma \vdash_{\text{RCF}} e \rightsquigarrow e'$, then $\llbracket \Delta, \Gamma \rrbracket \vdash e' \dashv \llbracket \Delta \rrbracket$ (with (\star) above).

Proof. (Of Lemma 21) Except for the LET case all other parts of the proof of lemma 12 remain unchanged.

$\forall y \in \Gamma_2. y \in \text{fv}(e_2)$ (1), given
 $\text{reach}(\text{fv}(e_2), \llbracket \Gamma_2, x \rrbracket \mid e_2)$ (2), by the definition of reach
 $\text{reach}(\Gamma_2, \llbracket \Gamma_2, x \rrbracket \mid e_2)$ (3), by (1) and (2)
 $\text{reach}(\llbracket \Gamma_2, \Delta \rrbracket - \llbracket \Delta \rrbracket, \llbracket \Gamma_2, x \rrbracket \mid e_2)$ (4), by (3)
 $\text{reach}(\llbracket \Gamma_2, \Delta \rrbracket - \llbracket \Delta \rrbracket, \llbracket \Gamma_2, x \rrbracket \mid \lceil e'_2 \rceil)$ (5), by lemma 1

Lemma 22. (Garbage-free at variables)

If $H_1 \vdash x \dashv H_2$, then $\text{reach}(H_1 - H_2, H_1 \mid \lceil x \rceil)$.

Proof. (Of Lemma 22) Follows from lemma 20, since $\lceil x \rceil = x$.

Lemma 23. (Garbage-free)

If $H_1 \vdash E[v] \dashv H_2$, then $\text{reach}(H_1 - H_2, H_1 \mid \lceil E[v] \rceil)$.

Proof. (Of Lemma 23) By induction on the evaluation context.

case E = \square .

subcase $v = C \bar{x}$.

$H_0 \vdash C x_1 \dots x_n \dashv H_n$ given
 $H_0 \vdash x_1 \dashv H_1 \dots H_{n-1} \vdash x_n \dashv H_n$ DRCON
 $\text{reach}(H_{i-1} - H_i, H_{i-1} \mid x_i)$ Lemma 22
 $\text{reach}(H_{i-1} - H_i, H_0 \mid x_i)$ Lemma 5
 $\text{reach}(H_n - H_0, H_0 \mid C x_1 \dots x_n)$ Follows
 $\text{reach}(H_n - H_0, H_0 \mid \lceil C x_1 \dots x_n \rceil)$ Follows

subcase $v = x$

$H_1 \vdash x \dashv H_2$ given
 $\text{reach}(H_1 - H_2, H_1 \mid \lceil x \rceil)$ Lemma 22

subcase $v = \lambda_{\bar{y}} x. e$

$H \vdash \lambda_{\bar{y}} x. e \dashv H_1$ given
 $H \vdash \bar{y} \dashv H_1$ DRLAM
 $\text{reach}(H - H_1, H \mid \bar{y})$ Lemma 22
 $\text{reach}(H - H_1, H \mid \lambda_{\bar{y}} x. e)$ follows
 $\text{reach}(H - H_1, H \mid \lambda_{\bar{y}} x. \lceil e \rceil)$ follows
 $\text{reach}(H - H_1, H \mid \lceil \lambda_{\bar{y}} x. e \rceil)$ follows

case E = $E_1 x$.

$H \vdash E_1[v] x \dashv H_2$ given
 $H \vdash E_1[v] \dashv H_1$ DRAPP
 $H_1 \vdash x \dashv H_2$ DRAPP
 $\text{reach}(H - H_1, H \mid \lceil E_1[v] \rceil)$ I.H.
 $\text{reach}(H_1 - H_2, H_1 \mid \lceil x \rceil)$ Lemma 22
 $\text{reach}(H_1 - H_2, H \mid \lceil x \rceil)$ Lemma 5
 $\text{reach}(H - H_2, H \mid \lceil E_1[v] x \rceil)$ follows

case E = $\text{val } x = E_1[v]; e_2$.

$H \vdash \text{val } x = E_1[v]; e_2 \dashv H_2$	given
$H \vdash E_1[v] \dashv H_1$	DRBIND
$H_1, x \mapsto^1 () \vdash e_2 \dashv H_2$	DRBIND
$x \notin H, H_2$	DRBIND
$\text{reach}(H - H_1, H \mid \lceil E_1[v] \rceil)$	I.H.
$\text{reach}(H_1 - H_2, (H_1, x \mapsto^1 ()) \mid \lceil e_2 \rceil)$	(*)
$\text{reach}(H_1 - H_2, (H, x \mapsto^1 ()) \mid \lceil e_2 \rceil)$	Lemma 5
$x \notin H$	known
$x \notin \text{dom}(H) - \text{dom}(H_2)$	follows
$\text{reach}(H - H_2, H \mid \lceil \text{val } x = e_1; e_2 \rceil)$	follows
\square	

Proof. (Of Theorem 3)

$\emptyset \mid \emptyset \vdash e \rightsquigarrow e'$	given
$H_i \vdash e_i \dashv \emptyset$	Theorem 7
$\text{reach}(H_i - \emptyset, H_i \mid \lceil E[v] \rceil)$	Lemma 23
\square	

D.7 Frame-limited

In this section we will define $(\star) = \exists H_1^1, H_1^2 = H_1. \text{reach}(H_1^1 - H_2, (H_1, x \mapsto^1 ()) \mid \lceil e_2 \rceil) \wedge |H_1^2| \leq c$.

Lemma 24. (linear resource calculus relates to reference counting (Frame-limited version))

If $\Delta \mid \Gamma \vdash_{\text{FL}} e \rightsquigarrow e'$, then $\llbracket \Delta, \Gamma \rrbracket \vdash e' \dashv \llbracket \Delta \rrbracket$ (with (\star) above).

Proof. (Of Lemma 21) Except for the LET case all other parts of the proof of lemma 12 remain unchanged.

$\Gamma_2 = \Gamma', \Gamma''$ where $\Gamma' \subseteq \text{fv}(e_2)$ and $\text{sizeof}(\Gamma'') \leq c$	(1), given
$H^1 = \llbracket \Gamma', \Delta \rrbracket$	(2), let
$\text{reach}(\text{fv}(e_2), \llbracket \Gamma_2, x \rrbracket \mid e_2)$	(3), by the definition of reach
$\text{reach}(H^1 - \llbracket \Delta \rrbracket, \llbracket \Gamma_2, x \rrbracket \mid e_2)$	(4), by (2) and (3)
$\text{reach}(H^1 - \llbracket \Delta \rrbracket, \llbracket \Gamma_2, x \rrbracket \mid \lceil e'_2 \rceil)$	(5), by lemma 1
$H^2 = \llbracket \Gamma_2, \Delta \rrbracket - H^1 = \Gamma''$	(6), let
$ H^2 \leq c$	(7), by (1)

With this (\star) rule it is not clear that evaluation preserves it: After all, we are substituting new values into an expression and the values may be big. But if the $\text{sizeof}(y)$ predicate is chosen such that no bigger value can be bound to y (for example, with help of a type system), then substitution preserves (\star) .

Lemma 25. (Substitution preserves dependencies (Frame-limited version))

If $y, z \notin H_1, H_2$ and $H_1, \llbracket y \rrbracket \vdash e \dashv H_2$ and $\text{sizeof}(y) = \text{sizeof}(z)$, then $H_1 \dashv z \vdash e[y:=z] \dashv H_2$.

Proof. (Of Lemma 25) Except for the LET case all other parts of the proof of lemma 14 remain unchanged.

$H, \llbracket y \rrbracket \vdash \text{val } x = e_1; e_2 \dashv H_2$	given
$\exists H_1^1, H_1^2 = H_1. \text{reach}(H_1^1 - H_2, (H_1, x \mapsto^1 ()) \mid \lceil e_2 \rceil) \wedge H_1^2 \leq c$	given
$y \in H_1$	else
$H \dashv z \vdash e_1[y:=z] \dashv (H_1 - \llbracket y \rrbracket) \dashv z$	by substitution and Lemma 13
$(H_1 - \llbracket y \rrbracket, \llbracket x \rrbracket) \dashv z \vdash e_2[y:=z] \dashv H_2$	I.H.
$\text{reach}((H_1^1 - \llbracket y \rrbracket) \dashv z - H_2, (H_1 - \llbracket y \rrbracket) \dashv z, x \mapsto^1 ()) \mid \lceil e_2[y:=z] \rceil$	follows (if $y \in H_1^1$)
$ H_1^2 - \llbracket y \rrbracket \dashv z \leq c$	follows (if $y \in H_1^2$)
$H \dashv z \vdash (\text{val } x = e_1; e_2)[y:=z] \dashv H_2$	DRBIND

Lemma 26. (Frame-limited)

If $H_1 \vdash E[v] \dashv H_2$, there are $H_1^1, H_1^2 = H_1$ such that $\text{reach}(H_1^1 - H_2, H_1 \mid \lceil E[v] \rceil)$ and $|H_1^2| \leq c \cdot |E|$

Proof. (Of Lemma 23) By induction on the evaluation context.

case $E = \square$.

By lemma 23, we can set $H_1^1 = H_1, H_1^2 = \emptyset$. **case** $E = E_1 x$.

$H \vdash E_1[v] x \dashv H_2$	given
$H \vdash E_1[v] \dashv H_1$	DRAPP
$H_1 \vdash x \dashv H_2$	DRAPP
$\text{reach}(H^1 - H_1, H \mid \lceil E_1[v] \rceil)$	I.H.
$ H^2 \leq c \cdot E_1 $	I.H.
$\text{reach}(H_1 - H_2, H_1 \mid \lceil x \rceil)$	Lemma 22
$\text{reach}(H_1 - H_2, H \mid \lceil x \rceil)$	Lemma 5
$H^1 = H'^1, (H_1 - H_2)$	let
$\text{reach}(H^1 - H_2, H \mid \lceil E_1[v] x \rceil)$	follows
$ H^2 \leq c \cdot E $	follows
case $E = \text{val } x = E_1[v]; e_2.$	
$H \vdash \text{val } x = E_1[v]; e_2 \dashv H_2$	given
$H \vdash E_1[v] \dashv H_1$	DRBIND
$H_1, x \mapsto^1 () \vdash e_2 \dashv H_2$	DRBIND
$x \notin H, H_2$	DRBIND
$\text{reach}(H^1 - H_1, H \mid \lceil E_1[v] \rceil)$	I.H.
$ H^2 \leq c \cdot E_1 $	I.H.
$\text{reach}(H_1^1 - H_2, (H_1, x \mapsto^1 ()) \mid \lceil e_2 \rceil)$	(*)
$ H_1^2 \leq c$	(*)
$\text{reach}(H_1^1 - H_2, (H, x \mapsto^1 ()) \mid \lceil e_2 \rceil)$	Lemma 5
$H_3 = (H^1 - H_1), H_1^1$	let
$H_4 = H^2, H_1^2$	let
$x \notin H$	known
$x \notin \text{dom}(H) - \text{dom}(H_2)$	follows
$\text{reach}(H_3 - H_2, H \mid \lceil \text{val } x = e_1; e_2 \rceil)$	follows
$ H_4 \leq c \cdot (E_1 + 1) = c \cdot E $	follows
□	

Proof. (Of Theorem 4)

$\emptyset \mid \emptyset \vdash e \rightsquigarrow e'$	given
$H_i \vdash e_i \dashv \emptyset$	Theorem 7
$\text{reach}(H_i^1 - \emptyset, H_i \mid \lceil E[v] \rceil)$	Lemma 26
$ H_i^2 \leq c \cdot E $	above
$H_i = H_i^1, H_i^2$	above
□	

D.8 Reuse analysis is frame-limited

Proof. (Of Theorem 5) We can prove this by induction on the reuse judgment $S \mid R \Vdash e' \rightsquigarrow e''$.

case RVAR.

$S \mid R \Vdash x \rightsquigarrow x$	(1), given
$\Delta \mid x \vdash_{\text{GF}} x \rightsquigarrow x$	(2), given
$\Delta \mid x, \emptyset \vdash_{\text{FL}} x \rightsquigarrow x$	(3), by <small>VAR</small>

case RCON.

$S \mid R \Vdash C \bar{x} \rightsquigarrow C \bar{x}$	(1), given
$\Delta \mid \bar{x} \vdash_{\text{GF}} C \bar{x} \rightsquigarrow C \bar{x}$	(2), given
$\Delta \mid \bar{x}, \emptyset \vdash_{\text{FL}} C \bar{x} \rightsquigarrow C \bar{x}$	(3), by <small>CON</small> .

case RLAM.

$e' = \lambda_{\Gamma} x. \underline{e'}$ (1), given
 $e'' = \lambda_{\Gamma} x. \underline{e''}$ (2), given
 $\Delta \mid \Gamma \vdash_{\text{GF}} \lambda x. \underline{e} \rightsquigarrow \lambda_{\Gamma} x. \underline{e'}$ (3), given and by (1)
 $\emptyset \mid \Gamma, x \vdash_{\text{GF}} \underline{e} \rightsquigarrow \underline{e'}$ (4), by (3)
 $S \mid \emptyset \Vdash \underline{e'} \rightsquigarrow \underline{e''}$ (5), given
 $\emptyset \mid \Gamma, x, \emptyset \vdash_{\text{FL}} \underline{e} \rightsquigarrow \underline{e''}$ (6), by induction
 $\Delta \mid \Gamma, \emptyset \vdash_{\text{FL}} \lambda x. \underline{e} \rightsquigarrow \lambda_{\Gamma} x. \underline{e''}$ (7), by LAM

case RAPP.

$e' = \underline{e'} x$ (1), given
 $e'' = \underline{e''} x$ (2), given
 $\Delta \mid \Gamma \vdash_{\text{GF}} \underline{e} x \rightsquigarrow \underline{e'} x$ (3), given and by (1)
 $\Delta, x \mid \Gamma \vdash_{\text{GF}} \underline{e} \rightsquigarrow \underline{e'}$ (4), by (3)
 $S \mid R \Vdash \underline{e'} \rightsquigarrow \underline{e''}$ (5), given
 $\Delta, x \mid \Gamma, R \vdash_{\text{FL}} \underline{e} \rightsquigarrow \underline{e''}$ (6), by induction
 $\Delta \mid \Gamma, R \vdash_{\text{FL}} \underline{e} x \rightsquigarrow \underline{e''} x$ (7), by APP

case RDROP-REUSE.

$\Delta \mid \Gamma \vdash_{\text{GF}} e \rightsquigarrow e'$ (1), given
 $S \mid R \Vdash e' \rightsquigarrow e''$ (2), given
 $\Delta \mid \Gamma, R \vdash_{\text{FL}} e \rightsquigarrow e''$ (3), by induction
 $\Delta \mid \Gamma, R, r \vdash_{\text{FL}} e \rightsquigarrow \text{drop } r; e''$ (4), by DROP

case RLET.

$S \mid R_1, R_2 \Vdash \text{let } x = e'_1 \text{ in } e'_2 \rightsquigarrow \text{let } x = e''_1 \text{ in } e''_2$ (1), given
 $S \mid R_1 \Vdash e'_1 \rightsquigarrow e''_1$ (2), by (1)
 $S \mid R_2 \Vdash e'_2 \rightsquigarrow e''_2$ (3), by (1)
 $\Delta \mid \Gamma_1, \Gamma_2 \vdash_{\text{GF}} \text{let } x = e_1 \text{ in } e_2 \rightsquigarrow \text{let } x = e'_1 \text{ in } e'_2$ (4), given
 $\Delta, \Gamma_2 \mid \Gamma_1 \vdash_{\text{GF}} e_1 \rightsquigarrow e'_1$ (5), by (4)
 $\Delta \mid \Gamma_2, x \vdash_{\text{GF}} e_2 \rightsquigarrow e'_2$ (6), by (4)
 $(\star)_{gf}$ for Γ_2 (7), by (4)
 $(\star)_{fl}$ for Γ_2 (8), by (7) and monotonicity
 $\Delta, \Gamma_2 \mid \Gamma_1, R_1 \vdash_{\text{FL}} e_1 \rightsquigarrow e'_1$ (9), by induction on (5) and (2)
 $\Delta \mid \Gamma_2, R_2, x \vdash_{\text{FL}} e_2 \rightsquigarrow e'_2$ (10), by induction on (6) and (3)
 $(\star)_{fl}$ for R_2 (11), since reuse tokens are small
 $\Delta \mid \Gamma_1, \Gamma_2, R \vdash_{\text{FL}} \text{let } x = e_1 \text{ in } e_2 \rightsquigarrow \text{let } x = e''_1 \text{ in } e''_2$ (12), by LET and (9), (10), (8) and (11)

case RMATCH.

$\Delta \mid \Gamma, x \vdash_{\text{GF}} \text{match } x \{p_i \vdash e_i\} \rightsquigarrow \text{match } x \{p_i \vdash \text{dup}(\overline{z_i}); e'_i\}$ (1), given
 $\Delta \mid \Gamma, x, \overline{z_i} \vdash_{\text{GF}} e_i \rightsquigarrow e'_i$ (2), by (1)
 $S \mid R \Vdash \text{match } x \{p_i \vdash \text{dup}(\overline{z_i}); e'_i\} \rightsquigarrow \text{match } x \{p_i \vdash e''_i\}$ (3), given
 $S, x : n \mid R \Vdash \text{dup}(\overline{z_i}); e'_i \rightsquigarrow e''_i$ (4), by (3)
 $S, x : n \mid R \Vdash e'_i \rightsquigarrow e''_i$ (5), by (3) and repeatedly invoking RDUP
 $\Delta \mid \Gamma, x, \overline{z_i}, R \vdash_{\text{FL}} e_i \rightsquigarrow e''_i$ (6), by induction on (2) and (5)
 $\Delta \mid \Gamma, x, R \vdash_{\text{FL}} \text{match } x \{p_i \vdash e_i\} \rightsquigarrow \text{match } x \{p_i \vdash \text{dup}(\overline{z_i}); e''_i\}$ (7), by MATCH

case RDUP.

$S \mid R \Vdash \text{dup } x; e' \rightsquigarrow \text{dup } x; e''$ (1), given
 $S \mid R \Vdash e' \rightsquigarrow e''$ (2), by (1)
 $\Delta \mid \Gamma \vdash_{\text{GF}} e \rightsquigarrow \text{dup } x; e'$ (3), given
 $\Delta \mid \Gamma, x \vdash_{\text{GF}} e \rightsquigarrow e'$ (4), by (2)
 $\Delta \mid \Gamma, x, R \vdash_{\text{FL}} e \rightsquigarrow e''$ (5), by induction
 $\Delta \mid \Gamma, R \vdash_{\text{FL}} e \rightsquigarrow \text{dup } x; e''$ (6), by DUP

case RDROP.

$S \mid R \Vdash \text{drop } x; e' \rightsquigarrow \text{drop } x; e''$ (1), given
 $S \mid R \Vdash e' \rightsquigarrow e''$ (2), by (1)
 $\Delta \mid \Gamma, x \vdash_{\text{GF}} e \rightsquigarrow \text{drop } x; e'$ (3), given
 $\Delta \mid \Gamma \vdash_{\text{GF}} e \rightsquigarrow e'$ (4), by (2)
 $\Delta \mid \Gamma, R \vdash_{\text{FL}} e \rightsquigarrow e''$ (5), by induction
 $\Delta \mid \Gamma, x, R \vdash_{\text{FL}} e \rightsquigarrow \text{drop } x; e''$ (6), by DROP

case RDROPRU.

$S \mid R \Vdash r \leftarrow \text{dropru } x; e' \rightsquigarrow r \leftarrow \text{dropru } x; e''$ (1), given
 $S \mid R, r \Vdash e' \rightsquigarrow e''$ (2), by (1)
 $\Delta \mid \Gamma, x \vdash_{\text{GF}} e \rightsquigarrow r \leftarrow \text{dropru } x; e'$ (3), given
 $\Delta \mid \Gamma, r \vdash_{\text{GF}} e \rightsquigarrow e'$ (4), by (2)
 $\Delta \mid \Gamma, R, r \vdash_{\text{FL}} e \rightsquigarrow e''$ (5), by induction
 $\Delta \mid \Gamma, R, r \vdash_{\text{FL}} e \rightsquigarrow r \leftarrow \text{dropru } x; e''$ (6), by DROPRU

case RREUSE-DROP.

$S \mid R \Vdash \text{drop } x; e' \rightsquigarrow r \leftarrow \text{dropru } x; e''$ (1), given
 $S \mid R, r : n \Vdash e' \rightsquigarrow e''$ (2), by (1)
 $\Delta \mid \Gamma, x \vdash_{\text{GF}} e \rightsquigarrow \text{drop } x; e'$ (3), given
 $\Delta \mid \Gamma \vdash_{\text{GF}} e \rightsquigarrow e'$ (4), by (2)
 $\Delta \mid \Gamma, R, r \vdash_{\text{FL}} e \rightsquigarrow e''$ (5), by induction
 $\Delta \mid \Gamma, x, R \vdash_{\text{FL}} e \rightsquigarrow r \leftarrow \text{dropru } x; e''$ (6), by DROPRU