First-Order Laziness

Technical Report, Jul 11, 2025 (v2).

ANTON LORENZEN, University of Edinburgh, United Kingdom
DAAN LEIJEN, Microsoft Research, USA
WOUTER SWIERSTRA, Utrecht University, Netherlands
SAM LINDLEY, University of Edinburgh, United Kingdom

In strict languages, laziness is typically modeled with explicit thunks that defer a computation until needed and memoize the result. Such thunks are implemented using a closure. Implementing *lazy data structures* using thunks thus has several disadvantages: closures cannot be printed or inspected during debugging; allocating closures requires additional memory, sometimes leading to poor performance; reasoning about the performance of such lazy data structures is notoriously subtle. These complications prevent wider adoption of lazy data structures, even in settings where they should shine. In this paper, we introduce *lazy constructors* as a simple first-order alternative to lazy thunks. Lazy constructors enable the thunks of a lazy data structure to be defunctionalized, yielding implementations of lazy data structures that are not only significantly faster but can easily be inspected for debugging.

Additional Key Words and Phrases: Laziness, Defunctionalization, Perceus

1 Introduction

Purely functional data structures have several important advantages. Data structures implemented in a purely functional language are persistent, thread safe, and may be verified using elementary methods. *Efficient* purely functional data structures, however, often require laziness to avoid recomputation, even when implemented in a strict language [Okasaki 1999]. In a strict language, like OCaml and Racket, computations may be deferred by creating an explicit thunk. Despite the apparent simplicity of implementing laziness in this fashion, using higher-order functions has its drawbacks: thunked computations cannot be printed or inspected; allocating closures requires additional memory; reasoning about the performance of such arbitrary closures is a subtle affair.

This paper explores how to add a dash of laziness to a strict language, where computations are deferred explicitly using a *first-order data constructor*, defunctionalizing the higher-order closures programmers would otherwise write by hand. As we will show, these techniques suffice to implement purely functional data structures efficiently, reducing the time and space used by traditional implementation techniques. To sketch the main idea, consider the following definition of a stream in Koka [Leijen 2014]:

```
type stream<a>
  SCons(head : a, tail : stream<a>)
  SNil
  lazy SAppend(s1 : stream<a>, s2 : stream<a>) ->
    match s1
      SCons(x,xx) -> SCons(x,SAppend(xx,s2))
      SNil        -> s2
```

Besides the familiar data constructors, `SNil` and `SCons`, there is a third *lazy constructor* `SAppend`. Whenever we append two streams using this constructor, the operation takes constant time.

However, in contrast to the regular constructors, we never match on a lazy constructor. Instead, whenever the run-time encounters an `SAppend` constructor, the associated right-hand side of the data declaration is executed, producing a single `SCons` cell if the first stream is non-empty. Written in this style, the append of the two streams happens *on-demand*, only traversing as much of the

first stream as is necessary. To illustrate this point, we define the `take` function on streams:[1]

```
fun stream/take(xs : stream<a>, n : int) : list<b>
  if n <= 0 then Nil
  else match xs
    SCons(x,xx) -> Cons(x, stream/take(xx,n - 1))
    SNil        -> Nil
```

As this definition shows, there is no need to write a case for the `SAppend` constructor. If there are any lazy constructors in the argument stream, these are forced on-demand as the `take` function traverses its input. This is best illustrated with an example:

```
val xs : stream<int> = SCons(0,SAppend(SCons(1,SNil),SCons(2,SNil)))
```

If we call `take(xs,1)` this produces a singleton list with the number `0`, leaving the tail of the stream unchanged. If we call `take(xs,2)`, however, this evaluates the lazy `SAppend` constructor – but only enough to discover that the second element of the resulting list should be `1`. Taking three or more elements forces the entire stream. This process happens entirely under the hood and the program cannot observe that thunks have been evaluated. Laziness preserves referential transparency: a lazy thunk is indistinguishable from the value it computes.

However, for debugging or educational purposes, it would be nice to be able to peek under the hood [Gill 2000]. With lazy constructors, this is possible: the unsafe primitive `debug-show` displays the lazy constructors without forcing any further evaluation. The informal description of the behaviour of `take` is visible in the command line:

```
> take(xs,1); debug-show(xs)
SCons(0,SAppend(SCons(1,SNil),SCons(2,SNil)))
> take(xs,2); debug-show(xs)
SCons(0,SCons(1,SAppend(SNil,SCons(2,SNil))))
> take(xs,3); debug-show(xs)
SCons(0,SCons(1,SCons(2,SNil)))
```

Lazy constructors are limiting: unlike unrestricted laziness as in Haskell or the explicit thunks in strict languages, our example stream only supports a single lazy operation. If we need other lazy operations, we need to add further lazy constructors to the stream data type. As we shall see, however, most implementations of lazy data structures (e.g. as given by Okasaki [1999]) rely only on a handful of lazy operations. By making the laziness first-order, we gain the ability to inspect and optimize thunked computations in new and interesting ways.

For example, the compiler can now statically determine the runtime size of each lazy constructor: the memory location associated with each forced `SAppend` cell can for example always be reused *in-place* for the resulting `SCons` cell, instead of overwriting it with an indirection node as in most implementations of laziness. Moreover, with Perceus reference counting [Lorenzen and Leijen 2022; Reinking, Xie et al. 2021], if the matched `SCons` of `s1` happens to be unique at runtime, the next `SAppend` can reuse that memory in-place as well.

Just as first-order data types are easier to manipulate and implement efficiently than their Church encoding, the first-order approach to laziness pioneered in this paper is both efficient and effective. This paper demonstrates the applicability of lazy constructors, nails down their semantics, and benchmarks the performance of our implementation in Koka. More specifically, this paper makes the following contributions:

- We illustrate the use of first-order laziness through a series of examples drawn from Okasaki's book on functional on functional data structures [Okasaki 1999], such as the Bankers Queue and Realtime Queue (Section 2). Our implementation using lazy constructors arises naturally from defunctionalizing the thunked closures used in Okasaki's original implementation (Section 4).
- We formalize the behaviour of lazy constructors as a modest extension of Launchbury's natural semantics for lazy evaluation [Launchbury 1993] and prove that this extension preserves type

[1] In Koka, we can *locally qualify* an identifier, as in `stream/take`. A bare `take` is usually resolved to the right definition based on the type context, but we can always use the fully qualified name as well to distinguish it for example from `list/take`.

soundness and referential transparency (Section 5).
- We present a small step semantics, which forms the basis of the implementation in Koka. The first-order nature of lazy constructors enables new compiler optimizations that are not possible in general: avoiding indirection nodes entirely; re-using memory; and running in constant stack space. We justify these compiler optimizations using equational reasoning (Section 6).
- We implement lazy constructors in Koka and benchmark all lazy queues and heaps given by Okasaki [1999]. Our benchmarks show that lazy data structures, implemented using lazy constructors, are always faster than the same data structures implemented using traditional thunks, and can come close to their strict implementations even in sequential settings where laziness provides no benefit (Section 8).

2 Programming with First-Order Laziness

To illustrate the importance of laziness, even in a strict language, we revisit the Bankers Queue example by Okasaki [1999]. It is a typical example of a functional data structure that uses laziness to obtain better amortized time complexity bounds in a persistent setting.

2.1 A Strict Bankers Queue using Lists

To warm up, we first define a *strict* Bankers Queue; the next section will give an alternative lazy implementation using streams. A Bankers Queue consists of a pair of lists, where new elements are appended to the rear list `ys`, and elements are removed from the front list `xs`:

```
struct queue<a>   // queue with elements ‘xs ++ reverse(ys)‘
  xs : list<a>    // front list
  n  : int        // length of the front
  ys : list<a>    // rear list (to be reversed)
  m  : int        // length of the rear
```

The queue maintains the invariants that `length(xs)==n`, `length(ys)==m`, and `n>=m`. As the rear list grows and the front list shrinks, the queue becomes unbalanced. To ensure the desired invariant is maintained, Okasaki defines a `balance` function that sometimes moves the rear list to the front list:

```
fun balance( Queue(xs,n,ys,m) : queue<a> ) : queue<a>
  if n >= m
    then Queue(xs,n,ys,m)
    else Queue(xs ++ reverse(ys), n + m, Nil, 0)
```

The the enqueue and dequeue operations ensure the result queues are always balanced:

```
fun snoc( Queue(xs,n,ys,m) : queue<a>, y : a ) : queue<a>
  balance(Queue(xs, n, Cons(y,ys), m + 1))

fun uncons( Queue(xs,n,ys,m) : queue<a> ) : maybe<(a,queue<a>)>
  match xs
    Cons(x,xx) -> Just((x, balance(Queue(xx, n - 1, ys, m))))
    Nil        -> Nothing
```

However, as noted by Okasaki, this implementation is not always very efficient. A rebalancing step may take time linear in the length of the queue (`xs ++ reverse ys`). In a *persistent* setting, there may be many shared references to a single queue; unless we ensure the rebalancing computation is also shared, each reference may need to redo the rebalancing work. To illustrate this point, consider the following code snippet:

```
val q = Queue(xs,n,xs,n)
for(1,n, fn(i) snoc(q,i))
```

In this example, we construct an ‘almost unbalanced’ queue `q`. Each `snoc` operation in the loop requires the queue to be rebalanced, where each rebalancing requires $2n$ steps. Consequently, the entire loop takes quadratic time. If, however, the rebalancing work is *shared* between the different calls to `snoc`, then the loop runs in linear time. Even in this strict and persistent setting, there is a

clear need for some memoization in order to avoid such recomputation.

2.2 A Lazy Bankers Queue using Streams

To obtain the optimal amortized time complexity of the Bankers Queue in a persistent setting, we need to ensure that the result of the rebalancing is shared between all copies of the queue. Rather than using lists, we use a variation of the streams from the introduction instead:

```
struct queue<a>
  xs : stream<a>  // the front stream
  n  : int        // length of the front
  ys : stream<a>  // rear stream
  m  : int        // length of the rear
```

However, unlike our streams from the introduction, we not only need an append operation, but also a reverse operation:

```
type stream<a>
  SNil
  SCons(head : a, tail : stream<a>)
  lazy SAppend(s1 : stream<a>, s2 : stream<a>) ->
    match s1
      SCons(x,xx) -> SCons(x,SAppend(xx,s2))
      SNil        -> s2
  lazy SReverse(s : stream<a>, acc : stream<a>) ->  // accumulating reverse
    match s
      SCons(x,xx) -> SReverse( xx, SCons(x,acc) )
      SNil        -> acc
```

The rebalancing function now uses the *lazy constructors* to defer and share rebalancing:

```
fun balance( Queue(xs,n,ys,m) ) : queue<a>
  if n >= m
    then Queue(xs,n,ys,m)
    else Queue(SAppend(xs,SReverse(ys,SNil)), n + m, SNil, 0)
```

Since the *only* operations we need to rebalance the queue are append and reverse, we only need two lazy constructors – `SAppend` and `SReverse`. Moreover, the definitions of `snoc` and `uncons` remain unchanged as we never match on lazy constructors. The `balance` function introduces lazy constructors, but defers the associated work. Consider the loop we saw previously:

```
val q = Queue(xs,n,xs,n)
for(1,1000, fn(i) snoc(q,i))
```

Each call to `snoc` simply creates a delayed computation for the rebalancing in constant time (as `SAppend(xs,SReverse(ys,SNil))`), which only takes constant time. In contrast, the `uncons` operation pattern matches on the front stream, which may trigger evaluation of the lazy constructors and can thus take linear time. Still, Okasaki shows that this implementation of the bankers queue has constant amortized time complexity.

2.3 Lazy Match

Unlike traditional implementations of laziness, lazy constructors remain first order. Consequently, they can be printed for the sake of debugging:

```
> val xs = SCons(1,SCons(0,SNil))
> val q0 = Queue(xs,2,xs,2)
> val q  = snoc(q,2)
> debug-show(q)
Queue( SAppend(SCons(1,SCons(0,SNil)),SReverse(SCons(2,SCons(1,SCons(0,SNil)))) ), 5, SNil, 0)
```

Of course, since `q` is persistent, we can `uncons` an element and still observe the original queue:

```
> val _ = uncons(q)
> debug-show(q)
Queue( SCons(1, SAppend(SCons(0,SNil),SReverse(SCons(2,SCons(1,SCons(0,SNil)))) ), 5, SNil, 0)
```

The reader may be startled at this point: clearly the queue `q` has changed! Doesn't this break referential transparency? The answer is no: though the stream has indeed changed, this cannot be observed since any attempt to match on `q` never yields a lazy constructor. In fact, this is exactly why any thunk can be overwritten with its value without breaking referential transparency.

However, it can be useful for debugging to peek under the hood during evaluation of a lazy data structure: this is what the `debug-show` function does. This function is implemented using an additional unsafe primitive, `lazy match`, that can observe lazy constructors without forcing evaluation.

The lazy match construct is used extensively to implement first-order laziness. In particular, the Koka compiler inserts an additional `eval` call whenever a programmer matches on a data type with lazy constructors. The compiler generated `eval` function evaluates the argument to weak head normal form. The code corresponding to the `uncons` function becomes:

```
fun uncons( Queue(xs,n,ys,m) )
  match stream/eval(xs)      // compiler inserts an ‘eval‘ automatically
    SCons(x,xx) -> Just((x, balance(Queue(xx, n - 1, ys, m))))
    SNil        -> Nothing
```

The `eval` function uses the lazy match primitive, inserting the code associated with each lazy constructor in the corresponding branch, roughly like:

```
// compiler generated
fun stream/eval(s : stream<a>)
  lazy match s
    SAppend(s1,s2)  -> match s1
      SCons(x,xx) -> lazy-update(s, SCons(x,SAppend(xx,s2)))
      SNil        -> lazy-update(s, eval(s2))
    SReverse(s1,acc) -> match s1
      SCons(x,xx) -> lazy-update(s, eval(SReverse(xx, SCons(x,acc))))
      SNil        -> lazy-update(s, eval(acc))
    _ -> s
```

where the `lazy-update` primitive updates the root node with the result. In practice, we generate a more efficient version where we do not use stack space unnecessarily. While this `stream/eval` function uses the stack in the last three branches (where `eval` is not a tail-call), we will derive a more efficient version in Section 6.

2.4 The Bankers Queue with Logarithmic Worst-Case Time Complexity

While our Bankers Queue has constant *amortized* time complexity, its worst-case time complexity is still linear in the size of the queue. The reversal is monolithic: once the `SAppend` is fully evaluated, we need to completely reverse the second list to find the last element. How can we ensure that these queues have better worst-case complexity?

Okasaki [1999] reimplements the Bankers Queue using *rotations* that combine appending with reversal. However, our first-order lazy constructors support a simpler solution – albeit one that requires further language support. The main idea is to evaluate at most one lazy constructor instead of recursively evaluating up to weak head normal form. If we do this for the reversed tail each time we evaluate an `SAppend` constructor, then we can ensure that by the time all `SAppend`'s are done, the tail is fully reversed. To evaluate only one lazy constructor, the compiler generates a second evaluation function, called `eval-one`:

```
// compiler generated
fun stream/eval-one(s : stream<a>) : stream<a>
  lazy match s
    ...
    SReverse(s1,acc) -> match s1
      SCons(x,xx) -> lazy-update(s, SReverse(xx, SCons(x,acc)))
      SNil        -> lazy-update(s, acc)
    _ -> s
```

As we can see, it is almost equivalent to the `eval` function. The key difference is that `eval-one` is no longer recursive. For example, if the first list is `Nil`, the `eval-one` function returns the accumulated list – whether it is in weak head normal form or not.

Using this primitive, we can reduce the worst-case time complexity from linear to logarithmic, by only adding a single line to the code associated with the `SAppend` constructor:

```
type stream<a>
  ...
  lazy SAppend( s1 : stream<a>, s2 : stream<a> ) ->
    stream/eval-one(s2)  // for each SAppend, evaluate also one SReverse
    match s1
      SCons(x,xx) -> SCons(x, SAppend(xx,s2))
      SNil        -> s2
```

The key here is that `s2` always contains the `SReverse` constructor. In this fashion, we evaluate one step of the reversal every time we invoke the `SAppend` constructor. As `s2` is only one element longer than `s1` when rebalance, by the time the `SAppend` hits the `SNil` case the `SReverse` is almost done – we pay a little more during each `SAppend` step, but gain more predictable performance overall.

2.5 Avoiding Stack Overflows from Recursive Evaluation

The code associated with the `SAppend` constructor seems entirely innocent. On closer inspection, however, when there are nested `SAppend` constructors, this may trigger recursive evaluation:

```
type stream<a>
  ...
  lazy SAppend( s1 : stream<a>, s2 : stream<a> ) ->
    match s1                                  // may evaluate an SAppend recursively!
      SCons(x,xx) -> SCons(x, SAppend(xx,s2))
      SNil        -> s2
```

As this call is not tail-recursive, it requires stack space. If the `s1` stream is a long sequence of unevaluated `SAppend` constructors, this may ultimately lead to a stack overflow.

This is not a purely theoretical concern. When working on our benchmarks, we discovered that the Physicists Queue, Implicit Queue, and Binomial Heaps as presented by Okasaki [1999] require stack space linear in the size of the queue. This can be a problem in practice: when two million elements were subsequently inserted into the Physicists Queue, the stack overflowed in both Koka 3.1.3 and OCaml 4.14.2 [2]. We are not aware of any technique that avoids this problem with traditional lazy thunks.

However, as we show in detail in Section 6.5, this problem *can* be resolved with lazy constructors. As we compile a specialized `eval` function for each data type, we can specialize `eval` to evaluate nested thunks tail-recursively using an in-place Schorr-Waite traversal [Leijen and Lorenzen 2023 2025; Lorenzen et al. 2023; Schorr and Waite 1967], illustrated in Figure 1. If a programmers wants to ensure that an argument of a lazy constructor is fully evaluated before the constructor is evaluated, they can put an exclamation mark, e.g. `!xs`, before the name of the argument. This indicates that Koka's runtime should force the argument before it begins to evaluate the lazy constructor[3]:

```
  lazy SAppend( !xs : stream<a>, ys : stream<a> ) -> ...
```

With the `!` annotation, before the `SAppend` constructor is evaluated, the runtime first evaluates the `xs` stream. This does not require any extra stack- or heap space, since `eval` can keep track of all lazy constructors under evaluation using an in-place *zipper* [Huet 1997; Lorenzen et al. 2024] stored in the lazy constructors themselves. In this fashion, we avoid allocating arbitrary stack space, even if `SAppend` constructors are nested.

[2] As measured on an Apple M1 which has a hard limit for stack size of 65mb.

[3] At the moment this is not yet implemented in Koka but we hope to have it working soon according to the rules in Section 6.5.

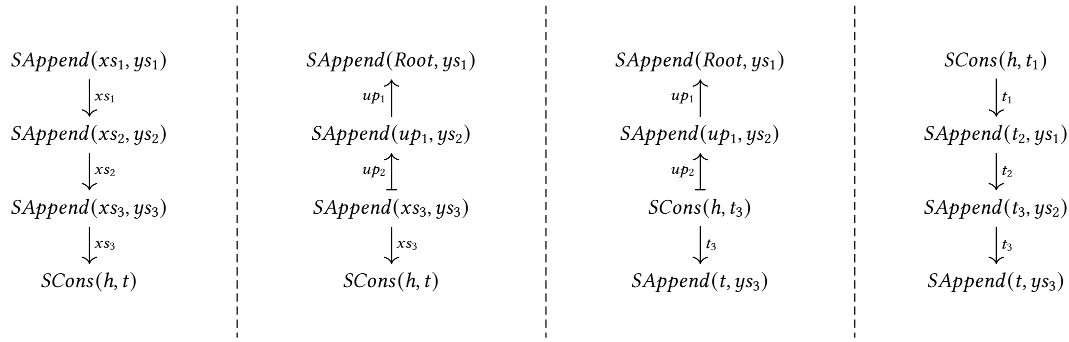

Fig. 1. In a Schorr-Waite traversal, we first descend and reverse the pointers until we find a normal constructor. Then we evaluate the lazy constructors from the bottom-up. From left to right: (1) nested lazy constructors: to force the topmost constructor we first have to force the ones below it, (2) the runtime has descended to the bottom-most constructor and reversed the pointers, (3) the bottom-most constructor has been evaluated and we follow the 'up' pointer, (4) the final result once evaluation is complete.

2.6 The Realtime Queue with Constant Time Complexity

The Bankers Queue operations are amortized constant-time, but still have logarithmic worst-case time complexity. When the queue rebalances, the front stream may start with another `SAppend` constructor. If this is the case, rebalancing creates nested `SAppend` constructors. Fortunately, the number of nested constructors is bounded: since we only rebalance when the front stream is shorter than the rear stream, we know that the number of nested `SAppend` thunks is logarithmic in the total length of the queue. A call to `uncons` may have to perform one step for each of the `SAppend` constructors in the front stream, leading to a logarithmic worst-case time complexity.

To ensure our operations have constant worst-case time complexity, we need the front stream to evaluate one step *every time* we perform an operation on the queue. Okasaki [1999] (Section 7.2) proposes to do this by adding a 'schedule parameter'. This schedule parameter is a suffix of the front stream; the schedule starts with the *only* `SAppend` constructor in the entire stream. Maintaining this invariant is the key ingredient to obtain constant worst-case time complexity.

```
struct queue<a>
  front : stream<a>
  rear  : list<a>
  sched : stream<a>
```

There are several other modifications to the queue type. Note that the rear queue is now a list rather than a stream; rather than store the lengths of the front and rear stream, the schedule determines when the queue must be rebalanced:

```
fun balance( Queue(front,rear,sched) : queue<a> ) : queue<a>
  match sched
    SCons(_,s) -> Queue(front,rear,s)
    SNil       -> val f = SAppend(front,SReverse(rear,SNil)) in Queue(f,Nil,f)
```

In each rebalancing operation, the schedule is evaluated to weak head normal form. If the schedule is non-empty, its tail becomes the new schedule. Rebalancing happens when the schedule is empty – and hence the `front` stream is fully evaluated. In that case, the new schedule is initialized to `f`, shared with the new front of the queue. As the front stream is fully evaluated, `f` will never contain nested `SAppend` constructors. This ensures the desired constant worst-case time complexity.

The only thing that remains to be done, is implement the `snoc` and `uncons` operations:

```
fun snoc( Queue(front,rear,sched) : queue<a>, x : a) : div queue<a>
  balance( Queue(front, Cons(x,rear), sched) )

fun uncons( Queue(front,rear,sched) : queue<a> ) : div maybe<(a,queue<a>)>
  match front
    SCons(x,xx) -> Just((x, balance(Queue(xx,rear,sched))))
    SNil        -> Nothing
```

As both operation call `balance`, they are guaranteed to advance the schedule, as required. In the previous examples we mostly relied on laziness to defer computation, but in the Realtime Queue we also rely on the computation being *shared* – in this case between the front stream and the schedule.

3 Advanced Examples of Lazy Constructors

In this section, we consider more advanced examples of using lazy constructors.

3.1 The Sieve of Eratosthenes

The Sieve of Eratosthenes is a classic example of a lazy algorithm. A simple implementation is featured on the Haskell homepage:

```
primes = filterPrime [2..] where
  filterPrime (p:xs) =
    p : filterPrime [x | x <- xs, x ‘mod‘ p /= 0]
```

In the program, `primes` is an infinite list of all prime numbers. It is generated by searching the numbers from 2 to infinity, where every time a prime number is found, all multiples of that number are filtered out of the list.

In Koka, we can implement this algorithm by defining a type for the lazy thunks that are involved in the computation above. We start by creating a type for infinite lists of prime numbers:

```
type primes
  Prime( x : int, xs : primes )
```

So far this type is relatively useless: Since this is an inductive definition, it is not possible to create an inhabitant of this type. However, we can use a lazy constructor to create an infinite stream:

```
  lazy From( n : int, _p : pad ) ->
    Prime( n, From(n + 1, Pad) )
```

This lazy constructor creates the infinite list of numbers starting from `n` and incrementing by 1 each time. Crucially, we can now create an inhabitant of the `primes` type by calling `From(2, Pad)`. We pass `Pad : pad` as a dummy value to ensure that the `From` constructor has the same size as the `Prime` constructor, which makes reuse possible.

While `From(2, Pad)` is an infinite list of numbers, it is not yet an infinite list of *prime* numbers. To filter out the non-prime numbers, we can define a `Filter` constructor:

```
  lazy Filter( p : int, xs : primes ) ->
    match xs
      Prime(x, xs) ->
        if x % p != 0
        then Prime(x, Filter(p, xs))
        else Filter(p, xs)
```

This constructor mirrors the common `filter` function, where the usual recursive calls to `filter` are replaced by re-constructions of the `Filter` constructor. Finally, we can define the `Sieve` constructor as the defunctionalization of the `filterPrime` function above:

```
  lazy Sieve(xs : primes, _p : pad) ->
    match xs
      Prime(p, xs) ->
        Prime(p, Sieve(Filter(p, xs), Pad))
```

Then we can define the infinite list of primes as `val primes = Sieve(From(2, Pad), Pad)`.

Compared to the Haskell definition, our implementation is significantly more verbose. Partly this is due to a lack of syntax: `From` and `Filter` have special support in Haskell by the `[2..]` syntax and list comprehensions. In contrast, the `Sieve` constructor is not more verbose than the `filterPrime` function in Haskell.

An advantage of our implementation is that it is more explicit about the laziness involved in the computation. For example, we can easily print the `primes` stream after evaluating the first five prime numbers:

```
Prime(2, Prime(3, Prime(5, Prime(7, Prime(11,
  Sieve(Filter(11, Filter(7, Filter(5, Filter(3, Filter(2, From(12, Pad)))))), Pad))))))
```

As you can see, the `Sieve` constructor accumulates a list of filters that are applied to the remaining `From` stream. That means that every time we try to evaluate another prime number, our implementation has to traverse as many filters as the number of primes we have already found. This makes this implementation of the sieve rather inefficient. However, the same issue is also present in the Haskell version [Gibbons 2025; O'Neill 2009]; just hidden by the opaqueness of implicit laziness.

3.2 Repmin: Tying the Knot

The classic repmin problem concerns the challenge of replacing all elements in a tree of integers with the minimum integer in the tree, in a single pass [Bird 1984]. The restriction to a single pass implies that the traversal has to change the elements to the minimum element before the exact value of the minimum element is known. Consider a tree such as:

```
type tree
  Leaf( n : elem )
  Node( l : tree, r : tree )

type elem
  Elem( n : int )
```

We have to make `elem` its own type instead of just using `int` directly in the tree. In Koka, an integer may be unboxed, in which case the integer is written directly into its surrounding structure without further indirection. But then we would have to know the precise value of the integer by the time we are updating the leaf. Instead, the repmin algorithm requires us to write a pointer into each leaf which points to a separate memory location, which can later be updated in a single step.

The `repmin` function is a simple traversal of the tree that writes `m` into each leaf and returns the minimum element:

```
fun repmin( m : elem, t : tree ) : (elem, tree)
  match t
    Leaf(n) -> (n, Leaf(m))
    Node( l, r ) ->
      val (nl, l') = repmin( m, l )
      val (nr, r') = repmin( m, r )
      (min( nl, nr ), Node( l', r' ))
```

The crucial part of the problem is now to connect `m` to the return value of `repmin`. Typically, this is done by "tying the knot", where we create a cyclic definition:

```
fun replace-min( t1 : tree ) : tree
  val (m, t2) = repmin(m, t1)
  t2
```

However, Koka does not support cyclic definitions nor tying the knot directly, which makes our `replace-min` function invalid. Instead, we can use references to tie the knot. We create a new lazy constructor `RepMin` in `elem`, which is passed a reference to the minimum element. We evaluate the lazy constructor by reading from the reference:

```
lazy type elem<h>
  Elem( n : int )
  lazy RepMin( m : ref<h,elem<h>> ) ->
    unsafe-total { !m }

fun replace-min( t : tree<h> ) : <div,st<h>> tree<h>
  val r = ref(Elem(0))
  val (m, t) = repmin(RepMin(r), t)
  set(r, m)
  t
```

Koka uses a type parameter `h` for reference cells (akin to Haskell's `ST` monad). Functions reading or writing from a reference get the `st<h>` effect, and we now have to pass this `h` parameter to the types `elem<h>` and `tree<h>`. In `replace-min`, we create a new reference `r` with a dummy value and use the `repmin` traversal to write `RepMin(r)` into all leaves of the tree. Then we set `r` to the minimum element, returned by `repmin`. Once a leaf of the tree is evaluated, the `RepMin` lazy constructor reads from the reference `!r` and overwrites itself with the content of the reference.

We have to call `unsafe-total` in `RepMin` since lazy constructors are currently restricted in the effects they may use and thus can not usually read from references. This design choice prevents lazy constructors from performing IO operations — which in Haskell give rise to the infamous "lazy IO" problem. However, we are working on loosening this restriction for selected algebraic effects.

3.3 Step-wise Evaluation with Unique Data in Thunks

The Hood-Melville Queue [Hood and Melville 1980] is another data structure discussed by Okasaki [1999]. Similar to the Realtime Queue, the Hood-Melville Queue reverses a `rear` list and appends it to a `front` list. But unlike the Realtime Queue, the Hood-Melville Queue is strict and uses no laziness. It can get away with this by performing the expensive append-and-reverse one step at a time: each `snoc` and `uncons` step performs a constant amount of work on the restructuring, which means that both operations use constant time in the *worst-case*.

However, it can be quite tricky to split a function like `append` into pieces that can be performed one step at a time. Consider the typical `append` function:

```
fun append( xs : list<a>, ys : list<a>) : list<a>
  match xs
    Cons( x, xx ) -> Cons( x, append( xx, ys ) )
    Nil -> ys
```

This function can not easily be split into pieces: to do so, we would have to record all internal state of this function, but that is impossible since much of it is on the stack. If we just recorded the tuple `(xx, ys)` as the internal state, we would lose the elements `x` that are on the call-stack. For this reason, Okasaki's implementation of the Hood-Melville Queue splits `append` into two functions:

```
fun reversing( xs : list<a>, ys : list<a>, acc : list<a>) : list<a>
  match xs
    Cons( x, xx ) -> reversing( xx, ys, Cons(x, acc) )
    Nil -> appending( acc, ys )

fun appending( acc : list<a>, ys : list<a> ) : list<a>
  match acc
    Cons( x, xx ) -> appending( xx, Cons( x, acc ) )
    Nil -> ys

fun append( xs : list<a>, ys : list<a> )
  reversing( xs, ys, Nil )
```

This implementation is easy to split into steps. (Mutally) tail-recursive functions act a bit like state-machines where all state is encapsulated in the arguments. We can thus create a data type with constructors `Reversing` and `Appending` which store the arguments to the respective functions. A single step corresponds to a single unfolding of the respective function. However, this design is

quite inefficient: it requires two traversals of the list `xs`.

In Koka, we can derive a better implementation using First-Class Constructor Contexts [Lorenzen et al. 2024] which are an efficient representation of a data structure with a hole. For example, we can create a one-element list without as `val l = ctx Cons( x, _ )`. We can then set the tail of the list later by writing, say, `l ++. Nil`. Similarly, we can append two constructor contexts using the `++` operation, which writes the root of the second context into the hole of the first context. Both operations take constant time if the (first) context is unique and otherwise require a copy.

We can use this feature to implement a one-step `append` function as:

```
fun append( xs : list<a>, ys : list<a>, acc : ctx<list<a>>)
  match xs
    Cons( x, xx ) -> append( xx, ys, acc ++ ctx Cons( x, _ ) )
    Nil -> acc ++. ys
```

Unlike the typical `append` function, we keep an extra constructor context `acc`, which is successively built-up from the list `xs`. But unlike the earlier tail-recursive version, we can assemble the result `acc ++. ys` in constant time and do not require an extra `appending(acc, ys)` traversal. For this reason, this version is significantly faster in practice. However, this version has a new problem: constructor contexts need an expensive, linear-time copy if they are shared. This can happen in the Hood-Melville Queues and if we just used the above design, we would lose the amortization guarantee of Hood-Melville Queues.

We can regain the amortization bound if we use the constructor context in a thunk. Even if there are several references to a thunk, the runtime guarantees that the thunk is only evaluated once. If a unique data structure is placed in a thunk, it will continue to be unique during evaluation of the thunk. This makes it possible to schedule computations involving data structures that would be expensive to copy in a persistent setting.

We can define a `thunk` type that contains an `Append` constructor for the partial computation:

```
type thunk<a>
  Done( done : list<a> )
  lazy Append( xs : list<a>, ys : list<a>, acc : ctx<list<a>> ) ->
    match xs
      Cons( x, xx ) -> Append( xx, ys, acc ++ ctx Cons( x, _ ) )
      Nil -> Done( acc ++. ys )
```

We can then easily perform a single step of the computation `t` by calling `eval-one(t)`. Even if several references to `t` exist, the `acc` value will retain its unique reference count.

3.3.1 Step-wise Evaluation using Traditional Thunks. The same effect can also be achieved with traditional thunks, but is significantly less elegant. To simulate this, we create a `schedule<a>` type. It is not possible to partially force a thunk, and so our only hope is by creating a new thunk in each step. We can define a datatype for schedules as follows:

```
type schedule<a>
  Continue( t : thunk<schedule<a>> )
  Done( done : list<a> )
```

Then we can create a function to create a schedule for a thunk. Whenever the computation has not finished, we create a new thunk for the remaining computation:

```
fun append( xs : list<a>, ys : list<a>, acc : ctx<list<a>> ) : schedule<list<a>>
  Continue(delay {
    match xs
      Cons( x, xx ) -> append( xx, ys, acc ++ ctx Cons( x, _ ) )
      Nil -> Done( acc ++. ys )
  })
```

In practice, the traditional lazy solution has a significant overhead in this case, but the version using a lazy constructor makes the schedule cheap. We believe that this technique could be especially useful for data structures that only obtain good performance when there is only a single reference such as arrays or hashtables.

3.4 Rewriting Laziness during Reduction

Let's return to our earlier `stream<a>` type:

```
type stream<a>
  SNil
  SCons(head : a, tail : stream<a>)
  lazy SAppend(s1 : stream<a>, s2 : stream<a>) ->
    match s1
      SCons(x,xx) -> SCons(x,SAppend(xx,s2))
      SNil        -> s2
  lazy SReverse(s : stream<a>, acc : stream<a>) ->
    match s
      SCons(x,xx) -> SReverse( xx, SCons(x,acc) )
      SNil        -> acc
```

The constructors of this type enjoy several algebraic laws, which derive from the underlying operations. All of these laws make for excellent rewrite rules and replacing the left-hand side expression with the right-hand side can often significantly improve a program:

(1) `SAppend(xs, SNil) = xs`: This law can be used to avoid the traversal of `xs` altogether.

(2) `SAppend(SAppend(xs, ys), zs) = SAppend(xs, SAppend(ys, zs))`: This law can be used to reduce the left-nesting of `SAppend`. The first $|xs|$-times that the left-hand side is forced, we have to recursively force the inner `SAppend`. The cost of the left-hand side is thus $2 * |xs| + |ys|$, while the cost of the right-hand side is just $|xs| + |ys|$.

(3) `SAppend(SReverse(xs, ys), zs) = SReverse(xs, SAppend(ys, zs))`: This law reduces the left-nesting similar to the second.

(4) `SReverse(SAppend(xs, ys), zs) = SReverse(ys, SReverse(xs, zs))`: This law makes the computation less monolithic: the right-hand side produces its first elements as soon as ys has been reversed, while the left-hand side produces elements only once both xs and ys have been reversed. This reduces pause-times from forcing laziness and can also save work if the reversal of xs is not needed.

(5) `SReverse(SReverse(xs, ys), zs) = SReverse(ys, SAppend(xs, zs))`: This law also makes the computation less monolithic and replaces the double reversal of xs by a single traversal to append.

In functional programming languages like Haskell, such rewrite rules can be applied at compile-time. But sometimes, a compile-time pass will miss optimizations. In this case, it can happen that lazy computations will end up looking like one of the left-hand side at runtime. Can we still apply such a rewrite rule at runtime?

It turns out that lazy constructors allow this! We can implement an improved `stream<a>` type as follows:

```
type stream<a>
  SNil
  SCons(head : a, tail : stream<a>)
  lazy SAppend(s1 : stream<a>, s2 : stream<a>) ->
    lazy match s2
      SNil -> s1
      _ -> lazy match s1
        SCons(x, xx)     -> SCons( x, SAppend(xx,s2) )
        SNil             -> s2
        SAppend(ys, zs)  -> SAppend( ys, SAppend(zs, s2) )
        SReverse(ys, zs) -> SReverse( ys, SAppend(zs, s2) )
```

```
lazy SReverse(s : stream<a>, acc : stream<a>) ->
  lazy match s
    SCons(x, xx)     -> SReverse( xx, SCons(x, acc) )
    SNil             -> acc
    SAppend(xs, ys)  -> SReverse( ys, SReverse(xs, acc) )
    SReverse(xs, ys) -> SReverse( ys, SAppend(xs, acc) )
```

The `SAppend` constructor checks *dynamically* at runtime whether `s2` is `SNil` and in that case simply returns `s1` as suggested the first law above (we use `lazy match s2` instead of `match s2` here to avoid forcing `s2` if it is lazy). Furthermore, it checks whether `s1` is a `SAppend` or `SReverse` constructor and if so, it applies the second or third law. The `SReverse` constructor checks dynamically whether `s` is `SAppend` or `SReverse` and if so applies the fourth or fifth law respectively.

In practice, the Bankers Queue benefits slightly from using our new rewriting `stream<a>` type. This seems to be since the `front` stream of the Bankers Queue consists of a sequence of left-nested `SAppend` constructors, which are flattened by the second law.

3.4.1 *Dangers of Lazy Match.*

When using `lazy match` in this way, most of the guarantees of laziness go out of the window.

If we got the laws wrong, our implementation would not be referentially transparent. For example, we might falsely write:

```
lazy SAppend(s1 : stream<a>, s2 : stream<a>) ->
  lazy match s2
    SNil -> SNil
```

Consider a term such as `SAppend(xs, SAppend(SNil, SNil))`. If we evaluate the inner `SAppend` first, we obtain `SAppend(xs, SNil)`, which reduces to `SNil`. However, if we (recursively) evaluate the outer `SAppend` first we obtain `xs`. A wrong use of `lazy match` can thus break referential transparency and the confluence of evaluation.

Furthermore, even a correct implementation can duplicate work. For example, consider:

```
val s1 = SAppend(xs, ys)
val s2 = SAppend(s1, zs)
```

In our simpler implementation, any work that is performed on evaluating `s1` can be shared with the computation `s2`. In fact, as we saw in Section 3.3, it is an important property that thunks are evaluated at most once. However, if we apply the second law to `s2`, the resulting value is:

```
val s2 = SAppend(xs, SAppend(ys, zs))
```

Here, we have created *new* thunks to append `xs` and `ys`. In particular, there is no more sharing with the thunk of `s1`, since all three thunks describe distinct computations.

In this case, the application of the algebraic law saves enough work that the result has to take no more steps than the original. But there can still be a performance degradation. Consider the case where `xs` is unique and we evaluate the streams as `eval(s1); eval-one(s2)`. If the law is not applied, `s1` can be evaluated fully in-place, even though there are several references to the thunk itself. But if we apply the law, then there will be several references to `xs`, which means that the evaluation of `s1` has to allocate new data.

4 Illuminating First-Order Laziness

To develop a deeper intuition for lazy constructors, we begin by showing how these arise naturally as the defunctionalized version of explicit thunks. To do so, we recreate the stream definition used in the previous section, starting from the more familiar implementation from the literature. We begin by defining the standard stream interface, without using any lazy constructors, but instead deferring computations with explicit thunks:

```
alias stream<a> = thunk<streamcell<a>>
type streamcell<a>
  SCons( head : a, tail : stream<a> )
  SNil
```

A stream is a list where the list cells are separated by lazy thunks (using the style of Wadler et al. [1998] and Okasaki [1999, section 4.2]). This makes it possible to perform operations like appending streams in constant time initially, where the linear-time append is only evaluated once that part of the list is reached. Such delayed computations are encapsulated in thunks, which take a closure that is only evaluated when needed. Thunks have the typical interface:

```
type thunk<a>
fun delay( f : () -> a )  : thunk<a>
fun force( t : thunk<a> ) : div a
```

The `delay` function creates a thunk from a computation and `force` evaluates the thunk. In Koka, we use the `div` effect to indicate that forcing a thunk may diverge. Under the hood, the implementation ensures that the computation is run at most once: after the first call to `force`, its result is memoized and returned in constant time upon every subsequent call. We can use this API to implement the familiar `append` and `reverse` functions on streams:

```
fun sappend( s1 : stream<a>, s2 : stream<a> ) : div streamcell<a>
  match s1.force
    SCons(x,xx) -> SCons(x, append(xx,s2))
    SNil        -> s2.force
fun append( s1 : stream<a>, s2 : stream<a> ) : div stream<a>
  delay{ sappend(s1,s2) }
fun sreverse( s1 : stream<a>, s2 : stream<a> ) : div streamcell<a>
  match s1.force
    SCons(x,xx) -> sreverse(xx, delay{ SCons(x,s2) })
    SNil        -> s2.force
fun reverse( s1 : stream<a> ) : div stream<a>
  delay{ sreverse(s1, delay{ SNil }) }
```

Lazy Function Syntax. Wadler et al. [1998] and Okasaki [1999, section 4.2] suggest a special syntax for functions like `append` and `reverse`. They note that these functions typically contain a top-level `delay { ... }` construct, followed by a `match` statement and forces in all branches. If their syntax was implemented in Koka, we could write the `append` function as:

```
lazy fun append( s1 : stream<a>, s2 : stream<a> ) : div stream<a>
  match s1.force
    SCons(x,xx) -> scons(x, append(xx,s2))
    SNil        -> s2
```

Here, we could omit the top-level `delay` construct and the `force` call on `s2`. However, we would still need to force `s1` manually. Unfortunately, the syntax adds a force call in the first branch, which would require us to delay the constructor manually:

```
fun scons( x : a, xx : stream<a> ) : stream<a>
  delay{ SCons(x,xx) }
```

A sufficiently smart compiler could remove the immediate force of the delay in `scons` in the first branch and compile this new `append` to the previous one. However, instead of supporting this syntax in Koka directly, we will see that lazy constructors naturally lead us to a very similar syntax.

4.1 Lazy Constructors for Thunks

As the first step towards recreating the streams used in the previous section, we begin by implementing the `thunk<a>` interface using lazy constructors:

```
type thunk<a>
  Memo( v : a )
  lazy Lazy( f : () -> a ) ->
    Memo( f() )
```

This type has two constructors. The `Memo` constructor stores a value of type `a`. The `Lazy` constructor is a *lazy constructor* that can be used to construct a values of type `thunk<a>`. Remember, the lazy

constructor is never observable in a match statement; any match on a value of type `thunk<a>` will force evaluation of the corresponding expression, `Memo( f() )`. Using this definition, we can easily simulate the typical thunk interface from strict languages:

```
fun delay( f : () -> a ) : thunk<a>
  Lazy(f)

fun force( t : thunk<a> ) : a
  match t
    Memo(v) -> v
```

The power of lazy constructors is that they allow us to *specialize* thunks to the computations they contain. By analyzing the possible closures `f` that may be stored in `Lazy`, we can specialize the higher-order `thunk<a>` type for our program. We define the first-order `stream<a>` type as:

```
type stream<a> =
  Memo( v : streamcell<a> )
  lazy SAppend( xs : stream<a>, ys : stream<a> ) ->
    Memo( sappend(xs, ys) )
  lazy SReverse( xs : stream<a>, acc : stream<a> ) ->
    Memo( sreverse(xs, acc) )
```

We keep the `Memo` constructor from the definition of `thunk<a>` but specialize `Lazy(f)` to the two computations that are used. This type now has two separate lazy constructors for these two computations, but we can still define `force` as before. We do need to update the corresponding code to append and reverse streams. The definition for `sappend`, for example, now reads:

```
fun sappend( s1 : stream<a>, s2 : stream<a> ) : div streamcell<a>
  match s1.force
    SCons(x,xx) -> SCons(x,SAppend(xx,s2))
    SNil        -> s2.force
```

Here we still need to force each stream to a `streamcell`. Where the previous definition deferred the recursive call to `sappend`, we now simply use the lazy `SAppend` constructor to the same effect.

4.2 The Cost of Laziness

Based on the last section, one might think that lazy constructors are just the defunctionalization of explicit thunks. But, in fact, they are a bit more powerful than that and allow us to fix a performance problem that arises when using explicit thunks. Let us consider a fully evaluated stream, such as:

```
> val nums = Memo(SCons(1, Memo(SCons(2, Memo(SCons(3, Memo(SNil)))))))
```

There is an indirection node between every pair of adjacent elements in the stream! This is a consequence of the encoding of streams, where the elements of `streamcell<a>` and those of `stream<a>` alternate. This means that traversing even a fully-evaluated stream will require twice as many pointer lookups as traversing a list: all `SCons` and `Memo` nodes live in different cells linked by pointers. In practice, languages like OCaml can mitigate this problem since they distinguish indirection nodes from all other values, which makes it possible to omit indirection nodes when the stream is fully evaluated on creation.

However, if a stream is the result of a lazy computation, the indirection nodes are unavoidable. For example, if we append a stream to `nums`, the `SAppend` constructor has to be rewritten as a `Memo` constructor to ensure that all references to `app` can share the memoized result:

```
> val nums' = SAppend(nums, Memo(SNil))
> debug-show(force(nums'))
Memo(SCons(1, SAppend(Memo(SCons(2, Memo(SCons(3, Memo(SNil))))), Memo(SNil))))
```

This implies that once we fully evaluate the append, all `Memo` nodes in the stream are necessary:

```
> debug-show(forceall(nums'))
Memo(SCons(1, Memo(SCons(2, Memo(SCons(3, Memo(SNil)))))))
```

These indirections are typical for the traditional approach to laziness, but they may introduce a significant performance cost compared to a strict program. Appending to a list of length n involves

only n allocations, but appending to a stream of length n requires $2n$ allocations to also create all the indirection nodes in between. In fact, without defunctionalization, this is usually even more expensive since each thunk involves another allocation for a closure and so $3n$ allocations can be necessary.

Furthermore, the indirection nodes stay in the `stream` even once the thunks are fully evaluated, where they introduce additional pointer lookups. Garbage collected languages like OCaml may remove this indirection during GC runs, but this optimization is not available in reference counted languages such as Koka [Leijen 2014] or Lean [Moura and Ullrich 2021].

This is one of the reasons why lazy data structures are often less efficient than their strict counterparts. As we show in our benchmarks, the classic lazy data structures of Okasaki [1999] are a factor of 2-3x less efficient than their strict counterparts when implemented using explicit thunks.

4.3 Fusing streams and stream cells

In contrast to traditional approaches to laziness, first-order lazy constructors allow us to remove most of these indirections. To obtain a better version of `stream<a>` with fewer indirections, we inline the definition of `streamcell<a>` in the type of streams:

```
type stream<a>
  SNil
  SCons( x : a, xx : stream<a> )
  lazy SAppend( xs : stream<a>, ys : stream<a> ) ->
    sappend(xs, ys)
  lazy SReverse( xs : stream<a>, acc : stream<a> ) ->
    sreverse(ys, acc)
```

Compared to the previous `stream<a>` declaration, we have replaced the `Memo` constructor with the `SNil` and `SCons` constructors. Furthermore, we do not return `Memo` from `SAppend` and `SReverse` and instead return the remaining stream immediately. This makes the structure of the type quite different: where the previous definition would alternate `SCons` cells with thunks, we can now mix normal and lazy constructors arbitrarily. For example, the fully evaluated stream `SCons(1, SCons(2, SNil))` is a valid inhabitant of this stream type, as is the stream containing lazy constructors `SAppend(SAppend(SNil, SNil), SNil)`.

To complete this definition, however, we need to update our `sappend` and `sreverse` functions. These no longer need to match on `Memo` constructors, but rather manipulate the streams directly. To illustrate this point, we redefine both `sreverse` and `sappend`:

```
fun sreverse( s1 : stream<a>, s2 : stream<a> ) : div stream<a>
  match s1
    SCons(x,xx) -> sreverse( xx, SCons(x,s2) )
    SNil        -> s2

fun sappend( s1 : stream<a>, s2 : stream<a> ) : div stream<a>
  match s1
    SCons(x,xx) -> SCons(x, SAppend(xx,s2) )
    SNil        -> s2
```

Note that, unlike our previous definition, we now build the accumulator of `sreverse` as a sequence of `SCons` cells with no more indirection nodes `Memo` in between.

Our new definitions of `sappend` and `sreverse` look quite similar to the lazy function definitions in Section 4.0.0.1. The main difference is that we now use constructors in places where the lazy function syntax uses (lazy) function calls that immediately delay their result. Our new definition is even a bit shorter since we can avoid a call to `force` in the pattern-match.

But where did the indirections go? We still need to memoize the result of evaluating an `SAppend` or `SReverse`. To ensure the results of evaluating these lazy constructors are still shared, Koka inserts an implicit indirection into the lazy constructor above:

```
  lazy SAppend( xs : stream<a>, ys : stream<a> ) ->
    Indirect(sappend(xs, ys))
```

When matching on a stream `s` it can now happen that `s` is an indirection node, pointing to some `s'`. In that case the runtime follows the indirection and keeps matching on `s'`. These indirection nodes, however, are only created when a lazy constructor is evaluated. For example, the accumulator built in the `sreverse` function is completely free from indirections. This reduces the memory overhead that typically arises from sharing lazy computations.

4.4 In-place Reuse of Lazy Constructors

As an additional optimization Koka avoids allocating an indirection node when it sees a constructor. This is exactly what made our earlier `thunk<a>` type work:

```
type thunk<a>
  Memo( v : a )
  lazy Lazy( f : () -> a ) ->
    Memo( f() )
```

Here, no implicit indirection node is created: instead the memory underlying the `Lazy` constructor is rewritten to contain a `Memo` constructor during evaluation.

By inlining the definitions of `sappend` and `sreverse` into the definition of the `stream` data type, we avoid indirection nodes altogether. The type of the `SAppend` constructor then becomes:

```
lazy SAppend( xs : stream<a>, ys : stream<a> ) ->
  match xs
    SCons(x,xx) -> SCons(x, SAppend(xx,ys))
    SNil        -> ys
```

This definition makes explicit that the `SAppend` constructor evaluates to an `SCons` constructor if the first branch is taken. Koka can detect this fact and will not create an indirection node in that case: instead the memory cell holding the `SAppend` constructor is overwritten to contain the `SCons` constructor. This is similar to how the original Spineless Tagless G-machine can sometimes perform in-place updates of closures instead of creating indirections [Peyton Jones 1992].

In those branches where the tail position is not a constructor of an appropriate size, we still generate an indirection node. In the `SNil` case above, we reuse the space of the `SAppend` constructor for an indirection to `ys`. To get rid of them, we can inline the `force` and match on the second stream `ys`:

```
lazy SAppend( xs : stream<a>, ys : stream<a> ) ->
  ...
    SNil -> match ys
      SCons(y,yy) -> SCons(y,yy)
      SNil        -> SNil        // an indirection is still necessary here
```

This rewrites `SAppend` into `SCons` in the first branch and only requires an indirection node in the last case. However, if the old `SCons` in `ys` still has a reference it will stay around, which can increase space usage [Peyton Jones 1992].

4.5 In-place Reuse with Reference Counting

It is important to note that our in-place reuse of *lazy constructors* is quite different to in-place reuse using reference counting [Reinking, Xie et al. 2021; Schulte and Grieskamp 1992; Ullrich and de Moura 2019]. In that setting, memory cells are reused in-place when their reference count is one. In contrast, the memoization of lazy constructors does not require reference count at all. In fact, memoization is only useful if the memory cell is shared among several references!

Nonetheless, the Koka compiler combines both techniques. Internally, the Koka compiler rewrites the `SAppend` constructor to the following code snippet:

```
lazy SAppend( xs : stream<a>, ys : stream<a> ) as _root ->
  match xs
    SCons( x, xx ) as cell ->
      reuse-always(_root, SCons( x, reuse-if-unique(cell, SAppend(xx,ys)) ))
    SNil -> reuse-always(_root, Indirect(ys))
```

That is, the `_root` memory cell holding the `SAppend` constructor is overwritten with the `SCons` cell in the first branch and the `Indirect` node in the second branch. This is independent of the reference count of the `stream` cell. Conversely, if the reference count of the `cell` of the front stream happens to be one (and only then), its memory location is reused for the new `SAppend` constructor.

When a programming languages combines lazy constructors with reference counting, this allows programmers to write code that runs with no fresh allocations at all. This is a key advantage of the first-order approach to laziness. Koka's memory re-use based on reference counts is limited to first-order data constructors: it cannot re-use memory locations associated with closures or traditional thunks. Our approach paves the way for adding laziness to the fully in-place calculus [Lorenzen et al. 2023], which promises to enable the first fully in-place lazy data structures.

4.6 Laziness with and without recursive forcing

As we could observe in the previous sections, lazy constructors differ from the laziness that is typically present in strict language. Forcing the traditional `thunk<a>` type always returns an value of type `a`. While `a` may itself be instantiated to the type `thunk<b>`, the initial `force` operation will not attempt to recursively force it. In contrast, even though a lazy constructor may return another lazy value, the `eval` operation will always attempt to evaluate the result to a strict constructor – quite similar to how lazy languages always force recursively to weak head normal form. This leads us to distinguish two separate designs in our formalization.

Our first design, presented in Section 4.1 corresponds closely to the way laziness is typically implemented in a strict language like OCaml. A rough translation of our terms to OCaml is given in the following table:

Lazy Constructors	OCaml	
$\mathsf{lazy}_F\ v$	`lazy e`	(where $F(v) = e$)
$\mathsf{match}\ x\ \{\ \mathsf{Memo}\ v \rightarrow \ldots\ \}$	`match x with lazy v -> ...`	(match on lazy value)
$\mathsf{step}\ x$	`Lazy.force x`	(one-step evaluation)
memo	`Forward_tag`	(memoized value)
locked	`Forcing_tag`	(lock during evaluation)
indirect	N/A	(indirection to other thunk)

Conversely, our second design, presented in Section 4.3, is more similar to the way laziness is typically implemented in a lazy functional language such as Haskell. In Haskell, forcing a thunk will always return a value in weak head normal form; just like matching on our (recursive) lazy constructors. A rough translation is as follows:

Recursive Lazy Constructors	Haskell	
$\mathsf{lazy}_F\ v$	`e`	(where $F(v) = e$)
$\mathsf{match}\ x\ \{\ Cons \ldots\ \}$	`case x of Cons ...`	(match on whnf)
$\mathsf{eval}\ x;\ e$	`Prelude.seq x e`	(evaluation to whnf)
N/A	`Control.Deepseq.force x`	(evaluate all thunks reachable from `x`)
memo	Indirection	(memoized value)
locked	Black hole	(lock during evaluation)
indirect	Indirection	(indirection to other thunk)

In our terminology, we avoid the term `force` altogether, since it has different meanings in Haskell and OCaml. Furthermore, we use the term indirect only for indirections that lead to another thunk (see Section 5.4) and use memo for indirections that lead directly to the memoized value [4].

[4]While Peyton Jones [1992] only writes indirections to values in whnf (like our memo), Marlow and Peyton Jones [2006] propose to add indirection nodes pointing to unevaluated thunks during GC.

5 Formalization

In this section we formalize a high-level view on lazy constructors that abstracts from implementation concerns such as in-place updates. First, we consider a model of lazy constructors based on Section 4.1, where forcing does not have to recurse. Our type $A \nrightarrow_F B$ states that a lazy constructor carrying A can be forced using the function F to yield a normal constructor of type B and will be memoized. This model of lazy constructors is quite similar to traditional laziness and we can adapt Launchbury [1993]'s semantic to reason about it.

Similar to how normal data types can be encoded as a sum-of-products, we propose to model lazy data types as a *thunked*-sum-of-products. For example, a lazy data type such as:

```
type example
  A(a1 : A1, ..., ai : Ai)
  B(b1 : B1, ..., bj : Bj)
  lazy C(c1 : C1, ..., ck : Ck) -> e_c
  lazy D(d1 : D1, ..., dl : Dl) -> e_d
```

might be encoded as

$$((C_1 \times \ldots \times C_k) + (D_1 \times \ldots \times D_l)) \nrightarrow_F ((A_1 \times \ldots \times A_i) + (B_1 \times \ldots \times B_j))$$

with: $F(lv) = \mathsf{case}\ lv\ \{\ \mathsf{inl}\ (c_1, \ldots, c_k) \rightarrow e_c;\ \mathsf{inr}\ (d_1, \ldots, d_l) \rightarrow e_d\ \}$.

However, this encoding only works if e_c and e_d are guaranteed to return one of the constructors `A` or `B`. As described in Section 4.3, we want to allow lazy constructors to also return further lazy constructors. To model this, we wrap the lazy type into a recursive type:

$$\mu\alpha.\ ((C_1 \times \ldots \times C_k) + (D_1 \times \ldots \times D_l)) \nrightarrow_F (((A_1 \times \ldots \times A_i) + (B_1 \times \ldots \times B_j)) + \alpha)$$

This allows the function F to return either a normal constructor (inl) or another lazy constructor (inr). As we show in this section, we can perform more optimizations if we fuse the recursive type and the lazy constructors into an abstract type of *recursive lazy constructors* $A \nRightarrow_F B := \mu\alpha.\ A \nrightarrow_F (B + \alpha)$. Our final encoding is then:

$$((C_1 \times \ldots \times C_k) + (D_1 \times \ldots \times D_l)) \nRightarrow_F ((A_1 \times \ldots \times A_i) + (B_1 \times \ldots \times B_j))$$

5.1 Core Calculus

Figure 2 shows the syntax and typing rules of a calculus with lazy constructors. The calculus is a standard lambda calculus restricted to be first-order. As such, we include units, sums, products and isorecursive types but not closures. We include top-level function declarations $F(x) = e : A \rightarrow B$. The restriction to first-order is not necessary for the soundness of lazy constructors (and indeed, you can store closures in lazy constructors in Koka), but it emphasises our point that laziness can exist in a purely first-order setting.

To model lazy constructors, we add a new type $A \nrightarrow_F B$. This type represents the lazy computation that arises from applying the top-level function F to an argument of type A, evaluating to a value of type B. We can create a new value of this type by supplying an already computed value $v : B$ as memo v or an input $v : A$ to the computation as $\mathsf{lazy}_F\ v$. While these introduction forms are similar to a sum type, the elimination form step v always returns a value of type B, either by evaluating the computation F or by returning the memoized value. The introduction forms $\mathsf{lazy}_F\ v$ and memo v are not part of the syntax for values, since in the semantics they involve the side-effect of allocating a new location in a store which can persist the result of the lazy evaluation.

5.2 Natural Semantics

In Figure 3, we present a big-step semantics for our calculus. The judgement $\Gamma : e \Downarrow \Delta : v$ means that under store Γ the expression e will evaluate to store Δ and value v. For the standard features of our calculus, the big-step rules are straightforward and they do not modify the store.

Types:

$$A, B \ ::= \ 1 \mid A + B \mid A \times B \mid \alpha \mid \mu\alpha.\, A \mid A \nrightarrow_F B$$

Values and Expressions:

$$\begin{array}{lcll lcll}
v & ::= & x, y, z & \text{(variables)} & e & ::= & v \mid lv & \text{((lazy) values)} \\
 & \mid & () & \text{(unit)} & & \mid & \mathsf{let}\ x = e\ \mathsf{in}\ e & \text{(let binding)} \\
 & \mid & \mathsf{in}_l\ v \mid \mathsf{in}_r\ v & \text{(sum)} & & \mid & \mathsf{case}\ v\ \{\ \mathsf{in}_l\ x \to e;\ \mathsf{in}_r\ y \to e\ \} & \text{(case split)} \\
 & \mid & (v,\ v) & \text{(pair)} & & \mid & \mathsf{split}\ v\ \{\ (x,\ y) \to e\ \} & \text{(splitting pairs)} \\
 & \mid & \mathsf{fold}\ v & \text{(fold rec. type)} & & \mid & \mathsf{unfold}\ v & \text{(unfold rec. type)} \\
lv & ::= & \mathsf{memo}\ v & \textbf{(memoized value)} & & \mid & F\ v & \text{(application)} \\
 & \mid & \mathsf{lazy}_F\ v & \textbf{(lazy computation)} & & \mid & \mathsf{step}\ v & \textbf{(single-step forcing)}
\end{array}$$

$$\Sigma \ ::= \ \varnothing \mid \Sigma, F(x) = e : A \to B \quad \text{(recursive top-level functions)}$$

$$\frac{}{\Gamma, x : A \vdash x : A}\ \textsc{var} \qquad \frac{\Gamma \vdash e_1 : A \quad \Gamma, x : A \vdash e_2 : B}{\Gamma \vdash \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 : B}\ \textsc{let}$$

$$\frac{\Gamma \vdash v : A_i \quad i \in \{l, r\}}{\Gamma \vdash \mathsf{in}_i\ v : A_l + A_r}\ \textsc{inl/inr} \qquad \frac{\Gamma \vdash v : A_l + A_r \quad \Gamma, x : A_i \vdash e_i : C}{\Gamma \vdash \mathsf{case}\ v\ \{\ \mathsf{in}_l\ x \to e_l;\ \mathsf{in}_r\ y \to e_r\ \} : C}\ \textsc{case}$$

$$\frac{\Gamma \vdash v : A \quad \Gamma \vdash w : B}{\Gamma \vdash (v,\ w) : A \times B}\ \textsc{pair} \qquad \frac{\Gamma \vdash v : A \times B \quad \Gamma, x : A, y : B \vdash e : C}{\Gamma \vdash \mathsf{split}\ v\ \{\ (x,\ y) \to e\ \} : C}\ \textsc{split}$$

$$\frac{\Gamma \vdash v : A[\mu\alpha.\, A/\alpha]}{\Gamma \vdash \mathsf{fold}\ v : \mu\alpha.\, A}\ \textsc{fold} \qquad \frac{\Gamma \vdash v : \mu\alpha.\, A}{\Gamma \vdash \mathsf{unfold}\ v : A[\mu\alpha.\, A/\alpha]}\ \textsc{unfold}$$

$$\frac{}{\Gamma \vdash () : 1}\ \textsc{unit} \qquad \frac{F : A \to B \in \Sigma \quad \Gamma \vdash v : A}{\Gamma \vdash F\ v : B}\ \textsc{app}$$

$$\frac{}{\Vdash \varnothing}\ \textsc{defbase} \qquad \frac{\Vdash \Sigma \quad x : A \vdash e : B}{\Vdash \Sigma, F(x) = e : A \to B}\ \textsc{deffun}$$

$$\frac{F : A \to B \in \Sigma \quad \Gamma \vdash v : A}{\Gamma \vdash \mathsf{lazy}_F\ v : A \nrightarrow_F B} \qquad \frac{F : A \to B \in \Sigma \quad \Gamma \vdash v : B}{\Gamma \vdash \mathsf{memo}\ v : A \nrightarrow_F B} \qquad \frac{\Gamma \vdash v : A \nrightarrow_F B}{\Gamma \vdash \mathsf{step}\ v : B}$$

Fig. 2. Syntax and Types for a First-order Calculus with Lazy Constructors

Following the Natural Semantics for Lazy Evaluation [Launchbury 1993], we use a store Γ to keep track of thunks. While Launchbury stores expressions e in the store, we store lazy constructors of the form $\mathsf{lazy}_F\ v$. Assuming that it is known in advance what possible expressions can appear, these representations correspond where F abstracts the expression e as $e = F(v)$ with $v = \mathsf{fv}(e)$. Launchbury stores a fully evaluated expression as a value w, whereas we use the more explicit $\mathsf{memo}\ w$.

The LAZY rule then follows Launchbury's Let-rule, where we create a new thunk in the store Γ and return a new reference to it. Similar to how the Let-rule applies both to expressions and values (since values are a subset of expressions), our LAZY rule applies to all lazy values (both $\mathsf{lazy}_F\ v$ and

$$\frac{}{\Gamma : v \Downarrow \Gamma : v}\ \textsc{value} \qquad \frac{\Gamma : e_1 \Downarrow \Delta : v \quad \Delta : e_2[v/x] \Downarrow \Theta : w}{\Gamma : \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 \Downarrow \Theta : w}\ \textsc{let}$$

$$\frac{F(x) = e \in \Sigma \quad \Gamma : e[v/x] \Downarrow \Delta : w}{\Gamma : F\ v \Downarrow \Delta : w}\ \textsc{app} \qquad \frac{\Gamma : e[v_1/x,\ v_2/y] \Downarrow \Delta : w}{\Gamma : \mathsf{split}\ (v_1,\ v_2)\ \{\ (x,\ y) \to e\ \} \Downarrow \Delta : w}\ \textsc{split}$$

$$\frac{z\ \text{fresh}}{\Gamma : lv \Downarrow (\Gamma,\ z \mapsto lv) : z}\ \textsc{lazy} \qquad \frac{}{\Gamma : \mathsf{unfold}\ (\mathsf{fold}\ v) \Downarrow \Gamma : v}\ \textsc{unfold}$$

$$\frac{z \mapsto \mathsf{memo}\ v \in \Gamma}{\Gamma : \mathsf{step}\ z \Downarrow \Gamma : v}\ \textsc{recall} \qquad \frac{\Gamma : e_i[v/x_i] \Downarrow \Delta : w}{\Gamma : \mathsf{case}\ (\mathsf{in}_i\ v)\ \{\ \mathsf{in}_l\ x_l \to e_l;\ \mathsf{in}_r\ x_r \to e_r\ \} \Downarrow \Delta : w}$$

$$\frac{\Gamma : F\ v \Downarrow \Delta : w}{(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) : \mathsf{step}\ x \Downarrow (\Delta,\ x \mapsto \mathsf{memo}\ w) : w}\ \textsc{step}$$

Fig. 3. Natural Semantics for Lazy Constructors

memo v). The STEP rule follows Launchbury's Variable-rule, where we remove x from the store, evaluate the computation and store the result in the new store. In the Variable-rule, the computed value is further transformed to rename all bound variables, but we can omit this step since our values do not contain lambdas and thus no bound variables. If the thunk happens to be evaluated already, we use the RECALL rule to access it.

5.3 Soundness

We show that our calculus is sound with respect to our semantics using a logical relation. We include a step-indexing parameter k since our calculus includes isorecursive types and write $\Downarrow_k$ for evaluations that can be performed in k steps. A store Δ extends Γ in k steps, written as $\Gamma \sqsubseteq_k \Delta$, if Δ only contains more or more-evaluated thunks than Γ. Concretely, we take the reflexive-transitive closure of the rules:

$$\frac{}{\Gamma \sqsubseteq_1 \Gamma,\ x \mapsto lv}\ \textsc{extend} \qquad \frac{\Gamma : F\ v \Downarrow_k \Delta : w}{\Gamma,\ x \mapsto \mathsf{lazy}_F\ v \sqsubseteq_{k+1} \Delta,\ x \mapsto \mathsf{memo}\ w}\ \textsc{eval}$$

As usual for logical relations we will argue that if an expression can evaluate to a value under a store Γ then it evaluates to the same value for all stores that extend Γ. In this interpretation, the LAZY rule thus encodes the notion of referential transparency: evaluating arbitrary thunks in the heap does not change whether (and to what) an expression evaluates.

Next, we define our meaning of values $\mathcal{V}_k[\![A]\!]$ and expressions $\mathsf{E}_{k,\Delta}[\![A]\!]$. Since we are working with a big-step semantics, we can not distinguish between stuck and diverging programs. However, we will prove that if a well-typed program evaluates to a value, then the value will have the correct type:

$$\mathsf{E}_{k,\,\Delta}[\![A]\!] := \{\ e \mid \forall j < k.\ \forall \Theta,\ v.\ (\Delta : e \Downarrow_j \Theta : v) \Rightarrow \Delta \sqsubseteq_j \Theta \text{ and } (\Theta,\ v) \in \mathcal{V}_{k-j}[\![\ A\ ]\!]\ \}$$

The interpretation of values is straightforward in all cases except for the lazy type. For a lazy value $\mathsf{lazy}_F\ v$, we need both that v is valid for type A and that the thunk can be evaluated to a value of type B at any time in the future:

$$
\begin{aligned}
\mathcal{V}_k[\![\, \mu\alpha.\, A \,]\!] &:= \{\, (\Delta,\ \mathsf{fold}\ v) \mid \forall j < k.\ (\Delta,\ v) \in \mathcal{V}_j[\![\, A[\mu\alpha.\, A/\alpha] \,]\!] \,\} \\
\mathcal{V}_k[\![\, A \nrightarrow_F B \,]\!] &:= \{\, ((\Delta,\ z \mapsto \mathsf{lazy}_F\ v),\ z) \mid (\Delta,\ v) \in \mathcal{V}_k[\![\, A \,]\!],\ \forall j \leqslant k,\ \Theta.\ \Delta \sqsubseteq_j \Theta \Rightarrow \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad F\ v \in \mathsf{E}_{k-j,\Theta}[\![B]\!] \,\} \\
&\ \ \cup \{\, ((\Delta,\ z \mapsto \mathsf{memo}\ v),\ z) \mid (\Delta,\ v) \in \mathcal{V}_k[\![\, B \,]\!] \,\}
\end{aligned}
$$

With this setup, we can prove:

Lemma 1. (*Store extension preserves types.*)
If $(\Delta,\ v) \in \mathcal{V}_k[\![\, A \,]\!]$ and $\Delta \sqsubseteq_j \Theta$, then $(\Theta,\ v) \in \mathcal{V}_{k-j}[\![\, A \,]\!]$.

We connect our semantics to the type system by defining the semantic soundness relation $\Gamma \vDash e : A$ which implies that for all substitutions σ of variables in Γ by values of the correct type that are valid for k more steps, the evaluation of e will yield a value of type A in k steps (if it converges):

$$\Gamma \vDash e : A := \forall k \geqslant 0,\ \Delta,\ \sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!].\ \sigma(e) \in \mathsf{E}_{k,\Delta}[\![A]\!]$$

Then we obtain our type soundness result:

Theorem 1. (*Type Soundness.*)
If $\Gamma \vdash e : A$, then $\Gamma \vDash e : A$.

As an intermediate result from our soundness proof, we also see that evaluating the heap further does not change the final computed value. This shows that the evaluation of lazy constructors in the store in referentially transparent:

Theorem 2. (*Lazy evaluation is referentially transparent.*)
If $\Gamma : e \Downarrow \Delta : v$ and $\Gamma \sqsubseteq \Gamma'$, then $\Gamma' : e \Downarrow \Delta' : v$ with $\Delta \sqsubseteq \Delta'$.

5.4 Recursive Lazy Constructors

To model a type like the `stream`, where `SReverse` can evaluate to another lazy constructor `SReverse`, we need to encode lazy data types using iso-recursive types. In the style of a delay monad [Altenkirch et al. 2017; Capretta 2005; Chapman et al. 2019] or trampoline [Ganz et al. 1999], we can define a type of recursive lazy constructors as:

$$A \twoheadrightarrow_F B := \mu\alpha.\, A \nrightarrow_F (B + \alpha)$$

where the `eval` function recursively steps the lazy constructors until we obtain a non-lazy constructor:

$$\mathsf{eval}\ x\ =\ \mathsf{case}\ (\mathsf{step}\ (\mathsf{unfold}\ x))\ \{\, \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y \,\}$$

This design is only a small extension of our first model of lazy constructors, yet allows us to encode the full power of lazy constructors as in Section 4.3. However, while the definition above is a correct description of recursive lazy constructors, it turns out that our implementation contains another subtlety: it does not allow us to distinguish how many iterations `eval` had to perform.

Consider an inhabitant of this type like $\mathsf{fold}\ (\mathsf{memo}\ (\mathsf{inr}\ (\mathsf{fold}\ (\mathsf{memo}\ (\mathsf{inl}\ v)))))$. At first glance, it might appear that the outer indirection is unnecessary and that the value is in fact equivalent to $\mathsf{fold}\ (\mathsf{memo}\ (\mathsf{inl}\ v))$, since both values yield v when passed to eval. However, in the definition above, there is no restriction that lazy constructors are always deconstructed using eval, which makes it possible to distinguish the two values by simply case-splitting after a call to step. This is a well-known problem for implementations of laziness that attempt to short-cut indirections. For example, consider the following OCaml code:

```
let nested = lazy (lazy (raise Not_found))
let eval l = match l with lazy v -> ()
let () = eval nested; eval nested; ()
```

In this code, we evaluate the outer lazy twice and expect the exception *not* to be thrown. However, after the first evaluation, the lazy is rewritten into an indirection. If the runtime system attempted to short-cut the indirection, this would make `nested` point to the inner lazy value. Then the second evaluation would throw the exception, thus changing the semantics of the program. In practice,

OCaml's runtime system still tries to short-cut indirections, but guarantees that this does not affect the semantics of the program; in particular, indirections are not removed if they lead to an unevaluated lazy value [5].

However, our example is a bit different from the OCaml example: while the two lazy values in OCaml belong to different thunks, in our case the two indirections morally belong to the same thunk. This suggests that if we were to make our encoding abstract, we could use this fact to short-cut the indirection.

5.5 Short-cutting Indirections

To be able to short-cut indirections, we thus need to ensure that the encoding of recursive lazy constructors stays abstract and make it a first-class type in our calculus. We introduce a new type $A \twoheadrightarrow\!\!\!\!\!\!+_F B$ and define its interpretation by fusing the $A \nrightarrow_F B$ type with the iso-recursive type wrapped around it:

$$\begin{aligned}\mathcal{V}_k[\![A \twoheadrightarrow_F B]\!] := \{ ((\Delta, z \mapsto \mathsf{lazy}_F\ v), z) \mid (\Delta, v) \in \mathcal{V}_k[\![A]\!], \forall j < k, \Theta.\, \Delta \sqsubseteq \Theta \Rightarrow \\ F\ v \in \mathsf{E}_{k-j,\Theta}[\![B + (A \twoheadrightarrow_F B)]\!] \} \\ \cup \{ ((\Delta, z \mapsto \mathsf{memo}\ v), z) \mid (\Delta, v) \in \mathcal{V}_k[\![B]\!] \} \\ \cup \{ ((\Delta, z \mapsto \mathsf{indirect}\ v), z) \mid \forall j < k.\, (\Delta, v) \in \mathcal{V}_j[\![A \twoheadrightarrow_F B]\!] \}\end{aligned}$$

Compared to earlier, a lazy value now evaluates to $B + (A \twoheadrightarrow_F B)$, but memo still only contains values of type B. Instead, if the lazy value evaluates to $A \twoheadrightarrow_F B$, then we create an indirection node pointing to the folded lazy constructor. The crucial aspect of this formalization is that we can now shortcut indirections. We formalize this using the following two rules:

$$\frac{y \mapsto \mathsf{indirect}\ z \in \Gamma}{\Gamma, x \mapsto \mathsf{indirect}\ y \sqsubseteq \Gamma, x \mapsto \mathsf{indirect}\ z}\ \textsc{cutI} \qquad \frac{y \mapsto \mathsf{memo}\ v \in \Gamma}{\Gamma, x \mapsto \mathsf{indirect}\ y \sqsubseteq \Gamma, x \mapsto \mathsf{memo}\ v}\ \textsc{cutM}$$

We do not allow shortcutting an indirection that points directly to a lazy constructor since that could duplicate work: if the lazy constructor would be duplicated into a different location in the store, its evaluation would be independent of the evaluation of the original location. Our referential transparency theorem still holds for this calculus:

Theorem 3. (*Short-cutting indirections is referentially transparent.*)
If $\Gamma : e \Downarrow \Delta : v$ and $\Gamma \sqsubseteq \Gamma'$, then $\Gamma' : e \Downarrow \Delta' : v$ with $\Delta \sqsubseteq \Delta'$.

5.6 Laziness via Quotient Types

As a helpful perspective on the result of this section, one might view our encoding as an application of quotient types. Our initial $A \nrightarrow_F B$ type can be viewed as a quotient type $[\![A + B]\!]_\sim$ which identifies $\mathsf{inl}\ \alpha \sim \mathsf{inr}\ b$ iff $F(\alpha) = b$. Similarly, the $A \twoheadrightarrow_F B$ type can be viewed as a quotient type $[\![A + B]\!]_\sim$ which identifies $\mathsf{inl}\ \alpha \sim c$ iff $F^n(\alpha) = c$ for some n. Since quotient types abstract from the concrete representative and only allow the programmer to access the value up to the equivalence relation, it is possible to replace the concrete representative at any time [Selsam et al. 2020]. This makes it sound to memoize lazy constructors and short-cut indirections.

However, the way we would use the quotient type as part of a larger type is perhaps different from the traditional use of quotient types. For example, you could define a stream type using lazy constructors as:

$\mu s.\ (\mathsf{SAppend}(s_1 : s,\ s_2 : s)) \nrightarrow_F (\mathsf{SCons}(x : \alpha,\ \overline{x} : s) + \mathsf{SNil})$
with $F(\mathsf{SAppend}(s_1, s_2)) = \ldots$

This encoding *interleaves* the lazy constructor with the iso-recursive structure s of streams. In

[5] https://github.com/ocaml/ocaml/blob/4.14/runtime/minor_gc.c#L236

contrast, consider what the encoding of streams would look like if we moved the quotient type outside of the recursive type:

$[\![\mu s.\ (\mathsf{SAppend}(s_1 : s,\ s_2 : s)) + (\mathsf{SCons}(x : \alpha,\ \overline{x} : s) + \mathsf{SNil})]\!]_\sim$
which identifies $\mathsf{SAppend}(s_1,\ s_2) \sim b$ iff $b = F(\mathsf{SAppend}(s_1,\ s_2))$.

This model would allow us to identify elements of the type "deeply" where you identify constructors depending on their children. For example, we might design the $\sim$ relation to identify $\mathsf{SAppend}(\mathsf{SAppend}(s_1,\ s_2),\ s_3) \sim \mathsf{SAppend}(s_1,\ \mathsf{SAppend}(s_2,\ s_3))$. This can be very powerful (see Section 3.4), but also makes it possible to break referential transparency.

In contrast, our encoding is "shallow" in that it ensures that the lazy constructor α can only be identified with its evaluation $F(\alpha)$. While the evaluation function F can match on the children of the lazy constructor, it can not inspect the representative of the quotient. For example, that implies that $F(\mathsf{SAppend}(s_1,\ s_2)) \sim F(\mathsf{SAppend}(s_1',\ s_2'))$ for all $s_1 \sim s_1'$, $s_2 \sim s_2'$. This property ensures the referential transparency of lazy evaluation.

6 Implementation

While Section 5 gives a high-level overview over lazy constructors, it does not give a direct strategy for implementing lazy constructors efficiently. In this section, we instead take a more low-level view. We propose several primitives that can be used to implement lazy constructors efficiently and derive an efficient recursive evaluation algorithm. Unlike the high-level step function of the previous section, our new low-level primitives do not preserve the referential transparency of lazy evaluation and should thus be exposed only as unsafe or kept hidden in the underbelly of a compiler.

In Section 2.3, we define a simple `stream/eval` function, but already noticed that it used too much stack space. In this section, we will show how to derive a more efficient implementation of `stream/eval`:

```
fun stream/eval( s : stream<a> )
  lazy match s
    SAppend( s1, s2 ) -> lazy-eval-sappend(s, s1, s2)
    SReverse( s1, acc ) -> lazy-eval-sreverse(s, s1, acc)
    Indirect(ind) -> eval(ind)
    _ -> s
fun lazy-eval-sappend( s : stream<a>, s1 : stream<a>, s2 : stream<a> )
  match s1
    SCons(x,xx) -> lazy-update(s, SCons(x, SAppend(xx, s2)))
    SNil        -> lazy-update(s, Indirect(s2)); eval(s2)
fun lazy-eval-sreverse( s : stream<a>, s1 : stream<a>, acc : stream<a> )
  match s1
    SCons(x,xx) -> lazy-eval-sreverse(s, xx, SCons(x, acc))
    SNil        -> lazy-update(s, Indirect(s2)); eval(s2)
```

Compared to our simpler implementation, we can see that we now write an indirection node if a lazy constructor returns another lazy constructor. This allows us to keep evaluating without using stack space, but means that we might have to follow an indirection chain from a previous evaluation. Furthermore, our simpler version generated an `eval(SReverse(xx, SCons(x,acc)))`, where a lazy constructor is created and immediatedly evaluated. In our new implementation, we instead directly jump to the correct evaluation function `lazy-eval-sreverse` which saves a branch and writes to memory.

Unfortunately, it is tricky to show that our final implementation is in fact correct, where a lazy constructor will be updated to the correct value. However, it turns out that we can derive this version by equational reasoning from the simpler implementation, which corresponds more clearly to our high-level calculus.

Expressions:

$$\begin{array}{rcll} e & ::= & \dots & \text{(as before)} \\ & | & \mathsf{lazy\ match}\ v\ \{\ \mathsf{lazy}_F\ l\ y \to e;\ \mathsf{memo}\ y \to e\ \} & \text{(lazy match and acquire lock)} \\ & | & \mathsf{memoize}\ l\ w & \text{(update cell and release lock)} \end{array}$$

$$\Sigma \ ::= \ \dots \mid \Sigma, F(l; x) \ = \ e : A \to B \to C \quad \text{(top-level functions)}$$

We keep all rules as in the high-level core calculus except STEP. We augment each rule with a linear environment L of locations. For the introduction rules, UNFOLD rule and APP rule L is empty, the LET rule splits L among its antecedents while the CASE and SPLIT rules pass L to their antecedents. Furthermore we add the rules:

$$\frac{\Vdash \Sigma \quad l : A \mid x : B \vdash e : C}{\Vdash \Sigma, F(l; x) \ = \ e : A \to B \to C}\ \text{DEFLAPP} \qquad \frac{F : A \to B \to C \in \Sigma \quad \varnothing \mid \Gamma \vdash v : B}{l : A \mid \Gamma \vdash F\ l\ v : C}\ \text{LAPP}$$

$$\frac{\varnothing \mid \Gamma \vdash w : B}{l : A \ \not\Rightarrow_F \ B \mid \Gamma \vdash \mathsf{memoize}\ l\ w : B}\ \text{MEMOIZE}$$

$$\frac{\varnothing \mid \Gamma \vdash v : A \ \not\Rightarrow_F \ B \quad L, l : A \ \not\Rightarrow_F \ B \mid x : A \vdash e_1 : C \quad L \mid \Gamma,\ y : B \vdash e_2 : C}{L \mid \Gamma \vdash \mathsf{lazy\ match}\ v\ \{\ \mathsf{lazy}_F\ l\ x \to e_1;\ \mathsf{memo}\ y \to e_2\ \} : C}\ \text{LAZYMATCH}$$

Fig. 4. Low-level core calculus

6.1 Implementation calculus

Our low-level calculus is a variation of the high-level calculus. We introduce two new primitives: $\mathsf{lazy\ match}\ v\ \{\ \mathsf{lazy}_F\ l\ x \to e_1;\ \mathsf{memo}\ y \to e_2\ \}$ allows us to inspect the value of a lazy constructor. In the first branch e_1 we additionally get access to the location l of the lazy constructor. Our second primitive $\mathsf{memoize}\ l\ w$ allows us to overwrite the cell l of a lazy constructor with memo w. A location $l : : A \ \not\Rightarrow_F \ B$ acts as a destination for a value of type B [Allain et al. 2025; Bagrel and Spiwack 2025; Shaikhha et al. 2017], which can filled using memoize. In our semantics, all locations l returned by lazy match are *locked* and thus can not be accessed until the lock is released by memoize.

To ensure the soundness of the low-level calculus, we need to ensure that locked locations are handled linearly. In particular, this guarantees that a locked location is overwritten using memoize exactly once. We achieve this by adding a second environment L to the calculus that contains all locked locations and ensure that no locked location can ever escape into a value held in Γ. The rules of our high-level calculus can be modified to treat the L environment linearly in the usual way.

In further preparation, we add a new type of top-level function $F(l; x)$ which also takes a location. This is necessary, since locked locations may not be stored in values and so we can not represent this by a product $F((l, x))$. Despite taking two arguments a function $F(l; x)$ has to be fully applied.

Given those primitive operations, we can implement the step operation as:

$$\begin{array}{ll} \mathsf{step}\ x \ = \mathsf{lazy\ match}\ x & \\ \qquad\qquad \mathsf{lazy}_F\ l\ v \to \mathsf{memoize}\ l\ (F\ v) & \\ \qquad\qquad \mathsf{memo}\ y \to y & \end{array}$$

It is easy to see that with this implementation of step, the original STEP rule of the high-level calculus becomes derivable. In particular, our low-level calculus strictly extends the high-level calculus:

Lemma 2. (*The low-level calculus implements the high-level calculus*)
If $\Gamma \vdash e : A$, then $\varnothing \mid \Gamma \vdash e : A$.

Store and evaluation context:

$v ::= a \mid \ldots$ (heap cells)
$\varphi ::= \mathsf{memo}\ v \mid \mathsf{lazy}_F\ v \mid \mathsf{locked}$
$\mathsf{S} ::= \varnothing \mid \mathsf{S}, a \mapsto \varphi$
$\mathsf{E} ::= \square \mid \mathsf{let}\ x = \mathsf{E}\ \mathsf{in}\ e$

$$\frac{\mathsf{S} \mid e_1 \longrightarrow \mathsf{S}' \mid e_2}{\mathsf{S} \mid \mathsf{E}[e_1] \longmapsto \mathsf{S}' \mid \mathsf{E}[e_2]}\ \textsc{step}$$

Evaluation steps:

$$\begin{array}{llll}
(\mathit{let}) & \mathsf{S} \mid \mathsf{let}\ x = v\ \mathsf{in}\ e & \longrightarrow & \mathsf{S} \mid e[v/x] \\
(\mathit{app}) & \mathsf{S} \mid F\ v & \longrightarrow & \mathsf{S} \mid e[v/x] \quad \text{where } F(x) = e \in \Sigma \\
(\mathit{split}) & \mathsf{S} \mid \mathsf{split}\ (v_1,\ v_2)\ \{\ (x,\ y) \to e\ \} & \longrightarrow & \mathsf{S} \mid e[v_1/x, v_2/y] \\
(\mathit{case}) & \mathsf{S} \mid \mathsf{case}\ (\mathsf{in}_i\ v)\ \{\ \mathsf{in}_l\ x_l \to e_l;\ \mathsf{in}_r\ x_r \to e_r\ \} & \longrightarrow & \mathsf{S} \mid e_i[v/x_i] \\
(\mathit{unfold}) & \mathsf{S} \mid \mathsf{unfold}\ (\mathsf{fold}\ v) & \longrightarrow & \mathsf{S} \mid v \\
(\mathit{lazy}) & \mathsf{S} \mid \mathsf{lazy}_F\ v & \longrightarrow & \mathsf{S},\ a \mapsto \mathsf{lazy}_F\ v \mid a \qquad a \text{ fresh} \\
(\mathit{memo}) & \mathsf{S} \mid \mathsf{memo}\ v & \longrightarrow & \mathsf{S},\ a \mapsto \mathsf{memo}\ v \mid a \qquad a \text{ fresh} \\
(\mathit{memoize}) & \mathsf{S},\ a \mapsto \mathsf{locked} \mid \mathsf{memoize}\ a\ w & \longrightarrow & \mathsf{S},\ a \mapsto \mathsf{memo}\ w \mid w
\end{array}$$

$$\begin{array}{lll}
(\mathit{lazy\ match}) & \mathsf{S},\ a \mapsto v \mid \mathsf{lazy\ match}\ a\ \{\ \mathsf{lazy}_F\ l\ x \to e_1;\ \mathsf{memo}\ y \to e_2\ \} & \\
\longrightarrow & \mathsf{S},\ a \mapsto \mathsf{locked} \mid e_1[a/l, w/x] & \text{if } v = \mathsf{lazy}_F\ w \\
\longrightarrow & \mathsf{S},\ a \mapsto \mathsf{memo}\ w \mid e_2[w/y] & \text{if } v = \mathsf{memo}\ w
\end{array}$$

Fig. 5. Small-step semantics of implementation

6.2 Small-step semantics

In Figure 5, we describe a small-step semantics for lazy constructors. As in the natural semantics, we only keep lazy constructors in the store. Each memory cell is either a lazy or memo value or locked. The small-semantics of the derived step function corresponds directly to the natural semantics of the primitive step: while the big-step semantics removes the cell x from the heap entirely during evaluation, the small-step semantics keeps it in the heap as locked and thus inaccessible. This allows us to prove that the small-step semantics faithfully implements the big-step semantics:

Lemma 3. (*The small-step semantics implements high-level semantics.*)
If $\Gamma : e \Downarrow \Delta : v$, then $\Gamma \mid e \longmapsto^* \Delta \mid v$.

Would it also be possible to show the reverse direction? In general this is not possible, since the implementation calculus allows us to memoize *any* value of the correct type, while the high-level calculus only allows us to memoize values that are produced by the evaluation of a lazy constructor. However, for expressions that can be checked in the high-level calculus, the small-step semantics and the big-step semantics are equivalent.

Our description of cells under evaluation as "locked" mirrors our implementation: If Koka detects that a lazy constructor is thread-shared (recorded using reference counts [Ullrich and de Moura 2019]), Koka will use an atomic compare-and-swap to overwrite the tag of the lazy constructor with a special tag that indicates that the cell is being evaluated. The lazy match primitive acquires this lock during matching and the memoize primitive releases it. This ensures that a cell is evaluated at most once, even in the presence of multiple threads. However, as in other implementations of laziness, the evaluation of a lazy constructor may deadlock if it tries to evaluate itself. Haskell and OCaml can detect this case and throw an exception.

6.3 Tail-recursive Evaluation

Using our new primitives, we can obtain a faster implementation of the lazy evaluation function. As before, we can define the recursive forcing function as:

$$
\begin{aligned}
\mathsf{eval}\ x &= \mathsf{case}\ (\mathsf{step}\ (\mathsf{unfold}\ x))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \} \\
&= \mathsf{lazy\ match}\ (\mathsf{unfold}\ x) \\
&\qquad \mathsf{lazy}_F\ l\ v \to \mathsf{case}\ (\mathsf{memoize}\ l\ (F\ v))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \} \\
&\qquad \mathsf{memo}\ y \to \mathsf{case}\ y\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}
\end{aligned}
$$

In the first branch of the lazy-match we run $F\ v$, memoize its result and then case-split to find out if the result can be returned (inl) or is a lazy constructor that needs to be evaluated further.

It turns out that we can often avoid the case-split if we specialize the eval function to the concrete function F that is evaluated in the lazy constructor. To achieve this we will use equational reasoning in the style of [Leijen and Lorenzen 2023 2025]. First, we define the translation $[\![e]\!]_l$ as:

$$[\![e]\!]_l = \mathsf{case}\ (\mathsf{memoize}\ l\ e)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}$$

For the function F that is evaluated in eval, we define a specialized function F' with the definition:

$$F'(l;\ v) = [\![F\ v]\!]_l$$

and we can use it in our eval function as:

$$
\begin{aligned}
\mathsf{eval}\ x &= \mathsf{lazy\ match}\ (\mathsf{unfold}\ x) \\
&\qquad \mathsf{lazy}_F\ l\ v \to F'\ l\ v \\
&\qquad \mathsf{memo}\ v \to \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}
\end{aligned}
$$

So far, nothing has happened: our new evaluation function corresponds exactly to the previous version. However, we have shifted the position of the case-split from the eval function into the translation. The key insight is now that we can improve the translation $[\![e]\!]_l$ by specializing it to different syntactic constructs. For example, several syntactic constructs permute with both memoize l and case:

$$
\begin{aligned}
[\![\mathsf{let}\ y = e_1\ \mathsf{in}\ e_2]\!]_l &= \mathsf{let}\ y = e_1\ \mathsf{in}\ [\![e_2]\!]_l \\
[\![\mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to e_1;\ \mathsf{inr}\ y \to e_2\ \}]\!]_l &= \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to [\![e_1]\!]_l;\ \mathsf{inr}\ y \to [\![e_2]\!]_l\ \} \\
[\![\mathsf{split}\ v\ \{\ (y,\ z) \to e\ \}]\!]_l &= \mathsf{split}\ v\ \{\ (y,\ z) \to [\![e]\!]_l\ \}
\end{aligned}
$$

This means that we can push down the memoization and case-split into the return values of the computation F. If the translated expression e is well-typed, there are few possible return values to F. For variables (and unfolds of variables), we have to perform the memoization and case-split. But in some cases we can do better:

$$
\begin{aligned}
[\![\mathsf{inl}\ w]\!]_l &= \mathsf{memoize}\ l\ (\mathsf{inl}\ w);\ w \\
[\![\mathsf{inr}\ w]\!]_l &= \mathsf{memoize}\ l\ (\mathsf{inr}\ w);\ \mathsf{eval}\ w
\end{aligned}
$$

If the function F ends in inl w, the evaluation ends at this point. We can thus memoize this result and return w without an extra case-split. If the function F ends in inr w, we also memoize the intermediate result and directly continue evaluating it.

In our implementation, these two cases enable the in-place update of lazy constructors: Since we know the size of the data that is memoized, we can often avoid creating an indirection node memo and instead write the data directly into the lazy cell.

Another interesting special case is if we can already see syntactically what the next lazy thunk will be. This happens for example when a `SReverse` constructor is evaluated to another `SReverse`. In that case, we can specialize the call to eval further:

$$
\begin{aligned}
[\![\mathsf{inr}\ (\mathsf{fold}\ (\mathsf{lazy}_F\ v))]\!]_l &= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ (\mathsf{lazy}_F\ v))))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \} \\
&= \mathsf{let}\ y = \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \mathsf{eval}\ (\mathsf{fold}\ y) \\
&= \mathsf{let}\ y = \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \mathsf{lazy\ match}\ y \\
&\qquad \mathsf{lazy}_F\ l\ v \to F'\ l\ v \\
&\qquad \mathsf{memo}\ v \to \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \} \\
&= \mathsf{let}\ y = \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \mathsf{lock}\ y\ \mathsf{in}\ F'\ y\ v
\end{aligned}
$$

In the last line, we now create a new lazy cell, memoize it in l and then immediately jump to F'.

$$\frac{\varnothing \mid \Gamma \vdash w : B}{l : A \twoheadrightarrow_F B \mid \Gamma \vdash \mathsf{memoize}\ l\ w : B} \qquad \frac{\varnothing \mid \Gamma \vdash w : A \twoheadrightarrow_F B}{l : A \twoheadrightarrow_F B \mid \Gamma \vdash \mathsf{indirect}\ l\ w : A \twoheadrightarrow_F B}$$

$$\frac{\begin{array}{cc}\varnothing \mid \Gamma \vdash v : A \twoheadrightarrow_F B & L, l : A \twoheadrightarrow_F B \mid x : A \vdash e_1 : C \\ L \mid \Gamma,\ y : A \twoheadrightarrow_F B \vdash e_2 : C & L \mid \Gamma,\ z : B \vdash e_3 : C\end{array}}{L \mid \Gamma \vdash \mathsf{lazy\ match}\ v\ \{\ \mathsf{lazy}_F\ l\ x \to e_1;\ \mathsf{indirect}\ y \to e_2;\ \mathsf{memo}\ z \to e_3\ \} : C}\ \textsc{lazymatch}$$

Evaluation steps:

$$\begin{array}{lll}
(\textit{memoize}) & \mathrm{S},\ a \mapsto \mathsf{locked} \mid \mathsf{memoize}\ a\ w & \longrightarrow\ \mathrm{S},\ a \mapsto \mathsf{memo}\ w \mid w \\
(\textit{indirect}) & \mathrm{S},\ a \mapsto \mathsf{locked} \mid \mathsf{indirect}\ a\ w & \longrightarrow\ \mathrm{S},\ a \mapsto \mathsf{indirect}\ w \mid w
\end{array}$$

$$\begin{array}{lll}
(\textit{lazy match}) & \mathrm{S},\ a \mapsto v \mid \mathsf{lazy\ match}\ a\ \{\ \mathsf{lazy}_F\ l\ x \to e_1;\ \mathsf{indirect}\ y \to e_2;\ \mathsf{memo}\ z \to e_3\ \} & \\
\longrightarrow & \mathrm{S},\ a \mapsto \mathsf{locked} \mid e_1[a/l, w/x] & \text{if } v = \mathsf{lazy}_F\ w \\
\longrightarrow & \mathrm{S},\ a \mapsto \mathsf{indirect}\ y \mid e_2[w/y] & \text{if } v = \mathsf{indirect}\ w \\
\longrightarrow & \mathrm{S},\ a \mapsto \mathsf{memo}\ w \mid e_3[w/z] & \text{if } v = \mathsf{memo}\ w
\end{array}$$

Fig. 6. Short-cutting indirections

Since F' expects y to be locked, we use the macro:

$\mathsf{lock}\ y\ \mathsf{in}\ e := \mathsf{lazy\ match}\ y\ \{\ \mathsf{lazy}_F\ y\ \_ \to e;\ \mathsf{memo}\ v \to \mathsf{impossible}\ \}$

However, this might seem slightly wasteful: Why do we write the result into a new cell y and write an indirection into l when we could just write the result into l? This is not quite possible so far, since we do not consider indirections specially and thus run into the problem of Section 5.4. However, by considering indirections as proposed in Section 5.5, we can use an additional reasoning step to reduce the last line to just $F'\ l\ v$.

6.4 Short-cuts during evaluation

To be able to short-cut indirections, we change our memoize and lazy match primitives and add a new indirect primitive as in Figure 6. The indirect instruction acts like memoize but puts an indirect into the store that can be safely shortcut as shown in Section 5.5. On the locations l we now have to record the full type of the lazy constructor, where a location of type of type $A \twoheadrightarrow_F B$ can be filled either with an indirection to another value of type $A \twoheadrightarrow_F B$ or with the memoized result B. Given those primitive operations, we can refine the eval operation from Section 6.3 as:

$$\begin{array}{ll}
\mathsf{eval}\ x\ = \mathsf{lazy\ match}\ x & \\
& \mathsf{lazy}_F\ l\ v \to \mathsf{case}\ (F\ v)\ \{\ \mathsf{inl}\ y \to \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \to \mathsf{eval}\ (\mathsf{indirect}\ l\ y)\ \} \\
& \mathsf{indirect}\ y \to \mathsf{eval}\ y \\
& \mathsf{memo}\ y \to y
\end{array}$$

We can now repeat the calculation from the previous section. The cases for let-bindings, case-statements and split-statements are the same and for inl w and inr w we obtain similar terms. The main difference is in the special inr ($\mathsf{lazy}_F\ v$) case. After calling $F'\ y\ v$, y points to a chain of indirections ending in a memo. The CUTM and CUTI short-cutting rules give us the laws:

$$\begin{array}{l}
\mathsf{indirect}\ l\ y;\ \mathsf{memoize}\ y\ z\ =\ \mathsf{memoize}\ l\ z;\ \mathsf{memoize}\ y\ z \\
\mathsf{indirect}\ l\ y;\ \mathsf{indirect}\ y\ z\ =\ \mathsf{indirect}\ l\ z;\ \mathsf{indirect}\ y\ z
\end{array}$$

This allows us to replace the evaluation of $F'\ y\ v$ by $F'\ l\ v$:

Expressions:

$$
\begin{array}{llll}
e & ::= & \ldots & \text{(as before)} \\
 & \mid & \mathsf{link}\ l\ (l',\ v)\ \mathsf{in}\ e & \text{(link cells and keep lock)} \\
 & \mid & \mathsf{unlink}\ \{\ \mathsf{inl}\ () \to e;\ \mathsf{inr}\ (l,\ x) \to e\ \} & \text{(unlink locked cells)}
\end{array}
$$

$$\varphi ::= \mathsf{memo}\ v\ \mid \mathsf{lazy}_F\ v\ \mid \mathsf{locked}\ \mid \mathsf{locked}\ (z,\ v)$$

$$
\begin{array}{rlll}
(\mathit{link}) & \mathrm{S},\ a \mapsto \mathsf{locked} \mid \mathsf{link}\ a\ (a',\ v)\ \mathsf{in}\ e & \longrightarrow \mathrm{S},\ a \mapsto \mathsf{locked}\ (a',\ v) \mid e & \\
(\mathit{unlink}) & \mathrm{S},\ a \mapsto v \mid \mathsf{unlink}\ a\ \{\ \mathsf{inl}\ () \to e_1;\ \mathsf{inr}\ (l,\ x) \to e_2\ \} & & \\
\longrightarrow & \mathrm{S},\ a \mapsto \mathsf{locked} \mid e_1 & \text{if } v\ =\ \mathsf{locked} & \\
\longrightarrow & \mathrm{S},\ a \mapsto \mathsf{locked} \mid e_2[a'/l, w/x] & \text{if } v\ =\ \mathsf{locked}\ (a',\ w) &
\end{array}
$$

$$\frac{\varnothing \mid \Gamma \vdash v : A \qquad L,\ l : \Box_A\ B \mid \Gamma \vdash e : C}{L,\ l : B,\ l' : \Box_A\ B \mid \Gamma \vdash \mathsf{link}\ l\ (l',\ v)\ \mathsf{in}\ e : C}\ \textsc{link}$$

$$\frac{L \mid \Gamma \vdash e_1 : C \qquad L,\ l : B,\ l' : \Box_A\ B \mid \Gamma,\ x : A \vdash e_2 : C}{L,\ l : \Box_A\ B \mid \Gamma \vdash \mathsf{unlink}\ l\ \{\ \mathsf{inl}\ () \to e_1;\ \mathsf{inr}\ (l',\ x) \to e_2\ \} : C}\ \textsc{unlink}$$

Fig. 7. Linking cells

$$
\begin{array}{lll}
[\![\mathsf{inl}\ w]\!]_l & = & \mathsf{memoize}\ l\ w \\
[\![\mathsf{inr}\ w]\!]_l & = & \mathsf{indirect}\ l\ w;\ \mathsf{eval}\ w \\
[\![\mathsf{inr}\ (\mathsf{lazy}_F\ v)]\!]_l & = & \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{indirect}\ l\ y;\ \mathsf{lock}\ y\ \mathsf{in}\ F'\ y\ v \\
 & = & \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{lock}\ y\ \mathsf{in}\ \mathsf{indirect}\ y\ l;\ F'\ l\ v \\
 & = & F'\ l\ v
\end{array}
$$

In the last case, we are creating an indirection in y that points to l. But then the cell y is unused and we can avoid allocating it altogether.

In practice, we also short-cut indirections in the indirect case of the lazy match construct. This important in practice to avoid long chains of indirections, which can change the time complexity if traversed repeatedly. We discuss in more detail in our tech report [Lorenzen et al. 2025].

6.5 Linking cells for Schorr-Waite traversal

Leading up to our discussion of the Schorr-Waite traversal, we make another addition to our calculus. While we have so far represented locked cells by the value locked, our implementation actually only sets a flag in the header to indicate that they are locked. This means that the storage space of the cell remains available even while it is in a locked state. In our calculus, we use a new value locked $(l,\ v)$ to indicate that a cell is in a locked state but contains both another location l and a value v.

We can write to a locked cell l using the link $l\ (l',\ v)$ in e construct and we can check whether a locked cell l contains a value using the unlink l construct. Cells that may contain a value have type $\Box_A\ B$, which means they contain a linked value of type A and are a destination for type B. To create a linked chain in the first place, we assume that there is a cell null $:\ \Box_A\ B$ in the environment (corresponding to a `NULL` pointer in the implementation).

6.6 Schorr-Waite Evaluation of Lazy Constructors

As we saw in Section 2.5, the evaluation of thunks can lead to stack overflow, if there are recursive calls to eval in F. We can avoid this by further transforming the lazy evaluation function as $[\![e]\!]_{z,l}$ with an additional zipper z that is stored in the lazy cells themselves.

For the i-th call to eval in F, we write E_i to denote its evaluation-context. Let A_i be the product type of the free variables of E_i. We define the zipper as the sum type of the A_i and define an unroll

function that for a given zipper extracts the correct evaluation context to continue:

$$\begin{array}{l} \mathsf{unroll}\ l\ v\ = \mathsf{unlink}\ l \\ \qquad\qquad\qquad \mathsf{inr}\ (z,\ \alpha) \to \mathsf{case}\ \alpha\ \{\ \overline{\mathsf{in}_i\ \alpha_i \to [\![\mathsf{E}_i[v]]\!]_{z,l}}\ \} \\ \qquad\qquad\qquad \mathsf{inl}\ () \to v\ \} \end{array}$$

Then we can extend our translation to find such evaluation contexts, construct the zipper and call the correct unroll function:

$$[\![\mathsf{E}_i[\mathsf{eval}\ v]]\!]_{z,l}\ =\ \mathsf{link}\ l\ (z,\ \mathsf{in}_i\ \alpha_i)\ \mathsf{in}\ \mathsf{unroll}\ l\ (\mathsf{eval}\ v)\ \mathsf{where}\ \alpha_i\ =\ \mathsf{fv}(\mathsf{E}_i)$$

Again, we see that nothing has changed: the call to unroll can be inlined to yield the left-hand-side, since the unlink in unroll just extract what we just linked into l. But this setup now provides us with a technique for making the call to eval tail-recursive. We define a new forcing function that includes the unroll:

$$\begin{array}{rcl} \mathsf{eval}'\ z\ x &=& \mathsf{unroll}\ z\ (\mathsf{eval}\ x) \\ &=& \mathsf{unroll}\ z\ (\mathsf{case}\ (\mathsf{step}\ (\mathsf{unfold}\ x))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}) \\ &=& \mathsf{case}\ (\mathsf{step}\ (\mathsf{unfold}\ x))\ \{\ \mathsf{inl}\ y \to \mathsf{unroll}\ z\ y;\ \mathsf{inr}\ y \to \mathsf{unroll}\ z\ (\mathsf{eval}\ y)\ \} \\ &=& \mathsf{case}\ (\mathsf{step}\ (\mathsf{unfold}\ x))\ \{\ \mathsf{inl}\ y \to \mathsf{unroll}\ z\ y;\ \mathsf{inr}\ y \to \mathsf{eval}'\ z\ y\ \} \\ &=& \mathsf{lazy\ match}\ (\mathsf{unfold}\ x) \\ && \quad \mathsf{lazy}_F\ l\ v \to \mathsf{case}\ (\mathsf{memoize}\ l\ (F\ v)\ )\ \{\ \mathsf{inl}\ y \to \mathsf{unroll}\ z\ y;\ \mathsf{inr}\ y \to \mathsf{eval}'\ z\ y\ \} \\ && \quad \mathsf{memo}\ y \to \mathsf{case}\ y\ \{\ \mathsf{inl}\ y \to \mathsf{unroll}\ z\ y;\ \mathsf{inr}\ y \to \mathsf{eval}'\ z\ y\ \} \end{array}$$

This suggests that we change our interpretation function to:

$$[\![e]\!]_{z,l}\ = \mathsf{case}\ (\mathsf{memoize}\ l\ e\ )\ \{\ \mathsf{inl}\ y \to \mathsf{unroll}\ z\ y;\ \mathsf{inr}\ y \to \mathsf{eval}'\ z\ y\ \}$$

which yields our tail-recursive forcing function:

$$\begin{array}{rcl} \mathsf{eval}'\ z\ x &=& \mathsf{lazy\ match}\ (\mathsf{unfold}\ x) \\ && \quad \mathsf{lazy}_F\ l\ v \to F'\ z\ l\ v \\ && \quad \mathsf{memo}\ y \to \mathsf{case}\ y\ \{\ \mathsf{inl}\ y \to \mathsf{unroll}\ z\ y;\ \mathsf{inr}\ y \to \mathsf{eval}'\ z\ y\ \} \\ \\ \mathsf{eval}\ x &=& \mathsf{eval}'\ \mathsf{null}\ x \end{array}$$

As before we have:

$$\begin{array}{lcl} [\![\mathsf{let}\ y\ =\ e_1\ \mathsf{in}\ e_2]\!]_{z,l} &=& \mathsf{let}\ y\ =\ e_1\ \mathsf{in}\ [\![e_2]\!]_{z,l} \\ [\![\mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to e_1;\ \mathsf{inr}\ y \to e_2\ \}]\!]_{z,l} &=& \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to [\![e_1]\!]_{z,l};\ \mathsf{inr}\ y \to [\![e_2]\!]_{z,l}\ \} \\ [\![\mathsf{split}\ v\ \{\ (y,\ z) \to e\ \}]\!]_{z,l} &=& \mathsf{split}\ v\ \{\ (y,\ z) \to [\![e]\!]_{z,l}\ \} \end{array}$$

But now our bases cases are:

$$\begin{array}{lcl} [\![\mathsf{inl}\ w]\!]_{z,l} &=& \mathsf{memoize}\ l\ (\mathsf{inl}\ w);\ \mathsf{unroll}\ z\ w \\ [\![\mathsf{inr}\ w]\!]_{z,l} &=& \mathsf{memoize}\ l\ (\mathsf{inr}\ w);\ \mathsf{eval}'\ z\ w \\ [\![\mathsf{inr}\ (\mathsf{lazy}_F\ v)]\!]_{z,l} &=& \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \mathsf{lock}\ y\ \mathsf{in}\ F'\ z\ y\ v \end{array}$$

Compared to our previous calculation, little has changed: we only pass the zipper on to eval′ and call unroll once evaluation is finished. In the remaining indirection in the last case can be short-cut to just $F'\ z\ l\ v$ as discussed in Section 6.4.

7 Compressing Indirections

Our implementation differs from typical implementations of laziness in the way it memoizes thunks. While most implementations use the stack to remember which thunks need to be memoized, we write indirection nodes. This has the advantage of yielding constant stack usage, but it can lead to cases where long chains of indirection nodes are present in memory. This can be problematic for performance, especially if these long chains are traversed repeatedly. In fact, this can even lead to cases where code using lazy constructors has worse asymptotic complexity than code using

$$
\begin{tikzcd}
list[0] \arrow[r] \arrow[d] & list[1] \arrow[r] \arrow[d] & \ldots \arrow[r] & list[n] \arrow[d] \\
Next(\_) \arrow[r] & Next(\_) \arrow[r] & \ldots \arrow[r] & Done(1)
\end{tikzcd}
\qquad
\begin{tikzcd}
list[0] \arrow[r] \arrow[d] & list[1] \arrow[r] \arrow[d] & \ldots \arrow[r] & list[n] \arrow[d] \\
Indirect(\_) \arrow[r] & Indirect(\_) \arrow[r] & \ldots \arrow[r] & Done(1)
\end{tikzcd}
$$

Fig. 8. The chain before and after evaluation.

traditional thunks. In this section, we discuss this problem and propose a solution.

To motivate our discussion, we define a `chain` of lazy constructors. A `Next` node points to another element of the chain, which ends in a `Done` node. We also define the helper functions `next` and `eval`, which we will instantiate with different implementations later:

```
type chain
  Done( v : int )
  lazy Next( c : chain ) -> c
fun done(v : int)
  Done(v)
fun next(c : chain)
  Next(c)
fun eval(c : chain)
  match c
    Done(v) -> v
```

We can use this interface to `build` a list of lazy constructors where each element of the list points to the next element:

```
fun build(i : int)
  if i == 0
    then Cons(done(1), Nil)
    else
      val xs = build(i - 1)
      Cons(next(head(xs)), xs)
```

The function `build(n)` returns a list of length `n+1`, where each element of the list except the last points to a `Next` constructor that itself points to the next element in the list (see the left side of Figure 8). To evaluate the lazy constructor n elements from the end of the list, we have to evaluate n other lazy constructors. This becomes a problem when we try to traverse the list, for example by summing all elements:

```
fun traverse(n : int)
  sum(map(build(n), eval))
```

Despite looking like a linear-time function, `traverse` will take quadratic time. The problem here is that each `Next` is rewritten to an indirection node, which gives us a chain of indirection nodes of length n, see the right side of Figure 8. Each call to `eval` takes time linear in the length of the remaining list, and `map` matches on all elements in the list. The total cost of `traverse` is thus $(n+1)\ +\ n\ +\ \ldots 2\ +\ 1\ \in O(n^2)$.

7.1 Traditional Thunks

We can replicate this issue using traditional thunks, when we model the `chain` by alternating `chaincell` and `thunk` similar to streams:

```
type chaincell { Done( v : int ); Next( next : chain ) }
alias chain = thunk<chaincell>
fun done(v : int)
  delay { Done(v) }
fun next(c : chain)
  delay { Next(c) }
```

```
fun eval(c : chain)
  match c.force()
    Done(v) -> v
    Next(n) -> eval(n)
```

Here, it is obvious that the `eval` function takes time linear in the length of the remaining chain. This is true even in a runtime that short-cuts indirections (like OCaml), since the `Next` elements themselves take linear time to traverse. But similar to our implementation, this example only uses constant stack space: each `c.force()` call returns immediately and can easily be memoized. Furthermore, `eval` can be a tail-call: the `Next` constructors function as indirections similar to our implementation.

We can fix the asymptotic behaviour of our example by not using `Next` constructors at all. Instead, we just use a single thunk and force the next element of the chain in `next`:

```
type chaincell { Done( v : int ) }
alias chain = thunk<chaincell>

fun done(v : int)
  delay { Done(v) }

fun next(c : chain)
  delay { c.force() }

fun eval(c : chain)
  c.force().v
```

In this version `eval` is a constant-time operation, and `traverse` will take linear time. However, we pay for this in stack-usage: unlike the previous two version, this version has linear stack usage. The `eval` function calls `force` on a thunk generated by `next`, which will call `force` again on the next thunk. While `force` is in tail-position in `next`, the evaluation of the thunk itself is not a tail-call. Looking back at our implementation of `force`, we see that the call to `f()` has to be followed by the overwrite of `Lazy` with `Memo`:

```
type thunk<a>
  Memo( v : a )
  lazy Lazy( f : () -> a ) ->
    Memo( f() )
fun force( t : thunk<a> ) : a
  match t
    Memo(v) -> v
```

In fact, it is necessary that `force` uses the stack here we have to memoize the result of `f()` after it is run. This seems to put us in a bind: we can either have constant stack usage and suffer from quadratic time complexity in this example. Or we can obtain the correct time complexity, but at the cost of linear stack usage.

7.2 Laziness in the STG Machine

Haskell's evalation strategy of laziness corresponds to the latter option. The Haskell code for our example is:

```
build 0 = [1]
build i = (head xs) : xs
  where xs = build (i-1)

traverse n = foldl' (+) 0 (build n)
```

In practice, this code takes linear time, but also uses linear stack space. This happens since the STG machine writes an update frame on the stack each time it enters a thunk, which is then overwritten with the result of the evaluation [Peyton Jones 1992]. To force the first thunk of the list, Haskell has to force all thunks in the list, recursively pushing update frames on the stack. However, once the first thunk is forced, all other thunks have been updated to the final result as well and so the

rest of the list can be summed without any further lazy evaluation.[6]

To illustrate the difference between our implementation and the STG machine, consider our implementation of `eval` from Section 6.4:

$$
\begin{aligned}
\mathsf{eval}\ x\ &= \mathsf{lazy\ match}\ x \\
&\qquad \mathsf{lazy}_F\ l\ v \to \mathsf{case}\ (F\ v)\ \{\ \mathsf{inl}\ y \to \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \to (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \} \\
&\qquad \mathsf{indirect}\ y \to \mathsf{eval}\ y \\
&\qquad \mathsf{memo}\ y \to y
\end{aligned}
$$

If we find a new lazy thunk after evaluating $F\ v$, we write an indirection node and keep evaluating. We can replace the indirection node with a stack frame, by first evaluating and then memoizing the result:

$$
\begin{aligned}
\mathsf{eval}\ x\ &= \mathsf{lazy\ match}\ x \\
&\qquad \mathsf{lazy}_F\ l\ v \to \mathsf{case}\ (F\ v)\ \{\ \mathsf{inl}\ y \to \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \to \mathsf{memoize}\ l\ (\mathsf{eval}\ y)\ \} \\
&\qquad \mathsf{indirect}\ y \to \mathsf{eval}\ y \\
&\qquad \mathsf{memo}\ y \to y
\end{aligned}
$$

or equivalently, moving out the `memoize`:

$$
\begin{aligned}
\mathsf{eval}\ x\ &= \mathsf{lazy\ match}\ x \\
&\qquad \mathsf{lazy}_F\ l\ v \to \mathsf{memoize}\ l\ (\mathsf{case}\ (F\ v)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}) \\
&\qquad \mathsf{indirect}\ y \to \mathsf{eval}\ y \\
&\qquad \mathsf{memo}\ y \to y
\end{aligned}
$$

This version of `eval` is much closer to the typical implementation strategy of laziness. Just like the STG machine, it first pushes an update frame on the stack (memoize l), then runs the thunk ($F\ v$), checks if the result is still a thunk ($\mathsf{case} \ldots \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}$) and if so keeps evaluating it.

7.3 Compression of Indirections

However, we would like to keep the constant stack usage of our original implementation. We can achieve this by compressing long chains of indirection nodes once we traverse them again. That is we keep our old eval function, but just change the indirect y case to compress the chain:

$$
\begin{aligned}
\mathsf{eval}\ x\ &= \mathsf{lazy\ match}\ x \\
&\qquad \mathsf{lazy}_F\ l\ v \to \mathsf{case}\ (F\ v)\ \{\ \mathsf{inl}\ y \to \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \to (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \} \\
&\qquad \mathsf{indirect}\ y \to \mathsf{compress}\ y \\
&\qquad \mathsf{memo}\ y \to y
\end{aligned}
$$

To state the compression function, we need to extend our syntax slightly by allowing lazy match to return the location of an indirection node. Then we can define the compression function as follows:

$$
\begin{aligned}
\mathsf{compress}\ x\ &= \mathsf{lazy\ match}\ x \\
&\qquad \mathsf{indirect}\ l\ v \to \mathsf{memoize}\ l\ (\mathsf{compress}\ v) \\
&\qquad \mathsf{memo}\ v \to v
\end{aligned}
$$

Note that we do not give a case for lazy_F here, since an indirection chain is guaranteed to end in a memoized value.

[6]While the STG machine only creates indirection chains of length one, Marlow and Peyton Jones [2006] note that the stack usage can be reduced by rewriting stack frames into an indirection chain:

> *at garbage collection time [..] a useful optimisation is to collapse sequences of adjacent update frames into a single frame, by choosing one of the objects to be updated and making all the others be indirections to it.*

This would have the effect of creating indirection chains whose length is bounded by the number of GCs that have occurred during the creation of the chain. They do not describe how those indirection chains are shortened to length one.

Of course, we have gained little: while the original evaluation of the chain uses constant stack space, the traversal of the indirection chain now uses linear stack space. The advantage of this approach is that it enables us to improve the compress function to use constant stack space as well.

7.4 Stackless Compression

Our compression problem for indirection chains is similar to the problem of path compression in disjoint-set data structures [Leeuwen and Weide 1977; Tarjan and Van Leeuwen 1984; Weide 1980]. This allows us to use common solutions for the path compression problem for the indirection compression problem. In particular, we consider the three algorithms given in Section 3 of [Tarjan and Van Leeuwen 1984]:

7.4.1 Explicit Scans. In the naive compression algorithm, we use the stack to update the indirection nodes in reverse order of the chain. But actually, there is no requirement on the order in which we update the indirection nodes. The insight of the explicit scans algorithm is that we can traverse the chain twice: first to find the memoized value at the end of the chain and then to update all indirection nodes in the chain. Crucially, both traversals can be done in a tail-recursive manner.

$$
\begin{array}{ll}
\mathsf{compress}\ x & = \mathsf{replace}\ x\ (\mathsf{find}\ x) \\
\\
\mathsf{find}\ x = & \mathsf{lazy\ match}\ x \\
& \quad \mathsf{indirect}\ v \longrightarrow \mathsf{find}\ v \\
& \quad \mathsf{memo}\ v \longrightarrow v \\
\\
\mathsf{replace}\ x\ w & = \mathsf{lazy\ match}\ x \\
& \quad \mathsf{indirect}\ l\ v \longrightarrow \mathsf{memoize}\ l\ w;\ \mathsf{replace}\ v\ w \\
& \quad \mathsf{memo}\ v \longrightarrow v
\end{array}
$$

A counter-point to the explicit scans algorithm is that it traverses the chain twice. However, the naive compression algorithm also traverses the chain twice: first to enter all indirection nodes and then, unrolling the stack, to update them. However, we can get away with a single traversal of the chain by using the path splitting or path halving algorithms.

7.4.2 Path splitting. The path splitting algorithm shortens the path from each node to the memoized value by short-cutting each indirection node by one step. The name of the algorithm comes from the property that it splits the indirection chain into two shorter chains: one of all the even nodes and one of all the odd nodes. Each odd node is updated to point to the next odd node, while each even node is updated to point to the next even node:

$$
\begin{array}{ll}
\mathsf{compress}\ x & = \mathsf{lazy\ match}\ x \\
& \quad \mathsf{indirect}\ l\ v \longrightarrow \mathsf{lazy\ match}\ v \\
& \quad\quad \mathsf{indirect}\ w \longrightarrow \mathsf{indirect}\ l\ w;\ \mathsf{compress}\ v \\
& \quad\quad \mathsf{memo}\ w \longrightarrow \mathsf{memoize}\ l\ w \\
& \quad \mathsf{memo}\ v \longrightarrow v
\end{array}
$$

Unlike the two previous algorithms, the path splitting algorithm traverses the chain only once. But a possible objection to this algorithm is that many indirections remain in the chain. Indeed, as Figure 8 of Tarjan and Van Leeuwen [1984] shows, this algorithm may use $n \log(n)$ steps to compress a chain of length n, whereas the two previous algorithms use $2n$ steps.

7.4.3 Path halving. A slight improvement on the path splitting algorithm can be achieved by only applying it to all odd nodes in the chain. As soon as one short-cut is applied, we immediately continue compressing the chain from the next odd node:

$$
\begin{aligned}
\text{compress } x \;=\; & \text{lazy match } x \\
& \quad \text{indirect } l\ v \longrightarrow \text{lazy match } v \\
& \qquad \text{indirect } w \longrightarrow \text{indirect } l\ w;\ \text{compress } w \\
& \qquad \text{memo } w \longrightarrow \text{memoize } l\ w \\
& \quad \text{memo } v \longrightarrow v
\end{aligned}
$$

This algorithm also halves the length of the chain, but it has the advantage of using fewer writes to memory. Furthermore, it avoids splitting the chain into two. This is particularly helpful in our running example, since it means that repeated traversals of the chain from different nodes can benefit from the compression. We conjecture that unlike the path splitting algorithm, this algorithm runs in linear time, but leave this as an open problem.

8 Benchmarks

To test the runtime performance of lazy constructors, we have implemented all the lazy queues and heaps presented by Okasaki [1999], both using the standard approach with lazy thunks and with our new approach using lazy constructors. We benchmark them in a sequential setting without sharing of the persistent data structures. Since the laziness has no performance benefits in this setting (even in a theoretical or amortized sense), this allows to isolate the performance overhead of laziness itself. We compare the following systems and implementations:

- Koka (lazy): Koka 3.1.3 (`-O2 --no-debug`) using the implementation given by Okasaki with Koka's traditional lazy type.
- Koka (strict): Same as Koka (lazy) but with all laziness removed from the implementations.
- Koka (lazy cons): Same as Koka (lazy) but using our custom implementation with lazy constructors.
- OCaml (lazy): OCaml 4.14.2 (`-O2`, `OCAMLRUNPARAM="s=16M"`) using the implementations as given by Okasaki using the `Lazy.t` type.
- OCaml (strict): Same as OCaml (lazy), but with all laziness removed from the implementations.
- Haskell (lazy): GHC 9.10.1 (`-O2 -threaded -fworker-wrapper-cbv` and `+RTS -N8 -A32M -qb0`) using the implementations given by Okasaki.
- Haskell (strict): Same as Haskell (lazy) but compiled using `-XStrict`.
- Koka (no reuse): As we discuss in the next section, reuse analysis [Lorenzen and Leijen 2022; Reinking, Xie et al. 2021] has a large impact on the performance of the benchmarks in Koka, so we also test the Koka benchmarks with the `-fno-reuse` flag which disables reuse analysis.

For queues, we iterate the following procedure 1000 times: we `snoc` 100 000 integers into a queue, where in the first iteration we generate the integers at random and in the following iterations we `uncons` the elements out of the previous queue. This setup means that at any time the memory contains up to two queues with about 100 000 elements combined. This keeps the RSS of the program stable over the course of a benchmark run which particularly helps garbage collected languages. For heaps, we use the same method where we `deleteMin` from one heap and `insert` into the next, but we only do this for 100 heaps.

The benchmarks results are shown in Figure 9 where we normalize against the run time of the strict versions. The figure shows from top-to-bottom the benchmarks for Koka, Koka with no-reuse, OCaml, and Haskell. The results support our three main claims:

1. Compared to the strict version, traditional lazy thunks have a significant average performance overhead of 150% in Koka, 110% in OCaml, and 140% in Haskell. In contrast, lazy constructors have a smaller average overhead of just 43% in Koka.
2. Lazy constructors are always faster than traditional lazy thunks.
3. In some benchmarks, the performance overhead of lazy constructors is less than 25%, thus yielding lazy data structures that are close in performance to their strict counterparts, while maintaining superior theoretical properties.

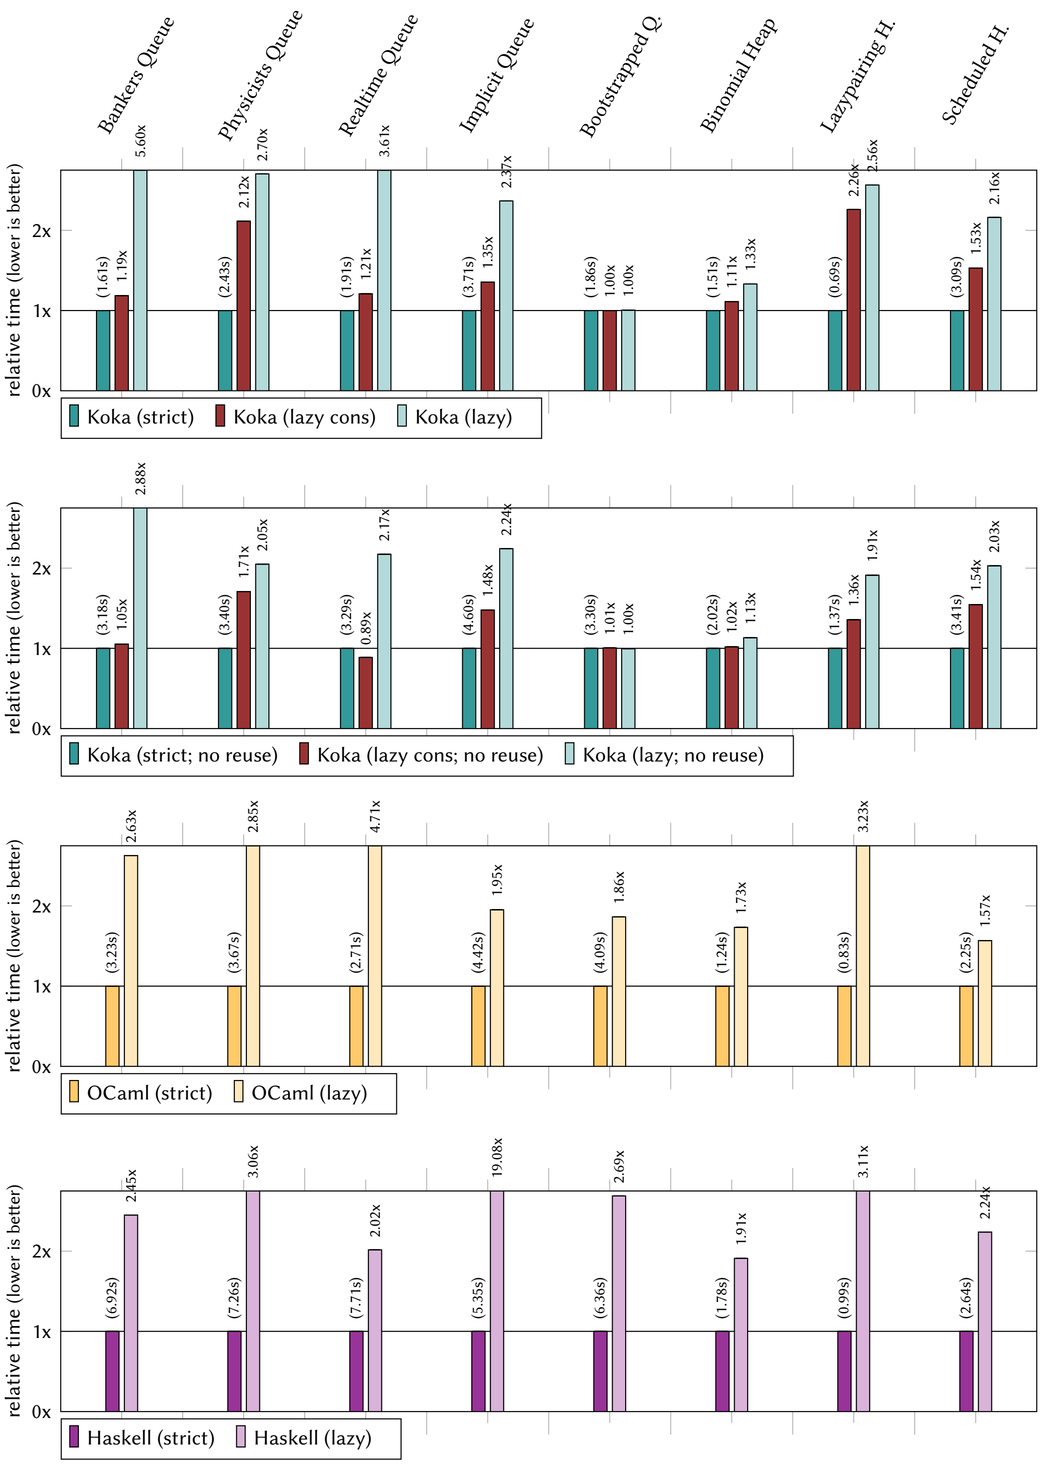

Fig. 9. Benchmarks on Apple M1 for Okasaki's queues and heaps in a sequential setting. Each graph shows results relative to the strict version of the benchmark. From top-to-bottom, the results are for Koka, Koka with no-reuse, OCaml, and Haskell.

8.1 Reuse analysis and laziness

We have split our analysis of Koka's performance into two parts: one with reuse analysis enabled and one with it disabled. Reuse analysis significantly improves the performance of strict data structures in Koka. However, it is much less useful in a lazy setting. To see this, consider the realtime queue, where we perform a lazy rotation using a `stream` type:

```
alias stream<a> = thunk<streamcell<a>>
type streamcell<a> =
  SNil
  SCons( x : a, xs : stream<a> )

fun rotate( front : stream<a>, rear : list<a>, sched : stream<a> ) : pure stream<a>
  match (front.eval(), rear)
    (SNil, Cons(y,_))         -> memo(SCons(y,sched ))
    (SCons(x,xs), Cons(y,ys)) -> delay{ SCons(x, rotate(xs,ys,memo(SCons(y,sched))) ) }
```

Here, `memo` creates a thunk from an existing value, while `delay` creates a thunk from a computation. Crucially, while reuse analysis can reuse the space from the `Cons` cell for the `SCons` cell passed to `memo` in the first branch, it does not apply to the constructors in the closure passed to `delay`. This is because reuse analysis never considers reuse opportunities under closures, since closures can be run arbitrarily often (although, in this case, this is a missing optimization since the `delay` function guarantees that the closure is only called once).

In contrast, when laziness is removed, the `stream` becomes a simple list and both the `front` and `rear` list can now be reused for the final result:

```
fun rotate( front : list<a>, rear : list<a>, sched : list<a> ) : pure list<a>
  match (front, rear)
    (Nil, Cons(y,_))          -> Cons(y, sched)
    (Cons(x,xs), Cons(y,ys)) -> Cons(x, rotate(xs,ys,Cons(y,sched)))
```

In practice, reuse analysis can significantly speed-up both the strict version of this data structure and our version with lazy constructors, while the traditional lazy version remains mostly unaffected.

9 Related Work

Codata. A lazy constructor can be viewed as a codefinition [Abel et al. 2013; Hagino 1989] for its lazy data type which combines with its strict constructors the codata operation `eval`. However, codefinitions are usually not memoized and we do not include an explicit function pointer in the runtime representation of lazy constructors.

Quotient Types. Lazy constructors $A \nrightarrow_F B$ can be viewed as a quotient type obtained by quotienting the sum-type $A + B$ by the step function, while recursive lazy constructors $A \nrightarrow\!\!\!\!\rightarrow_F B$ are the sum type quotiented by eval. In general, it is possible to mutate quotient types under the hood, where one element of an equivalence class can be swapped with any other element of the same class without breaking referential transparency. This trick was proposed by Selsam et al. [2020] and was a major inspiration for this work. They show how to implement pointer equality and hash-based memoization under a quotient so that they can be used in pure code, but do not consider the interaction with recursive types.

Semantics of Laziness. Nailing down the semantics of laziness was a longstanding problem. Launchbury [1993] was the first to give a natural semantics for lazy evaluation. The key insight was to require all function arguments to be let-bound; whenever an argument was evaluated, the corresponding heap location was updated. Sestoft [1997] has derived an abstract machine from Launchbury's lazy semantics. Even though the resulting machine is first order, it still needs to handle arbitrary lazy closures. Deriving a similar machine from our semantics for lazy constructors

would be interesting further work. More recently, Nakata and Hasegawa [2009] have given an alternative small-step and big-step semantics for laziness, that has been proven to be equivalent to the original natural semantics by Launchbury.

Laziness in OCaml. OCaml short-cuts indirection nodes during GC when the indirection points to a strict value. If the indirection node points to a lazy value or another indirection, it can not be short-cut to preserve the soundness of lazy pattern matching. In recent versions of OCaml, the short-cutting of indirection nodes is disabled when using instrumentation [Dolan 2018] and during major GC [Dolan 2021]. However, in practice, it appears that most lazy values are either forced early and their indirection short-cut during minor GC or not forced at all [Scherer et al. 2021]. OCaml's implementation of laziness currently does not support multicore [Scherer et al. 2021].

Laziness in Haskell. In Haskell, laziness is pervasive: all function arguments are evaluated lazily by default. As such, Haskell compilers use sophisticated techniques to make this efficient [Hartel 1991; Johnsson 1984; Marlow and Peyton Jones 2004 2006; Marlow et al. 2007; Peyton Jones 1992]. Just like our lazy constructors, early GHC implementations based on Spineless Tagless G-Machine [Peyton Jones 1992] used to update closures in-place if the closure was large enough. This approach was later abandoned [Marlow and Peyton Jones 1998,page 12] in favor of a more uniform return convention. As we showed in Section 6.3, we need to be careful to preserve lazy semantics when applying optimizations, and the GHC compiler takes great care to preserve laziness during its many program transformations [Peyton Jones and Lester 1991; Peyton Jones et al. 1996].

10 Limitations and Future work

The Global Nature of Defunctionalization. One drawback of lazy constructors is that they have to be declared up-front in the data type definition. This means that our approach does not support adding lazy constructors to a data type defined in a different module or library. To enable users of a library to define their own thunks, its author may add special lazy constructors such as `lazy SLazy(f : () -> stream<a>) -> f()` to their types. Users can then use these lazy constructors to define their own thunks, but they will be based on closures and not defunctionalized. To also benefit from defunctionalization, a programming language could combine lazy constructors with open data types [Löh and Hinze 2006]. However, open lazy constructors would have to carry a function pointer to their evaluation function, which makes it harder to see where reuse happens and we expect the performance may be slightly worse due to the indirect call.

GADTs and Effects. As typical for defunctionalized programs [Pottier and Gauthier 2004], some thunks can only be expressed as lazy constructors if the language supports GADTs [Cheney and Hinze 2003; Xi et al. 2003]. Koka currently assumes that for a type like `stream<a>`, all lazy constructors return a `stream<a>` and not, for example, a `stream<int>`. To lift this restriction, we would have to implement GADTs, so that the branches of the `lazy match` construct can make use of this type information. Furthermore, Koka assumes that lazy constructors perform no effects except divergence. We could allow lazy constructors to perform other effects (eg. throw an exception, write to a reference, I/O), but then every function that matches on a lazy data type would have to be annotated by those effects, leading to the expression problem if a new lazy constructor is added later.

Concurrency. Since Koka's reference counting scheme makes it possible to efficiently determine whether a thunk is thread-shared or not [Reinking, Xie et al. 2021; Ullrich and de Moura 2019], our implementation can use a fast-path and does not have to acquire or release locks in single-threaded programs. For thread-shared thunks, we implement blackholing using an atomic compare-and-swap operation, where other threads will busy-wait until the evaluation of the lazy constructor is complete. This is an inefficient strategy, which is likely to be slower than more advanced schemes [Harris

et al. 2005]. However, we plan to change the implementation in the future to block threads an a mutex. In particular, any lazy constructor has at least one field (to hold a possible indirection) and while being black-holed, we can use that field to store a mutex on which other threads block.

Syntax. Inspired by the syntax of lazy functions [1999, section 4.2; Wadler et al. 1998], we could define lazy constructors outside of data types:

```
lazy fun SAppend( s1 : stream<a>, s2 : stream<a> ) : div stream<a>
  match s1
    SCons(x,xx) -> SCons(x, SAppend(xx,s2))
    SNil        -> s2
```

But this syntax is more challenging to compile, since we need to assign a unique tag number to each lazy constructor and thus need to collect them from different parts of the file. In particular, lazy constructors defined in this manner outside of the file where their data type is defined could only be compiled using open data types.

Productivity. We are interested in understanding better under which circumstances lazy constructors are productive. All lazy constructors discussed in this paper recurse on their arguments and can thus easily be seen to be productive. But it is easy to come up with lazy constructors that are co-recursive and when we mix recursive and co-recursive lazy constructors, we can write programs that are no longer productive.

Acknowledgements

We are grateful to the anonymous reviewers and Gabriel Scherer for many helpful comments and suggestions. We thank Stephen Dolan and András Kovács for previous discussions about benchmarking. Sam Lindley was supported by UKRI Future Leaders Fellowship "Effect Handler Oriented Programming" (MR/T043830/1 and MR/Z000351/1).

References

Abel, Pientka, Thibodeau, and Setzer. 2013. Copatterns: Programming Infinite Structures by Observations. *ACM SIGPLAN Notices* 48 (1). ACM New York, NY, USA: 27–38.

Allain, Bour, Clément, Pottier, and Scherer. Jan. 2025. Tail Modulo Cons, OCaml, and Relational Separation Logic. *Proc. ACM Program. Lang.* 9 (POPL). Association for Computing Machinery, New York, NY, USA. doi:10.1145/3704915.

Altenkirch, Danielsson, and Kraus. 2017. Partiality, Revisited: The Partiality Monad as a Quotient Inductive-Inductive Type. In *International Conference on Foundations of Software Science and Computation Structures*, 534–549. Springer. doi:10.1007/978-3-662-54458-7_31.

Bagrel, and Spiwack. 2025. Destination Calculus: A Linear Lambda-Calculus for Purely Functional Memory Writes. *Proceedings of the ACM on Programming Languages* 9 (OOPSLA1). ACM New York, NY, USA: 253–279. doi:10.1145/3720423.

Bird. 1984. Using Circular Programs to Eliminate Multiple Traversals of Data. *Acta Informatica* 21. Springer: 239–250.

Capretta. 2005. General Recursion via Coinductive Types. *Log. Methods Comput. Sci.* 1 (2). doi:10.2168/LMCS-1(2:1)2005.

Chapman, Uustalu, and Veltri. 2019. Quotienting the Delay Monad by Weak Bisimilarity. *Mathematical Structures in Computer Science* 29 (1): 67–92. doi:10.1017/S0960129517000184.

Cheney, and Hinze. 2003. *First-Class Phantom Types*. Cornell University.

Dolan. 2018. AFL Stability Fixes for Objects (MPR#7725) and Lazy Values. `https://github.com/ocaml/ocaml/pull/1754`.

Dolan. 2021. Speed up GC by Prefetching during Marking. `https://github.com/ocaml/ocaml/pull/10195`.

Ganz, Friedman, and Wand. 1999. Trampolined Style. In *Proceedings of the Fourth ACM SIGPLAN International Conference on Functional Programming (ICFP '99)*, 18–27. doi:10.1145/317636.317779.

Gibbons. 2025. Turner, Bird, Eratosthenes: An Eternal Burning Thread. *Journal of Functional Programming* 35. Cambridge University Press: e5.

Gill. 2000. Debugging Haskell by Observing Intermediate Data Structures. In *Proceedings of the 2000 ACM SIGPLAN Workshop on Haskell, Haskell 2000, Satellite Event of PLI 2000, Montreal, Canada, September 17, 2000*, edited by Graham Hutton, 41:1. Electronic Notes in Theoretical Computer Science 1. Elsevier. doi:10.1016/S1571-0661(05)80538-9.

Hagino. 1989. Codatatypes in ML. *Journal of Symbolic Computation* 8 (6). Elsevier: 629–650.

Harris, Marlow, and Jones. 2005. Haskell on a Shared-Memory Multiprocessor. In *Proceedings of the 2005 ACM SIGPLAN Workshop on Haskell*, 49–61. doi:10.1145/1088348.1088354.

Hartel. 1991. Performance of Lazy Combinator Graph Reduction. *Software: Practice and Experience* 21 (3). Wiley Online Library: 299–329. doi:10.1002/SPE.4380210306.

Hood, and Melville. 1980. *Real Time Queue Operations in Pure Lisp*. Cornell University.

Huet. 1997. The Zipper. *Journal of Functional Programming* 7 (5): 549–554. doi:10.1017/S0956796897002864.

Johnsson. 1984. Efficient Compilation of Lazy Evaluation. In *Proceedings of the 1984 SIGPLAN Symposium on Compiler Construction*, 58–69. doi:10.1145/502874.502880.

Launchbury. 1993. A Natural Semantics for Lazy Evaluation. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 144–154. POPL '93. Charleston, South Carolina, USA. doi:10.1145/158511.158618.

Leeuwen, and Weide. 1977. Alternative Path Compression Techniques. Unknown Publisher.

Leijen. 2014. Koka: Programming with Row Polymorphic Effect Types. In *MSFP'14, 5th Workshop on Mathematically Structured Functional Programming*. doi:10.4204/EPTCS.153.8.

Leijen, and Lorenzen. Jan. 2023. Tail Recursion Modulo Context: An Equational Approach. *Proc. ACM Program. Lang.* 7 (POPL). doi:10.1145/3571233.

Leijen, and Lorenzen. Jan. 2025. Tail Recursion Modulo Context: An Equational Approach (extended Version). Microsoft Research. `https://www.microsoft.com/en-us/research/publication/tail-recursion-modulo-context-an-equational-approach-extended-version/`. Under submission.

Lorenzen, and Leijen. Sep. 2022. Reference Counting with Frame Limited Reuse. In *Proceedings of the 27th ACM SIGPLAN International Conference on Functional Programming (ICFP'2022)*. ICFP'22. Ljubljana, Slovenia. doi:10.1145/3547634.

Lorenzen, Leijen, and Swierstra. 2023. FP^2: Fully in-Place Functional Programming. In *Proc. ACM Program. Lang.*, 7:275–304. ICFP. doi:10.1145/3607840.

Lorenzen, Leijen, Swierstra, and Lindley. 2024. The Functional Essence of Imperative Binary Search Trees. *Proc. ACM Program. Lang.* 8 (PLDI): 518–542. doi:10.1145/3656398.

Lorenzen, Leijen, Swierstra, and Lindley. Jul. 2025. *First-Order Laziness*. `https://antonlorenzen.de/papers/lazycons.pdf`.

Löh, and Hinze. 2006. Open Data Types and Open Functions. In *Proceedings of the 8th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, 133–144. doi:10.1145/1140335.1140352.

Marlow, and Peyton Jones. Jan. 1998. The New GHC/Hugs Runtime System. `https://www.microsoft.com/en-us/research/publication/the-new-ghchugs-runtime-system/`.

Marlow, and Peyton Jones. 2004. Making a Fast Curry: Push/enter vs. Eval/apply for Higher-Order Languages. ACM New York, NY, USA, 4–15. doi:10.1145/1016850.1016856.

Marlow, and Peyton Jones. 2006. Making a Fast Curry: Push/enter vs. Eval/apply for Higher-Order Languages. *Journal of Functional Programming* 16 (4-5). Cambridge University Press: 415–449. doi:10.1017/S0956796806005995.

Marlow, Yakushev, and Peyton Jones. 2007. Faster Laziness Using Dynamic Pointer Tagging. ACM, 277–288. doi:10.1145/1291151.1291194.

Moura, and Ullrich. 2021. The Lean 4 Theorem Prover and Programming Language. In *Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, Virtual Event, July 12-15, 2021, Proceedings*, 12699:625–635. Lecture Notes in Computer Science. Springer. doi:10.1007/978-3-030-79876-5_37.

Nakata, and Hasegawa. 2009. Small-Step and Big-Step Semantics for Call-by-Need. *Journal of Functional Programming* 19 (6). Cambridge University Press: 699–722. doi:10.1017/S0956796809990219.

Okasaki. 1999. *Purely Functional Data Structures*. Cambridge University Press.

O'Neill. 2009. The Genuine Sieve of Eratosthenes. *Journal of Functional Programming* 19 (1). Cambridge University Press: 95–106.

Peyton Jones. 1992. Implementing Lazy Functional Languages on Stock Hardware: The Spineless Tagless G-Machine. *Journal of Functional Programming* 2 (2). Cambridge University Press: 127–202. doi:10.1017/S0956796800000319.

Peyton Jones, and Lester. 1991. A Modular Fully-Lazy Lambda Lifter in Haskell. *Software: Practice and Experience* 21 (5). Wiley Online Library: 479–506. doi:10.1002/spe.4380210505.

Peyton Jones, Partain, and Santos. 1996. Let-Floating: Moving Bindings to Give Faster Programs. In *Proceedings of the First ACM SIGPLAN International Conference on Functional Programming*, 1–12. doi:10.1145/232627.232630.

Pottier, and Gauthier. 2004. Polymorphic Typed Defunctionalization. ACM New York, NY, USA, 89–98. doi:10.1145/964001.964009.

Reinking, Xie, de Moura, and Leijen. 2021. 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. doi:10.1145/3453483.3454032.

Scherer, Bury, Leroy, Barnoy, Munch-Maccagnoni, White, Sivaramakrishnan, Bünzli, and Frisch. 2021. Discussing the Design of Lazy under Multicore. `https://github.com/ocaml-multicore/ocaml-multicore/issues/750`.

Schorr, and Waite. 1967. An Efficient Machine-Independent Procedure for Garbage Collection in Various List Structures. *Communications of the ACM* 10 (8). ACM New York, NY, USA: 501–506. doi:10.1145/363534.363554.

Schulte, and Grieskamp. 1992. Generating Efficient Portable Code for a Strict Applicative Language. In *Declarative Programming, Sasbachwalden 1991*, 239–252. Springer. doi:10.1007/978-1-4471-3794-8_16.

Selsam, Hudon, and Moura. 2020. Sealing Pointer-Based Optimizations behind Pure Functions. *Proceedings of the ACM on Programming Languages* 4 (ICFP). ACM New York, NY, USA: 1–20. doi:10.1145/3408997.

Sestoft. 1997. Deriving a Lazy Abstract Machine. *Journal of Functional Programming* 7 (3): 231–264. doi:10.1017/S0956796897002712.

Shaikhha, Fitzgibbon, Peyton Jones, and Vytiniotis. 2017. Destination-Passing Style for Efficient Memory Management. In *Proceedings of the 6th ACM SIGPLAN International Workshop on Functional High-Performance Computing*, 12–23. doi:10.1145/3122948.3122949.

Tarjan, and Van Leeuwen. 1984. Worst-Case Analysis of Set Union Algorithms. *Journal of the ACM (JACM)* 31 (2). ACM New York, NY, USA: 245–281.

Ullrich, and de Moura. Sep. 2019. 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. doi:10.1145/3412932.3412935.

Wadler, Taha, and MacQueen. 1998. How to Add Laziness to a Strict Language without Even Being Odd. In *SML'98, The SML Workshop*.

Weide. 1980. *Datastructures: An Axiomatic Approach and the Use of Binomial Trees in Developing and Analyzing Algorithms*. Mathematisch Centrum.

Xi, Chen, and Chen. 2003. Guarded Recursive Datatype Constructors. In *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 224–235. doi:10.1145/604131.604150.

A Data structures with Lazy Constructors

A.1 Bankers Queue

```
// Bankers queue  (see "Purely Functional Data Structures" by Chris Okasaki, Sect. 6.3.2)
// - worst-case  O(n), persistent O(n)
// - amortized   O(1), persistent O(1)
module queue/bankers

import queue

lazy type stream<a>
  SNil
  SCons( head : a, tail : stream<a> )
  lazy SAppRev( pre : stream<a>, post : list<a> ) ->
    match pre
      SNil        -> sreverse(post)
      SCons(x,xx) -> SCons(x,SAppRev(xx,post))

fun sreverse-acc( xs : list<a>, acc : stream<a> ) : stream<a>
  match xs
    Cons(x,xx) -> sreverse-acc(xx, SCons(x,acc))
    Nil        -> acc

fun sreverse( xs : list<a> ) : stream<a>
  sreverse-acc(xs,SNil)

pub fun list( s : stream<a> ) : div list<a>
  match s
    SCons(x,xx) -> Cons(x,xx.list)
    SNil -> Nil

fun stream(start : int, end : int)
  if (start>end) then SNil else SCons(start,stream(start+1,end))

pub fun test1()
  SAppRev(stream(1,3),[6,5,4]).list.println

struct bqueue<a>
  front : stream<a>
  front-len : int
  rear : list<a>
  rear-len : int

val bankers/empty = Bqueue( SNil, 0, Nil, 0 )

fun is-empty( ^q : bqueue<a> ) : bool
  q.front-len==0

fun size( ^q : bqueue<a> ) : int
  q.front-len + q.rear-len

fun check( Bqueue(front, front-len, rear, rear-len) : bqueue<a> ) : bqueue<a>
  if (front-len >= rear-len)
    then Bqueue(front, front-len, rear, rear-len)
    else Bqueue( SAppRev(front,rear), front-len + rear-len, Nil, 0)

fun snoc( Bqueue(front,front-len,rear,rear-len) : bqueue<a>, x : a ) : bqueue<a>
  Bqueue(front, front-len, Cons(x,rear), rear-len+1).check

fun uncons( Bqueue(front,front-len,rear,rear-len) : bqueue<a> ) : div maybe2<a,bqueue<a>>
  match front
    SCons(x,xx) -> Just2(x, Bqueue(xx,front-len - 1,rear,rear-len).check)
    SNil        -> Nothing2

pub fun main()
  benchmain(QueueI(bankers/empty, snoc, uncons))
```

A.2 Physicists Queue

```
module physicists-stack

import queue
```

```
lazy type thunk<a>
  Done( xs : list<a>, _p : pad )
  lazy Tail( thunk : thunk<a>, _p : pad ) ->
    match thunk
      Done( xs, p ) ->  Done( xs.tail, p )
  lazy AppRev( pre : list<a>, post : list<a> ) ->
    Done( list/(++)(pre, reverse(post)), Pad )
fun whnf( xs : thunk<a> ) : div list<a>
  match xs
    Done( xs', _ ) -> xs'
type queue<a>
  Queue( working : list<a>, lenf : int, front : thunk<a>, lenr : int, rear : list<a> )
fun checkw( q : queue<a> ) : div queue<a>
  match q
    Queue( working, lenf, front, lenr, rear ) ->
      match working
        Nil -> Queue( front.whnf(), lenf, front, lenr, rear )
        working -> Queue( working, lenf, front, lenr, rear )
val physicists/empty = Queue( Nil, 0, Done( Nil, Pad ), 0, Nil )
fun isEmpty( ^q : queue<a> ) : bool
  match q
    Queue( _, lenf, _, _, _ ) -> lenf == 0
fun size( ^q : queue<a> ) : int
  match q
    Queue( _, lenf, _, lenr, _ ) -> lenf + lenr
fun check( q : queue<a> ) : div queue<a>
  match q
    Queue( _, lenf, front, lenr, rear ) ->
      if (lenr <= lenf) then checkw(q)
      else
        val front' = front.whnf()
        checkw(Queue( front', lenf + lenr, AppRev( front', rear ), 0, Nil ))
fun snoc( q : queue<a>, x : a ) : div queue<a>
  match q
    Queue( working, lenf, front, lenr, rear ) ->
      check(Queue( working, lenf, front, lenr + 1, Cons( x, rear ) ))
fun uncons( q : queue<a> ) : div maybe2<a,queue<a>>
  match q
    Queue( working, lenf, front, lenr, rear ) ->
      match working
        Nil -> Nothing2
        Cons( x, xs ) -> Just2(x, check(Queue( xs, lenf - 1, Tail( front, Pad ), lenr, rear )))
pub fun main()
  benchmain(QueueI(physicists/empty, snoc, uncons))
```

A.3 Realtime Queue

```
// Realtime queue (see "Purely Functional Data Structures" by Chris Okasaki, Sect. 7.2)
// - worst-case  O(1), persistent O(1)
// - amortized   O(1), persistent O(1)
module queue/realtime
import std/core/unsafe
import queue
type stream<a>
  SNil
  SCons( head : a, tail : stream<a> )
```

```
  // `SRotate(pre,post,acc) == pre ++ reverse post ++ acc` with `|pre| + 1 == |post|`
  lazy SRotate( pre : stream<a>, post : list<a>, acc : stream<a> ) ->
    match pre
      SCons(x,xx) -> match post
                       Cons(y,yy) -> SCons(x, SRotate(xx, yy, SCons(y,acc)))
                       Nil        -> impossible("SRotate(SCons,Nil)")
      SNil        -> match post
                       Cons(y,_) -> SCons(y, acc)
                       Nil       -> impossible("SRotate(SNil,SNil)")
struct queue<a>
  front : stream<a>
  rear  : list<a>
  sched : stream<a>
val realtime/empty = Queue(SNil,Nil,SNil)
// evaluate one step into the front (by evaluating the schedule one step at a time)
fun queue( front : stream<a>, rear : list<a>, sched : stream<a> ) : div queue<a>
  match sched
    SCons(_,s) -> Queue(front,rear,s)
    SNil       -> val f = SRotate(front,rear,SNil ) in Queue(f,Nil,f)
fun snoc( Queue(front,rear,sched) : queue<a>, x : a ) : div queue<a>
  queue(front, Cons(x,rear), sched)
fun uncons( Queue(front,rear,sched) : queue<a> ) : div maybe2<a,queue<a>>
  match front
    SCons(x,xx) -> Just2(x, queue(xx,rear,sched))
    SNil        -> Nothing2
pub fun main()
  benchmain(QueueI(realtime/empty, snoc, uncons))
```

A.4 Implicit Queue

```
module implicit-stack
import std/core/unsafe
import queue
type digit<a>
  Zero
  One( x : a, p1 : pad, p2 : pad )
  Two( x : a, y : a, p : pad )
reference type pair3<a>
  Pair3( x : a, y : a, p : pad )
lazy type queue<a>
  Shallow( d : digit<a>, p1 : pad, p2 : pad )
  Deep( front : digit<a>, middle : queue<pair3<a>>, rear : digit<a> )
  lazy Tail( q : queue<a>, p1 : pad, p2 : pad ) ->
    match q
      Shallow(One(_, _, _), _, _) -> Shallow(Zero, p1, p2)
      Deep( Two(_, y, p), middle, rear ) -> Deep( One(y, Pad, p), middle, rear )
      Deep( One(_, _, _), middle, rear ) ->
        if isEmpty(middle)
          then Shallow(rear, p1, p2)
          else
            val Pair3(y, z, _) = head(middle)
            Deep( Two(y, z, Pad), Tail(middle, Pad, Pad), rear )
      _ -> unsafe-total { canthappen() }
```

```
  lazy Snoc( middle : queue<a>, y : a, p : pad ) ->
    match middle
      Shallow(Zero, p1, p2) ->
        Shallow(One(y, Pad, Pad), p1, p2)
      Shallow(One(x, _, p), p1, p2) ->
        Deep( Two( x, y, p ), Shallow( Zero, p1, p2 ), Zero )
      Deep( front, middle', Zero ) ->
        Deep( front, middle', One(y, Pad, Pad) )
      Deep( front, middle', One(x, _, p) ) ->
        Deep( front, Snoc(middle', Pair3(x, y, p), Pad), Zero )
      _ -> unsafe-total { canthappen() }
val implicit/empty = Shallow(Zero, Pad, Pad)
fun isEmpty( ^q : queue<a> ) : div bool
  match queue/lazy-force(q)
    Shallow(Zero) -> True
    _ -> False
fun snoc( q : queue<a>, y : a ) : div queue<a>
  lazy-step(Snoc(q, y, Pad))
fun head( ^q : queue<a> ) : div a
  match queue/lazy-force(q)
    Shallow(One(x, _, _), _, _) -> x
    Deep( One(x, _, _), _, _ ) -> x
    Deep( Two(x, _, _), _, _ ) -> x
    _ -> impossible()
fun uncons( q : queue<a> ) : div maybe2<a, queue<a>>
  match queue/lazy-force(q)
    Shallow(Zero) -> Nothing2
    Shallow(One(x), p1, p2) ->
      Just2( x, Shallow(Zero, p1, p2) )
    Deep( One(x), middle, rear ) ->
      if isEmpty(middle)
        then Just2( x, Shallow(rear, Pad, Pad) )
        else
          val Pair3(y, z, p) = head(middle)
          Just2( x, Deep( Two(y, z, p), Tail(middle, Pad, Pad), rear ) )
    Deep( Two(x, y, p), middle, rear ) ->
      Just2( x, Deep( One(y, Pad, p), middle, rear ) )
    _ -> canthappen()
pub fun main()
  benchmain(QueueI(implicit/empty, snoc, uncons))
```

A.5 Bootstrapped Queue

```
module bootstrapped
import queue
lazy type thunk<a>
  Done( x : list<a> )
  lazy Reverse( xs : list<a> ) ->
    Done( reverse( xs ) )
fun whnf( xs : thunk<a> ) : div list<a>
  match xs
    Done( xs' ) -> xs'
type queue<a>
  Empty
  Queue( lenfm : int, front : list<a>, m : queue<thunk<a>>, lenr : int, rear : list<a> )
val bootstrapped/empty = Empty
```

```
fun isEmpty( ^q : queue<a> ) : bool
  match q
    Empty -> True
    _ -> False
fun size( ^q : queue<a> ) : int
  match q
    Empty -> 0
    Queue( lenfm, _, _, lenr, _ ) -> lenfm + lenr
fun checkQ( q : queue<a> ) : div queue<a>
  match q
    Empty -> Empty
    Queue( lenfm, front, m, lenr, rear ) ->
      if lenr <= lenfm then
        Queue( lenfm, front, m, lenr, rear )
      else
        Queue( lenfm + lenr, front, snoc(m, Reverse(rear)), 0, Nil )
fun checkF( q : queue<a> ) : div queue<a>
  match q
    Empty -> Empty
    Queue( lenfm, front, m, lenr, rear ) ->
      match front
        Nil ->
          match uncons(m)
            Nothing2 ->
              Empty
            Just2( f, m' ) ->
              Queue( lenfm, f.whnf(), m', lenr, rear )
        front -> Queue( lenfm, front, m, lenr, rear )
fun snoc( q : queue<a>, x : a ) : div queue<a>
  match q
    Empty -> Queue( 1, Cons( x, Nil ), Empty, 0, Nil )
    Queue( lenfm, front, m, lenr, rear ) ->
      checkQ(Queue( lenfm, front, m, lenr + 1, Cons( x, rear ) ))
fun uncons( q : queue<a> ) : div maybe2<a,queue<a>>
  match q
    Empty -> Nothing2
    Queue( lenfm, front, m, lenr, rear ) ->
      match front
        Nil -> Nothing2 // impossible, queue would be Empty
        Cons( x, front' ) ->
          Just2( x, checkF(checkQ(Queue( lenfm - 1, front', m, lenr, rear ))) )
pub fun main()
  benchmain(QueueI(bootstrapped/empty, snoc, uncons))
```

A.6 Binomial Heap

```
module binomial-stack
import heap
alias elem = int
type tree
  Node( r : int, x : elem, c : list<tree> )
fun rank( ^t : tree ) : int
  match t
    Node( r, _, _ ) -> r
fun root( ^t : tree ) : elem
  match t
    Node( _, x, _ ) -> x
```

```
fun link( t1 : tree, t2 : tree, u : unit2 ) : tree
  match (t1, t2)
    (Node( r1, x1, c1 ), Node( r2, x2, c2 )) ->
      if x1 <= x2 then
        Node( r1 + 1, x1, Cons( Node( r2, x2, c2 ), c1 ) )
      else
        Node( r1 + 1, x2, Cons( Node( r1, x1, c1 ), c2 ) )
fun insTree( t : tree, ts : list<tree>, u : unit2 ) : list<tree>
  match ts
    Nil -> Cons( t, Nil )
    Cons( t', ts' ) ->
      if rank(t) < rank(t') then
        Cons( t, Cons( t', ts' ) )
      else
        insTree(link(t, t', Unit2( Pad, Pad )), ts', u)
fun mrg( ts1 : list<tree>, ts2 : list<tree> ) : div list<tree>
  match (ts1, ts2)
    (ts, Nil) -> ts
    (Nil, ts) -> ts
    (Cons( t1, ts1' ), Cons( t2, ts2' )) ->
      if rank(t1) < rank(t2) then
        Cons( t1, mrg(ts1', Cons( t2, ts2' )) )
      else if rank(t2) < rank(t1) then
        Cons( t2, mrg(Cons( t1, ts1' ), ts2') )
      else
        insTree(link(t1, t2, Unit2( Pad, Pad )), mrg(ts1', ts2'), Unit2( Pad, Pad ))
fun removeMinTree( ts : list<tree> ) : div maybe2<tree, list<tree>>
  match ts
    Nil -> Nothing2
    Cons( t, Nil ) -> Just2 ( t, Nil )
    Cons( t, ts ) ->
      match removeMinTree(ts)
        Just2 ( t', ts' )  ->
          if root(t) <= root(t') then
            Just2 ( t, ts )
          else
            Just2 ( t', Cons( t, ts' ) )
        Nothing2 -> Nothing2
lazy type heap
  Heap( ts : list<tree>, p : pad )
  lazy Insert( x : elem, ts : heap ) ->
    match ts
      Heap(ts, _) ->
        Heap( insTree(Node( 0, x, Nil ), ts, Unit2( Pad, Pad)), Pad )
  lazy Merge( ts1 : heap, ts2 : heap ) ->
    match ts1
      Heap(ts1, _) ->
        match ts2
          Heap(ts2, _) ->
            Heap( mrg(ts1, ts2), Pad )
  lazy MergeRev( ts1 : list<tree>, ts2 : list<tree> ) ->
    Heap( mrg(ts1.reverse(), ts2), Pad )
val binomial/empty = Heap( Nil, Pad )
fun isEmpty( h : heap ) : pure bool
  match h
    Heap( Nil ) -> True
    _ -> False
fun insert( x : elem, ts : heap ) : heap
  Insert( x, ts )
```

```
fun merge( ts1 : heap, ts2 : heap ) : heap
  Merge( ts1, ts2 )

fun splitMin( a : heap ) : pure maybe2<elem, heap>
  match a
    Heap( ts, _ ) ->
      match removeMinTree(ts)
        Just2 ( Node( _, x, ts1 ), ts2 )  ->
          Just2 ( x, MergeRev( ts1, ts2 ))
        Nothing2 -> Nothing2

pub fun main()
  benchmain(HeapI(binomial/empty, insert, splitMin))
```

A.7 Lazypairing Heap

```
module lazypairing

import heap

alias elem = int

lazy type heap
  Empty
  Heap( x : elem, h1 : heap, h2 : heap )
  lazy Link( a : heap, b : heap, m : heap ) ->
    merge( merge(a, b), m )

val pairing/empty = Empty

fun isEmpty( ^h : heap ) : div bool
  match heap/lazy-force(h)
    Empty -> True
    _ -> False

type neheap
  NEHeap( x : elem, h1 : heap, h2 : heap )

fun merge( a : heap, b : heap ) : div heap
  match a
    Empty -> b
    Heap( x, a1, a2 ) ->
      match b
        Empty -> a
        Heap( y, b1, b2 ) ->
          if x <= y then
            link(NEHeap( x, a1, a2 ), Heap( y, b1, b2 ))
          else
            link(NEHeap( y, b1, b2 ), Heap( x, a1, a2 ))

fun link( xbm : neheap, a : heap ) : div heap
  match xbm
    NEHeap( x, b, m ) ->
      match heap/lazy-force(b)
        Empty -> Heap( x, a, m )
        b’ -> Heap( x, Empty, Link(a, b’, m) )

fun insert( x : elem, a : heap ) : div heap
  merge(Heap( x, Empty, Empty ), a)

fun splitMin( a : heap ) : pure maybe2<elem, heap>
  match a
    Empty -> Nothing2
    Heap( x, a, m ) -> Just2 ( x, merge(a, m) )

pub fun main()
  benchmain(HeapI(pairing/empty, insert, splitMin))
```

A.8 Scheduled Heap

```
module scheduled

import heap
```

```
alias elem = int
type tree
  Node( x : elem, c : list<tree> )
lazy type stream
  SNil
  One( t : pair2<tree, pad>, xs : stream )
  Zero( p : pad, xs : stream )
  lazy InsTree( t : tree, p : pair2<stream, pad> ) ->
    match p
      Pair2( ds, _ ) ->
        // equivalent to: go_insTree(t, ds, Unit2(Pad, Pad), Unit2(Pad, Pad))
        // To get the best overwriting behaviour, we inline
        // go_insTree to rewrite InsTree into the head constructor
        match ds
          SNil ->
            One( Pair2( t, Pad ),SNil )
          Zero( _, ds' ) ->
            One( Pair2( t, Pad ), ds' )
          One( Pair2( t', _ ), ds' ) ->
            val t" = link(t, t', Unit2( Pad, Pad ))
            Zero( Pad, InsTree( t", Pair2( ds', Pad ) ))
  lazy Mrg( ts1 : stream, ts2 : stream ) ->
    go_mrg( ts1, ts2 )
value type heap
  Heap( working : stream, schedule : list<stream> )
val scheduled/empty = Heap( SNil, Nil )
fun isEmpty( h : heap ) : div bool
  match stream/lazy-force(h.working)
    SNil -> True
    _ -> False
fun root( ^t : tree ) : elem
  match t
    Node( x, _ ) -> x
fun link( t1 : tree, t2 : tree, u : unit2 ) : tree
  match (t1, t2)
    (Node( x1, c1 ), Node( x2, c2 )) ->
      if x1 <= x2 then
        Node( x1, Cons( Node( x2, c2 ), c1 ) )
      else
        Node( x2, Cons( Node( x1, c1 ), c2 ) )
fun go_insTree( t : tree, ts : stream, u1 : unit2, u2 : unit2 ) : div stream
  match stream/lazy-force(ts)
    SNil -> One( Pair2( t, Pad ),SNil )
    Zero( _, xs ) -> One( Pair2( t, Pad ), xs )
    One( Pair2( t', _ ), ds ) ->
      Zero( Pad, InsTree( link(t, t', Unit2( Pad, Pad )), Pair2( ds, Pad ) ))
```

```
fun go_mrg( ts1 : stream, ts2 : stream ) : div stream
  match stream/lazy-force(ts1)
    SNil -> ts2
    Zero( p, ts1' ) ->
      match stream/lazy-force(ts2)
        SNil -> Zero( p, ts1' )
        Zero( _, ts2' ) ->
          Zero( p, Mrg(ts1', ts2') )
        One( t2, ts2' ) ->
          One( t2, Mrg(ts1', ts2') )
    One( Pair2( t1, p ), ts1' ) ->
      match stream/lazy-force(ts2)
        SNil -> One( Pair2( t1, p), ts1' )
        Zero( _, ts2' ) ->
          One( Pair2( t1, p ), Mrg(ts1', ts2') )
        One( Pair2( t2, _ ), ts2' ) ->
          Zero( Pad, InsTree(link(t1, t2, Unit2( Pad, p )), Pair2( Mrg(ts1', ts2'), Pad )))
fun normalize( s : stream ) : div ()
  match stream/lazy-force(s)
    SNil -> ()
    Zero( _, xs ) -> normalize(xs)
    One( _, xs ) -> normalize(xs)
fun exec( sched : list<stream> ) : div list<stream>
  match sched
    Nil -> Nil
    Cons( s, sched' ) ->
      match stream/lazy-force(s)
        Zero( _, job ) -> Cons( job, sched' )
        _ -> sched'
fun insert( x : elem, h : heap ) : div heap
  match h
    Heap( ds, sched ) ->
      val ds' = go_insTree( Node( x, Nil ), ds, Unit2( Pad, Pad ), Unit2( Pad, Pad ) )
      Heap( ds', exec(exec(Cons( ds', sched ))) )
fun merge( h1 : heap, h2 : heap ) : div heap
  match (h1, h2)
    (Heap( ds1, _ ), Heap( ds2, _ )) ->
      val ds' = go_mrg( ds1, ds2 )
      normalize(ds')
      Heap( ds', Nil )
fun removeMinTree( ts : stream ) : div maybe2<tree, stream>
  match stream/lazy-force(ts)
    SNil -> Nothing2
    Zero( _, ds ) ->
      match removeMinTree(ds)
        Just2 ( t', ds' )  ->
          Just2 ( t', Zero( Pad, ds' ) )
        Nothing2 -> Nothing2
    One( Pair2( t, p ), ds ) -> match stream/lazy-force(ds)
      SNil -> Just2 ( t, SNil )
      _ ->
        match removeMinTree(ds)
          Just2 ( t', ds' )  ->
            if root(t) <= root(t') then
              Just2 ( t, Zero( Pad , ds ) )
            else
              Just2 ( t', One( Pair2( t, p ), ds' ) )
          Nothing2 -> Nothing2
```

```
fun listToStream( ts : list<tree> ) : stream
  match ts
    Nil -> SNil
    Cons( t, ts' ) -> One( Pair2( t, Pad ), listToStream( ts' ) )
fun mrgWithList( ts : list<tree>, ds : stream ) : div stream
  match ts
    Nil ->
      normalize(ds); ds
    Cons( t1, ts1 ) ->
      match stream/lazy-force(ds)
        SNil -> listToStream(ts)
        Zero( _, ds' ) ->
          One( Pair2( t1, Pad ), mrgWithList( ts1, ds' ) )
        One( Pair2( t2, _ ), ds' ) ->
          Zero( Pad, go_insTree(link(t1, t2, Unit2( Pad, Pad )),
                                mrgWithList(ts1, ds'),
                                Unit2( Pad, Pad ),
                                Unit2( Pad, Pad )) )
fun splitMin( h : heap ) : pure maybe2<elem, heap>
  match h
    Heap( ds, _ ) ->
      match removeMinTree(ds)
        Just2 ( Node( x, ts1 ), ts2 )  ->
          Just2 ( x, Heap( mrgWithList(ts1.reverse(), ts2), Nil ) )
        Nothing2 -> Nothing2
pub fun main()
  benchmain(HeapI(scheduled/empty, insert, splitMin))
```

B Benchmark

B.1 Queue

```
module queue
import std/core/undiv
import std/core/unsafe
import std/os/env
import std/num/random
import std/num/int32
import std/num/int64
pub reference type unit2
  Unit2(a : pad, b : pad)
pub reference type unit3
  Unit3(a : pad, b : pad, c : pad)
pub reference type unit4
  Unit4(a : pad, b : pad, c : pad, d : pad)
pub reference type pair2<a,b>
  Pair2(a : a, b : b)
pub fip fun canthappen(?kk-file-line : string) : div a
  impossible("canthappen")
pub alias rndstate = sfc
alias rndres = sfc-result
fun rnd-step( r : rndstate ) : rndres
  sfc-step(r)
fun rnd-init( s0 : int, s1 : int ) : rndstate
  (sfc-init32(s0.int32,s1.int32))
pub value struct queueI<q>
  qempty  : q
  qsnoc   : (q,int) -> div q
  quncons : q -> div maybe2<int,q>
```

```
fun bench-snoc( i : int32, qi : queueI<q>, rs : rndstate, queue : q ) : div q
  if i > 0.int32 then
    val step = rnd-step(rs)
    val q' = (qi.qsnoc)(queue, step.rnd.int)
    bench-snoc(i - 1.int32, qi, step.rstate, q')
  else queue
fun bench-pass-on( i : int32, from : q, to : q, qi : queueI<q>) : div q
  if i > 0.int32 then
    match (qi.quncons)(from)
      Just2( x, from' ) ->
        val to' = (qi.qsnoc)(to, x)
        bench-pass-on(i - 1.int32, from', to', qi)
      Nothing2 -> impossible("uncons failed")
  else
    to
fun bench-uncons( i : int32, queue : q, qi : queueI<q>, acc = 0.int32 ) : div int32
  if i > 0.int32 then
    match (qi.quncons)(queue)
      Just2( x, q ) -> bench-uncons(i - 1.int32, q, qi, acc + i*x.int32)
      Nothing2 -> impossible("uncons failed")
  else acc
fun bench-iterate( i : int32, n : int32, queue : q, qi : queueI<q>) : div int32
  if n > 1.int32 then
    val q = bench-pass-on(i, queue, qi.qempty, qi)
    bench-iterate(i, n - 1.int32, q, qi)
  else
    bench-uncons(i, queue, qi)
pub fun bench( ops : int32, n : int32, qi : queueI<q> ) : div int32
  val q = bench-snoc( ops, qi, rnd-init(42,43), qi.qempty)
  bench-iterate( ops, n, q, qi )
pub fun benchmain( qi : queueI<q>, ops : int = 100000, queues : int = 100) : io ()
  val n = get-args().head("").parse-int.default(queues).int32
  val sum = bench(ops.int32, n * 10.int32, qi) - 1097638789.int32
  println("Checksum: " ++ show(sum))
```

B.2 Heap

```
module heap
import std/core/undiv
import std/core/unsafe
import std/os/env
import std/num/random
import std/num/int32
import std/num/int64
pub ref type unit2
  Unit2(a : pad, b : pad)
pub ref type unit3
  Unit3(a : pad, b : pad, c : pad)
pub ref type unit4
  Unit4(a : pad, b : pad, c : pad, d : pad)
pub ref type pair2<a,b>
  Pair2(a : a, b : b)
pub fip fun canthappen() : div a
  canthappen()
pub alias rndstate = sfc
alias rndres = sfc-result
fun rnd-step( r : rndstate ) : rndres
  sfc-step(r)
```

```
fun rnd-init( s0 : int, s1 : int ) : rndstate
  (sfc-init32(s0.int32,s1.int32))
pub value struct heapI<h>
  empty : h
  insert : (int,h) -> pure h
  splitMin : h -> pure maybe2<int,h>
fun bench-insert( i : int32, hi : heapI<h>, rs : rndstate, heap : h ) : pure h
  if i > 0.int32 then
    val step = rnd-step(rs)
    val h' = (hi.insert)(step.rnd.int, heap)
    bench-insert(i - 1.int32, hi, step.rstate, h')
  else
    heap
fun bench-pass-on( i : int32, from : h, to : h, hi : heapI<h>) : pure h
  if i > 0.int32 then
    match (hi.splitMin)(from)
      Just2( x, from' ) ->
        val to' = (hi.insert)(x, to)
        bench-pass-on(i - 1.int32, from', to', hi)
      Nothing2 -> throw("splitMin failed")
  else
    to
fun bench-splitMin( i : int32, heap : h, hi : heapI<h>, acc = 0.int32 ) : pure int32
  if i > 0.int32 then
    match (hi.splitMin)(heap)
      Just2( x, heap' ) -> bench-splitMin(i - 1.int32, heap', hi, acc + i*x.int32)
      Nothing2 -> throw("splitMin failed")
  else
    acc
fun bench-iterate( i : int32, n : int32, heap : h, hi : heapI<h>) : pure int32
  if n > 1.int32 then
    val h = bench-pass-on(i, heap, hi.empty, hi)
    bench-iterate(i, n - 1.int32, h, hi)
  else
    bench-splitMin(i, heap, hi)
pub fun bench( ops : int32, n : int32, hi : heapI<h> ) : pure int32
  val h = bench-insert( ops, hi, rnd-init(42,43), hi.empty)
  bench-iterate( ops, n, h, hi )
pub fun benchmain( hi : heapI<h>) : io ()
  val n = get-args().head("").parse-int.default(1).int32
  val ops = 100000.int32
  val sum = bench(ops, n, hi) + 1973053443.int32
  println("Checksum: " ++ show(sum))
```

C Proofs

C.1 Soundness of Formalization

C.1.1 Steps. Since we use a step-indexed logical relation, we have to annotate our natural semantics with a step count. This is straightforward, although note that evaluating a lazy constructor can take between 1 step to recall the result or $k+1$ steps to evaluate it. Since we allow the store extension relation $\Gamma \sqsubseteq \Delta$ to evaluate lazy constructors, we also need to add a step count to it.

$$\frac{}{\Gamma : v \Downarrow_0 \Gamma : v}\ \textsc{value} \qquad \frac{\Gamma : e_1 \Downarrow_k \Delta : v \quad \Delta : e_2[v/x] \Downarrow_j \Theta : w}{\Gamma : \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 \Downarrow_{k+j+1} \Theta : w}\ \textsc{let}$$

$$\frac{F(x) = e \in \Sigma \quad \Gamma : e[v/x] \Downarrow_k \Delta : w}{\Gamma : F\ v \Downarrow_{k+1} \Delta : w}\ \textsc{app} \qquad \frac{\Gamma : e[v_1/x,\ v_2/y] \Downarrow_k \Delta : w}{\Gamma : \mathsf{split}\ (v_1,\ v_2)\ \{\ (x,\ y) \to e\ \} \Downarrow_{k+1} \Delta : w}\ \textsc{split}$$

$$\frac{z\ \text{fresh}}{\Gamma : lv \Downarrow_1 (\Gamma,\ z \mapsto lv) : z}\ \textsc{lazy} \qquad \frac{}{\Gamma : \mathsf{unfold}\ (\mathsf{fold}\ v) \Downarrow_1 \Gamma : v}\ \textsc{unfold}$$

$$\frac{z \mapsto \mathsf{memo}\ v \in \Gamma}{\Gamma : \mathsf{step}\ z \Downarrow_1 \Gamma : v}\ \textsc{recall} \qquad \frac{\Gamma : e_i[v/x_i] \Downarrow_{k_i} \Delta : w}{\Gamma : \mathsf{case}\ (\mathsf{in}_i\ v)\ \{\ \mathsf{in}_l\ x_l \to e_l;\ \mathsf{in}_r\ x_r \to e_r\ \} \Downarrow_{k_i+1} \Delta : w}$$

$$\frac{\Gamma : F\ v \Downarrow_k \Delta : w}{(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) : \mathsf{step}\ x \Downarrow_{k+1} (\Delta,\ x \mapsto \mathsf{memo}\ w) : w}\ \textsc{step}$$

$$\frac{}{\Gamma \sqsubseteq_k \Gamma}\ \textsc{refl} \qquad \frac{}{\Gamma \sqsubseteq_1 \Gamma,\ x \mapsto lv}\ \textsc{extend}$$

$$\frac{\Gamma \sqsubseteq_k \Delta \quad \Delta \sqsubseteq_j \Theta}{\Gamma \sqsubseteq_{k+j} \Theta}\ \textsc{trans} \qquad \frac{\Gamma : F\ v \Downarrow_k \Delta : w}{\Gamma,\ x \mapsto \mathsf{lazy}_F\ v \sqsubseteq_{k+1} \Delta,\ x \mapsto \mathsf{memo}\ w}\ \textsc{eval}$$

C.1.2 Store extension.

Lemma 4. (*Store extension is reflexive.*)
For all stores Γ, $\Gamma \sqsubseteq_0 \Gamma$.

Proof. By the REFL rule.

Lemma 5. (*Store extension is transitive.*)
If $\Gamma \sqsubseteq_k \Delta$ and $\Delta \sqsubseteq_j \Theta$, then $\Gamma \sqsubseteq_{k+j} \Theta$.

Proof. By the TRANS rule.

Lemma 6. (*Store extension may take more steps*)
If $\Gamma \sqsubseteq_k \Delta$ then $\Gamma \sqsubseteq_{k+1} \Delta$.

Proof. Use the REFL rule to obtain $\Gamma \sqsubseteq_1 \Gamma$ and then the TRANS rule to obtain the result.

Lemma 7. (*Growing the store does not change evaluation*)
If $\Gamma : e \Downarrow_k \Delta : v$, then $(\Gamma, x \mapsto lv) : e \Downarrow_k (\Delta, x \mapsto lv) : v$ for any $x \notin \Gamma$.

Proof. By induction on $\Gamma : e \Downarrow_k \Delta : v$, where we note that $\mathrm{dom}(\Delta) \setminus \mathrm{dom}(\Gamma)$ consists only of fresh variables.

Lemma 8. (*Growing the store does not change store extension.*)
If $\Gamma \sqsubseteq_k \Delta$, then $(\Gamma,\ x \mapsto lv) \sqsubseteq_k (\Delta,\ x \mapsto lv)$ for any $x \notin \Delta$.

Proof. By induction on $\Gamma \sqsubseteq_k \Delta$.

Case REFL: Use the REFL rule with the extra binding.

Case EXTEND: Use the EXTEND rule with the extra binding.

Case TRANS: Directly using the inductive hypothesis.

Case EVAL: By Lemma 7.

Lemma 9. (*Memos stay memos during evaluation*)
If $\Gamma : e \Downarrow_k \Delta : v$ and $x \mapsto \mathsf{memo}\ w \in \Gamma$, then $x \mapsto \mathsf{memo}\ w \in \Delta$.

Proof. By induction on k and case-split on $\Gamma : e \Downarrow_k \Delta : v$ where we note that no variables are removed from the store except those pointing to a $\mathsf{lazy}_F\ w$.

Lemma 10. (*Memos stay memos in the store*)
If $\Gamma \sqsubseteq_k \Delta$ and $x \mapsto \mathsf{memo}\ v \in \Gamma$, then $x \mapsto \mathsf{memo}\ v \in \Delta$.

Proof. By induction on $\Gamma \sqsubseteq_k \Delta$ using 9.

Lemma 11. (*Removing lazies from the store does not change evaluation*)
If $(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) : e \Downarrow_k (\Delta,\ x \mapsto \mathsf{lazy}_F\ v) : w$, then $\Gamma : e \Downarrow_k \Delta : w$.

Proof. By induction on $(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) : e \Downarrow_k (\Delta,\ x \mapsto \mathsf{lazy}_F\ v) : v$, where we note that only RECALL and STEP check for the existence of a value in the store, but RECALL only fires for lazy values and STEP replaces them with a memo that stays in the store by Lemma 9.

Lemma 12. (*Removing lazies from the store does not change store extension.*)
If $(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k (\Delta,\ x \mapsto \mathsf{lazy}_F\ v)$, then $\Gamma \sqsubseteq_k \Delta$.

Proof. By induction on $\Gamma \sqsubseteq_k \Delta$.

Case REFL: Use the REFL rule without the extra binding.

Case EXTEND: Use the EXTEND rule without the extra binding.

Case TRANS: Directly using the inductive hypothesis.

Case EVAL: By Lemma 11.

Lemma 13. (*Evaluation extends the store*)
If $\Gamma : e \Downarrow_k \Delta : v$, then $\Gamma \sqsubseteq_k \Delta$.

Proof. By induction on k. We use the inductive hypothesis that for all Γ, e and $j < k$, $\Gamma : e \Downarrow_j \Delta : v$ implies $\Gamma \sqsubseteq_j \Delta$. If $k = 0$, then only VALUE applies, which corresponds to REFL. If $k > 0$:

Case LAZY: By EXTEND.

Case UNFOLD: By REFL.

Case RECALL: By REFL.

Case LET: By the inductive hypothesis and TRANS.

Case APP: By the inductive hypothesis.

Case SPLIT: By the inductive hypothesis.

Case CASE: By the inductive hypothesis.

Case STEP: By EVAL.

C.1.3 Logical relation. Our logical relation is formally defined as:

$$
\begin{array}{ll}
\mathcal{V}_k[\![1]\!] & := \{ (\Delta,\ ()) \} \\
\mathcal{V}_k[\![A + B]\!] & := \{ (\Delta,\ \mathsf{inl}\ v) \mid (\Delta,\ v) \in \mathcal{V}_k[\![A]\!] \} \cup \{ (\Delta,\ \mathsf{inr}\ v) \mid (\Delta,\ v) \in \mathcal{V}_k[\![B]\!] \} \\
\mathcal{V}_k[\![A \times B]\!] & := \{ (\Delta,\ (v,\ w)) \mid (\Delta,\ v) \in \mathcal{V}_k[\![A]\!],\ (\Delta,\ w) \in \mathcal{V}_k[\![B]\!] \} \\
\mathcal{V}_k[\![\mu\alpha.\ A]\!] & := \{ (\Delta,\ \mathsf{fold}\ v) \mid \forall j < k.\ (\Delta,\ v) \in \mathcal{V}_j[\![A[\mu\alpha.\ A/\alpha]]\!] \} \\
\mathcal{V}_k[\![A \nrightarrow_F B]\!] & := \{ ((\Delta,\ z \mapsto \mathsf{lazy}_F\ v),\ z) \mid (\Delta,\ v) \in \mathcal{V}_k[\![A]\!],\ \forall j \leqslant k, \Theta.\ \Delta \sqsubseteq_j \Theta \Rightarrow F\ v \in \mathsf{E}_{k-j,\Theta}[\![B]\!] \} \\
& \quad \cup \{ ((\Delta,\ z \mapsto \mathsf{memo}\ v),\ z) \mid (\Delta,\ v) \in \mathcal{V}_k[\![B]\!] \}
\end{array}
$$

$$\mathsf{E}_{k,\Delta}[\![A]\!] := \{ e \mid \forall j < k.\ \forall \Theta,\ v.\ (\Delta : e \Downarrow_j \Theta : v) \Rightarrow \Delta \sqsubseteq_j \Theta \text{ and } (\Theta,\ v) \in \mathcal{V}_{k-j}[\![A]\!] \}$$

$$
\begin{array}{ll}
\mathcal{G}_{k,\Delta}[\![\varnothing]\!] & := \{ \cdot \} \\
\mathcal{G}_{k,\Delta}[\![\Gamma,\ x : A]\!] & := \{ \sigma[x \mapsto v] \mid \sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!],\ (\Delta,\ v) \in \mathcal{V}_k[\![A]\!] \}
\end{array}
$$

We have the usual properties:

Lemma 14. (*Values are sound expressions.*)
If $(\Delta,\ v) \in \mathcal{V}_k[\![A]\!]$, then $v \in \mathsf{E}_{k,\Delta}[\![A]\!]$.

Proof.

$\Delta : v \Downarrow_0 \Delta : v$	(1), by VALUE
$\Delta \sqsubseteq_0 \Delta$	(2), by Lemma 21
$(\Delta,\ v) \in \mathcal{V}_{k-0}[\![A]\!]$	(3), by assumption
$v \in \mathsf{E}_{k,\Delta}[\![A]\!]$	(4), by (1), (2), (3)

Lemma 15. (*On values, the expression denotation is the value denotation.*)
If $v \in \mathsf{E}_{k,\Delta}[\![A]\!]$, then $(\Delta,\ v) \in \mathcal{V}_k[\![A]\!]$.

Proof.

$\Delta : v \Downarrow_0 \Delta : v$	(1), by VALUE
$(\Delta,\ v) \in \mathcal{V}_{k-0}[\![A]\!]$	(3), by assumption

Lemma 16. (*Downward closure*)
If $(\Delta,\ v) \in \mathcal{V}_k[\![A]\!]$, then $(\Delta,\ v) \in \mathcal{V}_j[\![A]\!]$ for all $j \leqslant k$.

Proof. By induction on k. If $k = 0$, then obvious. Else: induction on A.

Case $A = 1$: Obvious

Case $A = A' + B$ or $A = A' \times B$: Follows directly from the inner inductive hypothesis.

Case $A = \mu\alpha.\ A'$: Follows directly from the outer inductive hypothesis.

Case $A = A' \nrightarrow_F B$. Let $(\Delta,\ z) \in \mathcal{V}_k[\![A]\!]$ with $z \mapsto \mathsf{memo}\ v \in \Delta$. Then the claim follows directly from the inner inductive hypothesis.

Case $A = A' \nrightarrow_F B$. Let $(\Delta,\ z) \in \mathcal{V}_k[\![A]\!]$ with $z \mapsto \mathsf{lazy}_F\ v \in \Delta$.

$(\Delta,\ z) \in \mathcal{V}_k[\![A' \nrightarrow_F B]\!]$	(1), by assumption
$z \mapsto \mathsf{lazy}_F\ v \in \Delta$	(2), by assumption
$(\Delta,\ v) \in \mathcal{V}_k[\![A]\!]$	(3), by (1) and (2)
$\forall j \leqslant k, \Theta.\ \Delta \sqsubseteq_j \Theta \Rightarrow F\ v \in \mathsf{E}_{k-j,\Theta}[\![B]\!]$	(4), by (1) and (2)
$j' \leqslant k$	(5), by assumption
$(\Delta,\ v) \in \mathcal{V}_{j'}[\![A]\!]$	(6), apply inner inductive hypothesis to (3) and (5)
$\forall j \leqslant j', \Theta.\ \Delta \sqsubseteq_j \Theta \Rightarrow F\ v \in \mathsf{E}_{k-j,\Theta}[\![B]\!]$	(7), by (4) and (5)
$(\Delta,\ z) \in \mathcal{V}_{j'}[\![A' \nrightarrow_F B]\!]$	(8), by (6) and (7)

Lemma 17. (*Growing the Store preserves types.*)
If $(\Delta,\ v) \in \mathcal{V}_k[\![A]\!]$, then $((\Delta,\ x \mapsto lv),\ v) \in \mathcal{V}_k[\![A]\!]$ for any $x \notin \Delta$.

Proof. By induction on k and then on A.

Case $A = 1$ or $A = A' + B$ or $A = A' \times B$: Follows directly from the inner inductive hypothesis.

Case $A = \mu\alpha.\, A'$: Follows directly from the outer inductive hypothesis.

Case $A = A' \nrightarrow_F B$: Follows from the inner inductive hypothesis and by using the EXTEND rule to obtain $\Delta \sqsubseteq (\Delta,\ x \mapsto lv)$.

Lemma 18. (*Evaluation preserves types.*)
If $((\Gamma, x \mapsto \mathsf{lazy}_F\ v'),\ v) \in \mathcal{V}_k[\![\ A\ ]\!]$ and $\Gamma : F\ v' \Downarrow_j \Delta : w$ for $j < k$ and $\forall v.\ (\Gamma,\ v) \in \mathcal{V}_k[\![A]\!] \Rightarrow (\Delta,\ v) \in \mathcal{V}$
then $((\Delta, x \mapsto \mathsf{memo}\ w),\ v) \in \mathcal{V}_{k-j}[\![\ A\ ]\!]$.
Proof. By induction on k and then on A.

Case $A = 1$ or $A = A' + B$ or $A = A' \times B$: Follows directly from the inner inductive hypothesis.

Case $A = \mu\alpha.\, A'$: Follows directly from the outer inductive hypothesis.

Case $A = A' \nrightarrow_F B$: Let $v = z$. Case $x = z$:

$((\Gamma, x \mapsto \mathsf{lazy}_F\ v'),\ x) \in \mathcal{V}_k[\![\ A' \nrightarrow_F B\ ]\!]$	(1), by assumption
$(\Gamma,\ v') \in \mathcal{V}_k[\![\ A'\ ]\!]$	(2), by (1)
$\forall j \leqslant k,\ \Theta.\ \Gamma \sqsubseteq_j \Theta \Rightarrow F\ v' \in \mathsf{E}_{k-j,\,\Theta}[\![B]\!]$	(3), by (1)
$F\ v' \in \mathsf{E}_{k,\,\Gamma}[\![B]\!]$	(4), instantiate (3)
$\forall j < k.\ \forall \Delta,\ v.\ (\Gamma : F\ v' \Downarrow_j \Delta : w) \Rightarrow \Gamma \sqsubseteq_j \Delta$ and $(\Delta,\ w) \in \mathcal{V}_{k-j}[\![\ B\ ]\!]$	(5), unroll (4)
$(\Delta,\ w) \in \mathcal{V}_{k-j}[\![\ B\ ]\!]$	(6), simplify (5)
$((\Delta,\ x \mapsto \mathsf{memo}\ w),\ x) \in \mathcal{V}_{k-j}[\![\ A' \nrightarrow_F B\ ]\!]$	(7), by definition with (6)

Case $x \neq z$:

$((\Gamma, x \mapsto \mathsf{lazy}_F\ v', z \mapsto \mathsf{lazy}_F\ v''),\ z) \in \mathcal{V}_k[\![\ A' \nrightarrow_F B\ ]\!]$	(1), by assumption
$((\Gamma, x \mapsto \mathsf{lazy}_F\ v'),\ v'') \in \mathcal{V}_k[\![\ A'\ ]\!]$	(2), by (1)
$((\Delta, x \mapsto \mathsf{lazy}_F\ v'),\ v'') \in \mathcal{V}_{k-j}[\![\ A'\ ]\!]$	(3), by (2) and inductive hypothesis
$\forall l \leqslant k,\ \Theta.\ (\Gamma, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_l \Theta \Rightarrow F\ v \in \mathsf{E}_{k-l,\,\Theta}[\![B]\!]$	(4), by (1)
$\Gamma \sqsubseteq_j \Delta$	(5), by Lemma 13
$(\Gamma, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_j (\Delta, x \mapsto \mathsf{lazy}_F\ v')$	(6), by Lemma 24
$\forall l \leqslant k,\ \Theta.\ (\Delta, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_l \Theta \Rightarrow (\Gamma, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_{j+l} \Theta$	(7), by Lemma 22
$\forall l \leqslant k,\ \Theta.\ (\Delta, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_l \Theta \Rightarrow F\ v \in \mathsf{E}_{k-j-l,\,\Theta}[\![B]\!]$	(8), by (4) and (7)
$((\Delta, x \mapsto \mathsf{lazy}_F\ v', z \mapsto \mathsf{lazy}_F\ v''),\ z) \in \mathcal{V}_{k-j}[\![\ A' \nrightarrow_F B\ ]\!]$	(4), by definition with (3) and (8)

Lemma 19. (*Store extension preserves types.*)
If $(\Delta,\ v) \in \mathcal{V}_k[\![\ A\ ]\!]$ and $\Delta \sqsubseteq_j \Theta$, then $(\Theta,\ v) \in \mathcal{V}_{k-j}[\![\ A\ ]\!]$.
Proof. By induction on j and then on $\Delta \sqsubseteq_j \Theta$.

Case REFL: Follows from Lemma 29.

Case TRANS:

$\Delta \sqsubseteq_j \Delta' \sqsubseteq_l \Theta$	(1), by assumption
$(\Delta',\ v) \in \mathcal{V}_{k-j}[\![\ A\ ]\!]$	(2), by outer inductive hypothesis
$(\Theta,\ v) \in \mathcal{V}_{k-j-l}[\![\ A\ ]\!]$	(3), by outer inductive hypothesis

Case EXTEND: Follows from Lemma 30.

Case EVAL:

$\Gamma,\ x \mapsto \mathsf{lazy}_F\ v' \sqsubseteq_{j+1} \Delta,\ x \mapsto \mathsf{memo}\ w$	(1), by assumption
$((\Gamma, x \mapsto \mathsf{lazy}_F\ v'),\ v) \in \mathcal{V}_k[\![\ A\ ]\!]$	(1), by assumption
$\Gamma : F\ v' \Downarrow_j \Delta : w$	(3), by (1)
$\Gamma \sqsubseteq_j \Delta$	(4), by (3) and Lemma 13
$\forall v.\ (\Gamma,\ v) \in \mathcal{V}_k[\![A]\!] \Rightarrow (\Delta,\ v) \in \mathcal{V}_{k-j}[\![\ A\ ]\!]$	(5), by (4) and outer inductive hypothesis
$((\Delta, x \mapsto \mathsf{memo}\ w),\ v) \in \mathcal{V}_{k-j}[\![\ A\ ]\!]$	(6), by Lemma 31

C.1.4 Type Soundness.

$\Gamma \vDash e : A := \forall k \geqslant 0,\ \Delta,\ \sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!].\ \sigma(e) \in \mathsf{E}_{k,\Delta}[\![A]\!]$

Theorem 4. (*Semantic Type Soundness.*)
If $\varnothing \vDash e : A$ and $\varnothing : e \Downarrow_k \Delta : v$, then $(\Delta, v) \in \mathcal{V}_j[\![A]\!]$ for any $j > 0$.

Proof. Let $\Gamma = \varnothing$. We have $\cdot \in \mathcal{G}_{k+j,\Gamma}[\![\varnothing]\!]$. Then $\varnothing \vDash e : A$ implies that $\cdot(e) = e \in \mathsf{E}_{k+j,\Gamma}[\![A]\!]$. Then $\varnothing \sqsubseteq_k \Delta$ and $(\Delta, v) \in \mathcal{V}_j[\![A]\!]$.

Theorem 5. (*Type soundness of Environment.*)
If $\Vdash \Sigma$, then for all $F(x) = e : A \to B \in \Sigma$, $k \geqslant 0$ and Δ, we have that $(\Delta, v) \in \mathcal{V}_k[\![A]\!]$ implies $e[v/x] \in \mathsf{E}_{k,\Delta}[\![B]\!]$ and $F\ v \in \mathsf{E}_{k,\Delta}[\![B]\!]$.

Proof. By induction on Σ.

Case DEFBASE: Obvious.

Case DEFFUN: The first claim follows by:

$\Vdash \Sigma$	(1), by assumption
$x : A \vdash e : B$	(2), by assumption
$x : A \vDash e : B$	(3), by Lemma 6 below
$(\Delta, v) \in \mathcal{V}_k[\![A]\!]$	(4), by assumption
$\sigma = [x \mapsto v] \in \mathcal{G}_{k,\Delta}[\![x : A]\!]$	(5), by (4)
$\sigma(e) = e[v/x] \in \mathsf{E}_{k,\Delta}[\![B]\!]$	(6), by (3) and (5)

The second claim follows by:

$(\Delta, v) \in \mathcal{V}_k[\![A]\!]$	(1), by assumption
$e[v/x] \in \mathsf{E}_{k,\Delta}[\![B]\!]$	(2), as before
$j < k, \Theta, w. (\Delta : F\ v \Downarrow_j \Theta : w)$	(3), assume
$\Delta : e[v/x] \Downarrow_{j-1} \Theta : w$	(4), by (3) and APP
$\Delta \sqsubseteq_{j-1} \Theta$	(5), by (2) and (4)
$(\Theta, w) \in \mathcal{V}_{k-j-1}[\![B]\!]$	(6), by (2) and (4)
$\Delta \sqsubseteq_j \Theta$	(6), by (5) and Lemma 23
$(\Theta, w) \in \mathcal{V}_{k-j}[\![B]\!]$	(7), by (6) and Lemma 29
$F\ v \in \mathsf{E}_{k,\Delta}[\![B]\!]$	(8), by (7)

Theorem 6. (*Type soundness.*)
If $\Gamma \vdash e : A$, then $\Gamma \vDash e : A$.

Proof. By induction on $\Gamma \vdash e : A$.

Case VAR:

$\Gamma, x : A \vdash x : A$	(1), by assumption
$\sigma[x \mapsto v] \in \mathcal{G}_{k,\Delta}[\![\Gamma, x : A]\!]$	(2), assume
$\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!]$	(3), by (2)
$(\Delta, v) \in \mathcal{V}_k[\![A]\!]$	(4), by (2)
$v \in \mathsf{E}_{k,\Delta}[\![A]\!]$	(5), by (2) and Lemma 27
$(\sigma[x \mapsto v])(x) \in \mathsf{E}_{k,\Delta}[\![A]\!]$	(6), by (5)
$\Gamma, x : A \vDash x : A$	(7), by (6)

Case $[inl/inr]$:

$$
\begin{array}{ll}
\Gamma \vdash \mathsf{in}_i\ v : A_l + A_r & \text{(1), by assumption} \\
\Gamma \vDash v : A_i & \text{(2), by assumption} \\
\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!] & \text{(3), assume} \\
\sigma(v) \in \mathsf{E}_{k,\Delta}[\![A_i]\!] & \text{(4), by (2)} \\
(\Delta,\ \sigma(v)) \in \mathcal{V}_k[\![A_i]\!] & \text{(5), by (4) and Lemma 28} \\
(\Delta,\ \mathsf{in}_i\ \sigma(v)) \in \mathcal{V}_k[\![A_l + A_r]\!] & \text{(6), by (5) and definition} \\
(\Delta,\ \sigma(\mathsf{in}_i\ v)) \in \mathcal{V}_k[\![A_l + A_r]\!] & \text{(7), by (6)} \\
\sigma(\mathsf{in}_i\ v) \in \mathsf{E}_{k,\Delta}[\![A_l + A_r]\!] & \text{(8), by (7) and Lemma 27} \\
\Gamma \vDash \mathsf{in}_i\ v : A_l + A_r & \text{(9), by (8)}
\end{array}
$$

Case PAIR:

$$
\begin{array}{ll}
\Gamma \vdash (v,\ w) : A \times B & \text{(1), by assumption} \\
\Gamma \vDash v : A & \text{(2), by assumption} \\
\Gamma \vDash w : B & \text{(3), by assumption} \\
\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!] & \text{(4), assume} \\
\sigma(v) \in \mathsf{E}_{k,\Delta}[\![A]\!] & \text{(5), by (2)} \\
(\Delta,\ \sigma(v)) \in \mathcal{V}_k[\![A]\!] & \text{(6), by (5) and Lemma 28} \\
\sigma(w) \in \mathsf{E}_{k,\Delta}[\![B]\!] & \text{(7), by (2)} \\
(\Delta,\ \sigma(w)) \in \mathcal{V}_k[\![B]\!] & \text{(8), by (7) and Lemma 28} \\
(\Delta,\ (\sigma(v),\ \sigma(w))) \in \mathcal{V}_k[\![A \times B]\!] & \text{(9), by (6),(8) and definition} \\
(\Delta,\ \sigma((v,\ w))) \in \mathcal{V}_k[\![A \times B]\!] & \text{(10), by (9)} \\
\sigma((v,\ w)) \in \mathsf{E}_{k,\Delta}[\![A \times B]\!] & \text{(11), by (10) and Lemma 27} \\
\Gamma \vDash (v,\ w) : A \times B & \text{(12), by (11)}
\end{array}
$$

Case FOLD:

$$
\begin{array}{ll}
\Gamma \vdash \mathsf{fold}\ v : \mu\alpha.\ A & \text{(1), by assumption} \\
\Gamma \vDash v : A[\mu\alpha.\ A/\alpha] & \text{(2), by assumption} \\
\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!] & \text{(3), assume} \\
\sigma \in \mathcal{G}_{j,\Delta}[\![\Gamma]\!] & \text{(4), by (3) and Lemma 29} \\
\sigma(v) \in \mathsf{E}_{j,\Delta}[\![A[\mu\alpha.\ A/\alpha]]\!] & \text{(5), by (2)} \\
(\Delta,\ \sigma(v)) \in \mathcal{V}_j[\![A[\mu\alpha.\ A/\alpha]]\!] & \text{(6), by (5) and Lemma 28} \\
(\Delta,\ \mathsf{fold}\ \sigma(v)) \in \mathcal{V}_k[\![\mu\alpha.\ A]\!] & \text{(7), by (6) and definition} \\
(\Delta,\ \sigma(\mathsf{fold}\ v)) \in \mathcal{V}_k[\![\mu\alpha.\ A]\!] & \text{(8), by (7)} \\
\sigma(\mathsf{fold}\ v) \in \mathsf{E}_{k,\Delta}[\![\mu\alpha.\ A]\!] & \text{(9), by (8) and Lemma 27} \\
\Gamma \vDash \mathsf{fold}\ v : \mu\alpha.\ A & \text{(10), by (9)}
\end{array}
$$

Case UNIT:

$$
\begin{array}{ll}
\Gamma \vdash () : 1 & \text{(1), by assumption} \\
\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!] & \text{(2), assume} \\
(\Delta,\ ()) \in \mathcal{V}_k[\![1]\!] & \text{(3), by definition} \\
() \in \mathsf{E}_{k,\Delta}[\![1]\!] & \text{(4), by (3) and Lemma 27} \\
\sigma(()) \in \mathsf{E}_{k,\Delta}[\![A]\!] & \text{(5), by (4)} \\
\Gamma \vDash () : 1 & \text{(6), by (5)}
\end{array}
$$

Case LET:

$\Gamma \vdash \mathsf{let}\ x\ =\ e_1\ \mathsf{in}\ e_2\ :\ B$	(1), by assumption
$\Gamma\ \vDash\ e_1\ :\ A$	(2), by (1)
$\Gamma,\ x\ :\ A\ \vDash\ e_2\ :\ B$	(3), by (1)
$\sigma \in \mathcal{G}_{k+j+1,\Delta}[\![\Gamma]\!]$	(4), assume
$\mathsf{let}\ x\ =\ e_1\ \mathsf{in}\ e_2\ \Downarrow_{k+j+1}\ \Theta\ :\ w$	(5), assume
$\Delta\ :\ e_1\ \Downarrow_k\ \Delta'\ :\ v$	(6), by LET
$\Delta'\ :\ e_2[v/x]\ \Downarrow_j\ \Theta\ :\ w$	(7), by LET
$\sigma(e_1)\ \in \mathsf{E}_{k+j+1,\Delta}[\![A]\!]$	(8), by (2)
$\Delta\ \sqsubseteq_k\ \Delta'$	(9), by (2) and (6)
$(\Delta',\ v)\ \in \mathsf{E}_{k+j+1\ -\ j}[\![A]\!]$	(10), by (2) and (6)
$\sigma \in \mathcal{G}_{k+1,\Delta'}[\![\Gamma]\!]$	(11), by (4) and Lemma 36
$\sigma[x\ \mapsto\ v]\ \in \mathcal{G}_{k+1,\Delta'}[\![\Gamma,\ x\ :\ A]\!]$	(12), by (11) and (10)
$\sigma[x\ \mapsto\ v](e_2)\ \in \mathsf{E}_{k+1,\Delta}[\![B]\!]$	(13), by (3) and (12)
$\sigma(e_2[v/x])\ \in \mathsf{E}_{k+1,\Delta}[\![B]\!]$	(14), by (13)
$\sigma(\mathsf{let}\ x\ =\ e_1\ \mathsf{in}\ e_2)\ \in \mathsf{E}_{k,\Delta}[\![B]\!]$	(15), by (14)
$\Gamma\ \vDash\ \mathsf{let}\ x\ =\ e_1\ \mathsf{in}\ e_2\ :\ B$	(16), by (15)

Case CASE:

$\Gamma \vdash \mathsf{case}\ v\ \{\ \mathsf{inl}\ x \longrightarrow e_l;\ \mathsf{inr}\ y \longrightarrow e_r\ \}\ :\ C$	(1), by assumption
$\Gamma\ \vDash\ v\ :\ A_l\ +\ A_r$	(2), by (1)
$\Gamma,\ x\ :\ A_i\ \vDash\ e_i\ :\ C$	(3), by (1)
$\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!]$	(4), assume
$\Gamma\ :\ \mathsf{case}\ (\mathsf{in}_i\ v)\ \{\ \mathsf{in}_l\ x_l \longrightarrow e_l;\ \mathsf{in}_r\ x_r \longrightarrow e_r\ \}\ \Downarrow_{k_i\ +\ 1}\ \Delta\ :\ w$	(5), assume
$\sigma(v)\ \in \mathsf{E}_{k,\Delta}[\![A_l\ +\ A_r]\!]$	(6), by (2)
$(\Delta,\ \sigma(v))\ \in \mathcal{V}_k[\![A_l\ +\ A_r]\!]$	(7), by (6) and Lemma 28
$\sigma(v)\ =\ \mathsf{in}_i\ \sigma(w)$	(8), by (7)
$(\Delta,\ \sigma(w))\ \in \mathcal{V}_k[\![A_i]\!]$	(9), by (7)
$\sigma'\ =\ \sigma[x \mapsto \sigma(w)]\ \in \mathcal{G}_{k,\Delta}[\![\Gamma,\ x\ :\ A_i]\!]$	(10), by (4), (9)
$\sigma'(e_i)\ \in \mathsf{E}_{k,\Delta}[\![C]\!]$	(11), by (3),(9)
$\sigma(\mathsf{case}\ v\ \{\ \mathsf{inl}\ x \longrightarrow e_l;\ \mathsf{inr}\ y \longrightarrow e_r\ \})\ \in \mathsf{E}_{k,\Delta}[\![C]\!]$	(12), by (10)
$\Gamma\ \vDash\ \mathsf{case}\ v\ \{\ \mathsf{inl}\ x \longrightarrow e_l;\ \mathsf{inr}\ y \longrightarrow e_r\ \}\ :\ C$	(13), by (12)

Case SPLIT:

$\Gamma \vdash \mathsf{split}\ v\ \{\ (x,\ y) \longrightarrow e\ \}\ :\ C$	(1), by assumption
$\Gamma\ \vDash\ v\ :\ A\ \times B$	(2), by (1)
$\Gamma,\ x\ :\ A,\ y\ :\ B\ \vDash\ e\ :\ C$	(3), by (1)
$\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!]$	(4), assume
$\Gamma\ :\ \mathsf{split}\ (v_1,\ v_2)\ \{\ (x,\ y) \longrightarrow e\ \}\ \Downarrow_{k'+1}\ \Delta\ :\ w$	(5), assume
$\sigma(v)\ \in \mathsf{E}_{k,\Delta}[\![A\ \times B]\!]$	(6), by (2)
$(\Delta,\ \sigma(v))\ \in \mathcal{V}_k[\![A\ \times B]\!]$	(7), by (6) and Lemma 28
$\sigma(v)\ =\ (\sigma(v_1),\ \sigma(v_2))$	(8), by (7)
$(\Delta,\ \sigma(v_1))\ \in \mathcal{V}_k[\![A]\!]$	(9), by (7)
$(\Delta,\ \sigma(v_2))\ \in \mathcal{V}_k[\![B]\!]$	(10), by (7)
$\sigma'\ =\ \sigma[x \mapsto \sigma(v_1),\ y \mapsto \sigma(v_2)]\ \in \mathcal{G}_{k,\Delta}[\![\Gamma,\ x\ :\ A,\ y\ :\ B]\!]$	(11), by (4), (9), (10)
$\sigma'(e)\ \in \mathsf{E}_{k,\Delta}[\![C]\!]$	(12), by (3),(11)
$\sigma(\mathsf{split}\ v\ \{\ (x,\ y) \longrightarrow e\ \})\ \in \mathsf{E}_{k,\Delta}[\![C]\!]$	(13), by (12)
$\Gamma\ \vDash\ \mathsf{split}\ v\ \{\ (x,\ y) \longrightarrow e\ \}\ :\ C$	(14), by (13)

Case UNFOLD:

$$\begin{array}{ll}
\Gamma \vdash \mathsf{unfold}\ v\ :\ A[\mu\alpha.\ A/\alpha] & (1), \text{by assumption} \\
\Gamma\ \vDash\ v\ :\ \mu\alpha.\ A & (2), \text{by } (1) \\
\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!] & (3), \text{assume} \\
\Gamma\ :\ \mathsf{unfold}\ (\mathsf{fold}\ v)\ \Downarrow_1\ \Gamma\ :\ v & (4), \text{assume} \\
\sigma(v)\ \in \mathsf{E}_{k,\Delta}[\![\mu\alpha.\ A]\!] & (5), \text{by } (2) \\
(\Delta,\ \sigma(v))\ \in \mathcal{V}_k[\![\mu\alpha.\ A]\!] & (6), \text{by (4) and Lemma 28} \\
\sigma(v)\ =\ \mathsf{fold}\ w & (7), \text{by } (6) \\
(\Delta,\ \sigma(w))\ \in \mathcal{V}_{k-1}[\![A[\mu\alpha.\ A/\alpha]]\!] & (8), \text{by } (6) \\
\mathsf{unfold}\ (\mathsf{fold}\ \sigma(v))\ \in \mathsf{E}_{k-1}[\![A[\mu\alpha.\ A/\alpha]]\!] & (9), \text{by } (8) \\
\Gamma\ \vDash\ \mathsf{unfold}\ v\ :\ A[\mu\alpha.\ A/\alpha] & (10), \text{by } (9)
\end{array}$$

Case APP:

$$\begin{array}{ll}
\Gamma \vdash F\ v\ :\ B & (1), \text{by assumption} \\
F\ :\ A \longrightarrow B\ \in \Sigma & (2), \text{by } (1) \\
\Gamma\ \vDash\ v\ :\ A & (3), \text{by } (1) \\
\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!] & (4), \text{assume} \\
\sigma(v)\ \in \mathsf{E}_{k,\Delta}[\![A]\!] & (5), \text{by } (2) \\
(\Delta,\ \sigma(v))\ \in \mathcal{V}_k[\![A]\!] & (6), \text{by (5) and Lemma 28} \\
F\ \sigma(v)\ \in \mathsf{E}_{k,\Delta}[\![B]\!] & (7), \text{by (6) and Lemma 5} \\
\Gamma\ \vDash\ F\ v\ :\ B & (8), \text{by } (7)
\end{array}$$

Case LAZY:

$$\begin{array}{ll}
\Gamma \vdash \mathsf{lazy}_F\ v\ :\ A\ \nrightarrow_F\ B & (1), \text{by assumption} \\
\Gamma\ \vDash\ v\ :\ A & (2), \text{by assumption} \\
\forall k\ \geqslant\ 0,\ \Delta,\ w.\ (\Delta,\ w)\ \in \mathcal{V}_k[\![A]\!]\ \implies\ F\ w\ \in \mathsf{E}_{k,\Delta}[\![B]\!] & (3), \text{by Lemma 5} \\
\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!] & (4), \text{assume} \\
\sigma(v)\ \in \mathsf{E}_{k,\Delta}[\![A]\!] & (5), \text{by } (2) \\
(\Delta,\ \sigma(v))\ \in \mathcal{V}_k[\![A]\!] & (6), \text{by (5) and Lemma 28} \\
(\Theta,\ \sigma(v))\ \in \mathcal{V}_{k\ -\ j}[\![A]\!] & (7), \text{by (6) and Lemma 36} \\
F\ w\ \in \mathsf{E}_{k\ -\ j,\ \Theta}[\![B]\!] & (8), \text{by (7) and (3)} \\
((\Delta,\ z\ \mapsto \mathsf{lazy}_F\ \sigma(v)),\ z)\ \in \mathcal{V}_k[\![A\ \nrightarrow_F\ B]\!] & (9), \text{by (5) and (8)} \\
\Delta\ :\ \sigma(\mathsf{lazy}_F\ v)\ \Downarrow_1\ (\Delta,\ z\ \mapsto \mathsf{lazy}_F\ \sigma(v))\ :\ z & (10), \text{by } \textsc{lazy} \\
\Delta\ \sqsubseteq_1\ (\Delta,\ z\ \mapsto \mathsf{lazy}_F\ \sigma(v)) & (11), \text{by } \textsc{extend} \\
\Gamma\ \vDash\ \mathsf{lazy}_F\ v\ :\ A\ \nrightarrow_F\ B & (12), \text{by } (9),(10),(11)
\end{array}$$

Case MEMO:

$$\begin{array}{ll}
\Gamma \vdash \mathsf{memo}\ v\ :\ A\ \nrightarrow_F\ B & (1), \text{by assumption} \\
\Gamma\ \vDash\ v\ :\ B & (2), \text{by assumption} \\
\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!] & (3), \text{assume} \\
\sigma(v)\ \in \mathsf{E}_{k,\Delta}[\![B]\!] & (4), \text{by } (2) \\
(\Delta,\ \sigma(v))\ \in \mathcal{V}_k[\![B]\!] & (5), \text{by (4) and Lemma 28} \\
((\Delta,\ z\ \mapsto \mathsf{memo}\ \sigma(v)),\ z)\ \in \mathcal{V}_k[\![A\ \nrightarrow_F\ B]\!] & (6), \text{by } (5) \\
\Delta\ :\ \sigma(\mathsf{memo}\ v)\ \Downarrow_1\ (\Delta,\ z\ \mapsto \mathsf{memo}\ \sigma(v))\ :\ z & (7), \text{by } \textsc{lazy} \\
\Delta\ \sqsubseteq_1\ (\Delta,\ z\ \mapsto \mathsf{memo}\ \sigma(v)) & (8), \text{by } \textsc{extend} \\
\Gamma\ \vDash\ \mathsf{memo}\ v\ :\ A\ \nrightarrow_F\ B & (9), \text{by } (8)
\end{array}$$

Case STEP:

$\Gamma \vdash \mathsf{step}\ v : B$ (1), by assumption
$\Gamma \vDash v : A \nrightarrow_F B$ (2), by assumption
$\sigma \in \mathcal{G}_{k,\Delta}[\![\Gamma]\!]$ (3), assume
$\sigma(v) \in \mathsf{E}_{k,\Delta}[\![A \nrightarrow_F B]\!]$ (4), by (2)
$(\Delta, \sigma(v)) \in \mathcal{V}_k[\![A \nrightarrow_F B]\!]$ (5), by (4) and Lemma 28

Case $\sigma(v) = z$ and $\Delta = \Delta', z \mapsto \mathsf{lazy}_F\ v'$:

$(\Delta', v') \in \mathcal{V}_k[\![A]\!]$ (6), by unfolding (5)
$\forall j \leqslant k, \Theta.\ \Delta' \sqsubseteq_j \Theta \Rightarrow F\ v' \in \mathsf{E}_{k-j,\Theta}[\![B]\!]$ (7), by unfolding (5)
$(\Delta', z \mapsto \mathsf{lazy}_F\ v') : \mathsf{step}\ z \Downarrow_{k+1} (\Theta, z \mapsto \mathsf{memo}\ w) : w$ (8), by EVAL
$\Delta' : F\ v' \Downarrow_k \Theta : w$ (9), by (8)
$\Delta' \sqsubseteq_0 \Delta'$ (10), by REFL
$F\ v' \in \mathsf{E}_{k,\Delta'}[\![B]\!]$ (11), by (7) and (10)
$\Delta' \sqsubseteq_k \Theta$ (12), by (9) and (11)
$(\Theta, w) \in \mathcal{V}_0[\![B]\!]$ (13), by (9) and (11)
$\Delta \sqsubseteq_{k+1} (\Theta, z \mapsto \mathsf{memo}\ w)$ (14), by (12) and EXTEND
$(\Theta, z \mapsto \mathsf{memo}\ w, w) \in \mathcal{V}_0[\![B]\!]$ (15), by (13) and Lemma 30
$\mathsf{step}\ v \in \mathsf{E}_k[\![B]\!]$ (16), by (14) and (15)
$\Gamma \vDash \mathsf{step}\ v : B$ (17), by (16)

Case $\sigma(v) = z \mapsto \mathsf{memo}\ v' \in \Delta$:

$(\Delta, v') \in \mathcal{V}_k[\![B]\!]$ (6), by unfolding (5)
$v' \in \mathsf{E}_{k,\Delta}[\![B]\!]$ (7), by (6) and Lemma 27
$\Gamma : \mathsf{step}\ z \Downarrow_1 \Gamma : v'$ (8), by RECALL
$\Gamma \sqsubseteq_1 \Gamma$ (9), by REFL
$\mathsf{step}\ z \in \mathsf{E}_{k,\Delta}[\![B]\!]$ (10), by (7)
$\mathsf{sigma}(\mathsf{step}\ v) \in \mathsf{E}_{k,\Delta}[\![B]\!]$ (11), by (9) and (10)
$\Gamma \vDash \mathsf{step}\ v : B$ (12), by (11)

C.1.5 Referential Transparency.

Lemma 20. (*Evaluation in store extension*)
If $(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k (\Delta, x \mapsto \mathsf{memo}\ w)$, then there are stores Γ', Δ' such that $\Gamma' : F\ v \Downarrow_j \Delta' : w$ and $j \leqslant k$ and $(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Gamma', x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Delta', x \mapsto \mathsf{memo}\ w) \sqsubseteq (\Delta, x \mapsto \mathsf{memo}\ w)$.

Proof. By induction on $(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k (\Delta, x \mapsto \mathsf{memo}\ w)$.

Case REFL: Impossible.

Case EXTEND: Impossible.

Case TRANS: Let $(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k \Delta \sqsubseteq_j (\Theta, x \mapsto \mathsf{memo}\ w)$. If $\Delta = \Delta_1, x \mapsto \mathsf{lazy}_F\ v$:

$(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k (\Delta_1, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_j (\Theta, x \mapsto \mathsf{memo}\ w)$ (1), by assumptio
$(\Delta_1, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_j (\Theta, x \mapsto \mathsf{memo}\ w)$ (2), restrict (1)
$\Gamma' : F\ v \Downarrow_j \Delta' : w$ and $j \leqslant k$ (3), by the induct
$(\Delta_1, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Gamma', x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Delta', x \mapsto \mathsf{memo}\ w) \sqsubseteq (\Theta, x \mapsto \mathsf{memo}\ w)$ (4), by inductive
$(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Gamma', x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Delta', x \mapsto \mathsf{memo}\ w) \sqsubseteq (\Theta, x \mapsto \mathsf{memo}\ w)$ (5), by (4) and tra

If $\Delta = \Delta_1, x \mapsto \mathsf{memo}\ w$:

$(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k (\Delta_1, x \mapsto \mathsf{memo}\ w) \sqsubseteq_j (\Theta, x \mapsto \mathsf{memo}\ w)$ (1), by assumptio
$(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k (\Delta_1, x \mapsto \mathsf{memo}\ w)$ (2), restrict (1)
$\Gamma' : F\ v \Downarrow_j \Delta' : w$ and $j \leqslant k$ (3), by the induct
$(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Gamma', x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Delta', x \mapsto \mathsf{memo}\ w) \sqsubseteq (\Delta_1, x \mapsto \mathsf{memo}\ w)$ (4), by inductive l
$(\Gamma, x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Gamma', x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Delta', x \mapsto \mathsf{memo}\ w) \sqsubseteq (\Theta, x \mapsto \mathsf{memo}\ w)$ (5), by (4) and tra

Case EVAL: The statement follow directly from the EVAL rule with $\Gamma = \Gamma'$ and $\Delta = \Delta'$.

Theorem 7. (*Lazy constructors are referentially transparent.*)
If $\Gamma : e \Downarrow_k \Delta : v$ and $\Gamma \sqsubseteq_j \Gamma'$, then $\Gamma' : e \Downarrow_l \Delta' : v$ with $\Delta \sqsubseteq_{l'} \Delta'$.

Proof. By induction on k and then case-split on the evaluation. If $k = 0$, then only VALUE applies. Then the claim follows directly. If $k > 0$:

Case UNFOLD: The claim follows directly.

Case LET: By the inductive hypothesis.

Case APP: By the inductive hypothesis.

Case SPLIT: By the inductive hypothesis.

Case CASE: By the inductive hypothesis.

Case LAZY:

$\Gamma : lv \Downarrow_1 (\Gamma, z \mapsto lv) : z$	(1), by LAZY
$\Gamma \sqsubseteq_j \Gamma'$	(2), by assumption
$\Gamma' : lv \Downarrow_1 (\Gamma', z \mapsto lv) : z$	(3), by LAZY
$(\Gamma, z \mapsto lv) \sqsubseteq_j (\Gamma', z \mapsto lv)$	(4), by (2) and Lemma 24

Case RECALL:

$\Gamma : \mathsf{step}\ z \Downarrow_1 \Gamma : w$	(1), by RECALL
$z \mapsto \mathsf{memo}\ w \in \Gamma$	(2), by RECALL
$\Gamma \sqsubseteq_j \Gamma'$	(3), by assumption
$z \mapsto \mathsf{memo}\ w \in \Gamma'$	(4), by (3) and Lemma 25
$\Gamma' : \mathsf{step}\ z \Downarrow_1 \Gamma' : w$	(5), by (4) and RECALL

Case STEP: We have $\Gamma = \Gamma_1, z \mapsto \mathsf{lazy}_F\ v$ and $\Delta = \Delta_1, z \mapsto \mathsf{memo}\ w$. If $\Gamma' = \Gamma_1', z \mapsto \mathsf{lazy}_F\ v$:

$\Gamma : \mathsf{step}\ z \Downarrow_{k+1} \Delta : w$	(1), by STEP
$\Gamma_1 : F\ v \Downarrow_k \Delta_1 : w$	(2), by STEP
$\Gamma_1' : F\ v \Downarrow_l \Delta_1' : w$	(3), by the inductive hypothesis
$\Delta_1 \sqsubseteq \Delta_1'$	(4), by the inductive hypothesis
$\Gamma_1', z \mapsto \mathsf{lazy}_F\ v : \mathsf{step}\ z \Downarrow_1 \Delta_1', z \mapsto \mathsf{memo}\ w : w$	(5), by (3) and STEP
$\Delta \sqsubseteq \Delta_1', z \mapsto \mathsf{memo}\ w$	(6), by (4) and Lemma 24

Else $\Gamma' = \Gamma_1', z \mapsto \mathsf{memo}\ w$:

$\Gamma : \mathsf{step}\ z \Downarrow_{k+1} \Delta : w$	(1), by STEP
$\Gamma_1 : F\ v \Downarrow_k \Delta_1 : w$	(2), by STEP
$\Gamma_1'' : F\ v \Downarrow_j \Delta_1' : w'$	(3), by Lemma 37
$\Gamma \sqsubseteq (\Gamma_1'', z \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Delta_1', z \mapsto \mathsf{memo}\ w) \sqsubseteq \Gamma'$	(4), by Lemma 37
$\Gamma_1 \sqsubseteq \Gamma_1''$	(5), by Lemma 26
$\Gamma_1'' : F\ v \Downarrow_j \Delta_1' : w$	(6), by the inductive hypothesis
$\Delta_1 \sqsubseteq \Delta_1'$	(7), by the inductive hypothesis
$\Delta \sqsubseteq (\Delta_1', z \mapsto \mathsf{memo}\ w)$	(8), by Lemma 24
$\Delta \sqsubseteq \Gamma'$	(9), by transitivity

C.2 Lazy Constructors

C.2.1 Typing rules. We extend the typing rules as follows:

$$\frac{F : A \to B + (A \twoheadrightarrow_F B) \in \Sigma \qquad \Gamma \vdash v : A}{\Gamma \vdash \mathsf{lazy}_F\ v : A \twoheadrightarrow_F B}\ \textsc{lazy}$$

$$\frac{F : A \to B + (A \nrightarrow_F B) \in \Sigma \qquad \Gamma \vdash v : B}{\Gamma \vdash \text{memo } v : A \nrightarrow_F B}\ \textsc{memo} \qquad \frac{\Gamma \vdash v : A \nrightarrow_F B}{\Gamma \vdash \text{eval } v : B}\ \textsc{eval}$$

C.2.2 Store extension. We extend the store relation with:

$$\frac{y \mapsto \text{indirect } z \in \Gamma}{\Gamma, x \mapsto \text{indirect } y \sqsubseteq_1 \Gamma, x \mapsto \text{indirect } z}\ \textsc{cuti} \qquad \frac{y \mapsto \text{memo } v \in \Gamma}{\Gamma, x \mapsto \text{indirect } y \sqsubseteq_1 \Gamma, x \mapsto \text{memo } v}\ \textsc{cutm}$$

$$\frac{y \mapsto \text{indirect } z \in \Gamma}{\Gamma, x \mapsto \text{indirect } z \sqsubseteq_1 \Gamma, x \mapsto \text{indirect } y}\ \textsc{indi} \qquad \frac{y \mapsto \text{memo } v \in \Gamma}{\Gamma, x \mapsto \text{memo } v \sqsubseteq_1 \Gamma, x \mapsto \text{indirect } y}\ \textsc{indm}$$

Lemma 21. (*Store extension is reflexive.*)
For all stores Γ, $\Gamma \sqsubseteq_0 \Gamma$.

Lemma 22. (*Store extension is transitive.*)
If $\Gamma \sqsubseteq_k \Delta$ and $\Delta \sqsubseteq_j \Theta$, then $\Gamma \sqsubseteq_{k+j} \Theta$.

Lemma 23. (*Store extension may take more steps*)
If $\Gamma \sqsubseteq_k \Delta$ then $\Gamma \sqsubseteq_{k+1} \Delta$.

Lemma 24. (*Growing the store does not change store extension.*)
If $\Gamma \sqsubseteq_k \Delta$, then $(\Gamma, x \mapsto lv) \sqsubseteq_k (\Delta, x \mapsto lv)$ for any $x \notin \Delta$.

Proof. By induction on $\Gamma \sqsubseteq_k \Delta$.

Case cutm: Use the cutm rule with the extra binding.

Case cuti: Use the cuti rule with the extra binding.

Case indm: Use the indm rule with the extra binding.

Case indi: Use the indi rule with the extra binding.

Lemma 25. (*Memos stay memos in the store*)
If $\Gamma \sqsubseteq_k \Delta$ and $x \mapsto \text{memo } v \in \Gamma$, then $x \mapsto \text{memo } v \in \Delta$.

Proof. By induction on $\Gamma \sqsubseteq_k \Delta$ using 9.

Lemma 26. (*Removing lazies from the store does not change store extension.*)
If $(\Gamma, x \mapsto \text{lazy}_F\ v) \sqsubseteq_k (\Delta, x \mapsto \text{lazy}_F\ v)$, then $\Gamma \sqsubseteq_k \Delta$.

Proof. By induction on $\Gamma \sqsubseteq_k \Delta$.

Case cutm: Use the cutm rule without the extra binding.

Case cuti: Use the cuti rule without the extra binding.

Case indm: Use the indm rule without the extra binding.

Case indi: Use the indi rule without the extra binding.

C.2.3 Logical relation. We extend the logical relation with:

$$\begin{aligned}\mathcal{V}_k[\![A \nrightarrow_F B]\!] := \{ ((\Delta, z \mapsto \text{lazy}_F\ v), z) \mid (\Delta, v) \in \mathcal{V}_k[\![A]\!], \forall j < k, \Theta. \Delta \sqsubseteq \Theta \Rightarrow \\ F\ v \in \mathsf{E}_{j,\Theta}[\![B + (A \nrightarrow_F B)]\!] \} \\ \cup \{ ((\Delta, z \mapsto \text{memo } v), z) \mid (\Delta, v) \in \mathcal{V}_k[\![B]\!] \} \\ \cup \{ ((\Delta, z \mapsto \text{indirect } v), z) \mid \forall j < k. (\Delta, v) \in \mathcal{V}_j[\![A \nrightarrow_F B]\!] \}\end{aligned}$$

We maintain the properties from before:

Lemma 27. (*Values are sound expressions.*)
If $(\Delta, v) \in \mathcal{V}_k[\![A]\!]$, then $v \in \mathsf{E}_{k,\Delta}[\![A]\!]$.

Lemma 28. (*On values, the expression denotation is the value denotation.*)
If $v \in \mathsf{E}_{k,\Delta}[\![A]\!]$, then $(\Delta,\ v) \in \mathcal{V}_k[\![A]\!]$.

Lemma 29. (*Downward closure*)
If $(\Delta,\ v) \in \mathcal{V}_k[\![A]\!]$, then $(\Delta,\ v) \in \mathcal{V}_j[\![A]\!]$ for all $j \leqslant k$.

Proof. By induction on k. If $k = 0$, then obvious. Else: induction on A.

Case $A = A' \not\twoheadrightarrow_F B$. Let $(\Delta,\ z) \in \mathcal{V}_k[\![A]\!]$ with $z \mapsto \mathsf{memo}\ v \in \Delta$. Then the claim follows directly from the inner inductive hypothesis.

Case $A = A' \not\twoheadrightarrow_F B$. Let $(\Delta,\ z) \in \mathcal{V}_k[\![A]\!]$ with $z \mapsto \mathsf{indirect}\ v \in \Delta$. Then the claim follows directly from the outer inductive hypothesis.

Case $A = A' \not\twoheadrightarrow_F B$. Let $(\Delta,\ z) \in \mathcal{V}_k[\![A]\!]$ with $z \mapsto \mathsf{lazy}_F\ v \in \Delta$.

$(\Delta,\ z) \in \mathcal{V}_k[\![\, A' \not\twoheadrightarrow_F B \,]\!]$	(1), by assumption
$z \mapsto \mathsf{lazy}_F\ v \in \Delta$	(2), by assumption
$(\Delta,\ v) \in \mathcal{V}_k[\![\, A \,]\!]$	(3), by (1) and (2)
$\forall j \leqslant k,\ \Theta.\ \Delta \sqsubseteq_j \Theta \Rightarrow F\ v \in \mathsf{E}_{k-j,\,\Theta}[\![B + (A' \not\twoheadrightarrow_F B)]\!]$	(4), by (1) and (2)
$j' \leqslant k$	(5), by assumption
$(\Delta,\ v) \in \mathcal{V}_{j'}[\![\, A \,]\!]$	(6), apply inner inductive hypothesis to (3) and
$\forall j \leqslant j',\ \Theta.\ \Delta \sqsubseteq_j \Theta \Rightarrow F\ v \in \mathsf{E}_{k-j,\,\Theta}[\![B + (A' \not\twoheadrightarrow_F B)]\!]$	(7), by (4) and (5)
$(\Delta,\ z) \in \mathcal{V}_{j'}[\![\, A' \not\twoheadrightarrow_F B \,]\!]$	(8), by (6) and (7)

Lemma 30. (*Growing the Store preserves types.*)
If $(\Delta,\ v) \in \mathcal{V}_k[\![\, A \,]\!]$, then $((\Delta,\ x \mapsto lv),\ v) \in \mathcal{V}_k[\![\, A \,]\!]$ for any $x \notin \Delta$.

Proof. By induction on k and then on A.

Case $A = A' \not\twoheadrightarrow_F B$: Follows from the inner inductive hypothesis and by using the EXTEND rule to obtain $\Delta \sqsubseteq (\Delta,\ x \mapsto lv)$.

Lemma 31. (*Evaluation preserves types.*)
If $((\Gamma, x \mapsto \mathsf{lazy}_F\ v'),\ v) \in \mathcal{V}_k[\![\, A \,]\!]$ and $\Gamma : F\ v' \Downarrow_j \Delta : w$ for $j < k$ and $\forall v.\ (\Gamma,\ v) \in \mathcal{V}_k[\![A]\!] \Rightarrow (\Delta,\ v) \in \mathcal{V}$
then $((\Delta, x \mapsto \mathsf{memo}\ w),\ v) \in \mathcal{V}_{k-j}[\![\, A \,]\!]$.

Proof. By induction on k and then on A.

Case $A = A' \not\twoheadrightarrow_F B$: Let $v = z$. Case $x = z$:

$((\Gamma, x \mapsto \mathsf{lazy}_F\ v'),\ x) \in \mathcal{V}_k[\![\, A' \not\twoheadrightarrow_F B \,]\!]$	(1), by assumpti
$(\Gamma,\ v') \in \mathcal{V}_k[\![\, A' \,]\!]$	(2), by (1)
$\forall j \leqslant k,\ \Theta.\ \Gamma \sqsubseteq_j \Theta \Rightarrow F\ v' \in \mathsf{E}_{k-j,\,\Theta}[\![B + (A' \not\twoheadrightarrow_F B)]\!]$	(3), by (1)
$F\ v' \in \mathsf{E}_{k,\,\Gamma}[\![B + (A' \not\twoheadrightarrow_F B)]\!]$	(4), instantiate (
$\forall j < k.\ \forall \Delta,\ v.\ (\Gamma : F\ v' \Downarrow_j \Delta : w) \Rightarrow \Gamma \sqsubseteq_j \Delta$ and $(\Delta,\ w) \in \mathcal{V}_{k-j}[\![\, B + (A' \not\twoheadrightarrow_F B) \,]\!]$	(5), unroll (4)
$(\Delta,\ w) \in \mathcal{V}_{k-j}[\![\, B + (A' \not\twoheadrightarrow_F B) \,]\!]$	(6), simplify (5)
$(\Delta,\ w') \in \mathcal{V}_{k-j}[\![\, B \,]\!]$	(7), case (6)
$((\Delta,\ x \mapsto \mathsf{memo}\ w'),\ x) \in \mathcal{V}_{k-j}[\![\, A' \not\twoheadrightarrow_F B \,]\!]$	(8), by definitio
$(\Delta,\ w') \in \mathcal{V}_{k-j}[\![\, A' \not\twoheadrightarrow_F B \,]\!]$	(7), case (6)
$((\Delta,\ x \mapsto \mathsf{indirect}\ w'),\ x) \in \mathcal{V}_{k-j}[\![\, A' \not\twoheadrightarrow_F B \,]\!]$	(8), by definitio

Case $x \neq z$:

$((\Gamma, x \mapsto \mathsf{lazy}_F\ v', z \mapsto \mathsf{lazy}_F\ v''),\ z) \in \mathcal{V}_k[\![A' \twoheadrightarrow_F B]\!]$ (1), by assumption
$((\Gamma, x \mapsto \mathsf{lazy}_F\ v'),\ v'') \in \mathcal{V}_k[\![A']\!]$ (2), by (1)
$((\Delta, x \mapsto \mathsf{lazy}_F\ v'),\ v'') \in \mathcal{V}_{k-j}[\![A']\!]$ (3), by (2) and inductive hypo
$\forall l \leqslant k, \Theta.\ (\Gamma, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_l \Theta \Rightarrow F\ v \in \mathsf{E}_{k-l, \Theta}[\![B + (A' \twoheadrightarrow_F B)]\!]$ (4), by (1)
$\Gamma \sqsubseteq_j \Delta$ (5), by Lemma 13
$(\Gamma, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_j (\Delta, x \mapsto \mathsf{lazy}_F\ v')$ (6), by Lemma 24
$\forall l \leqslant k, \Theta.\ (\Delta, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_l \Theta \Rightarrow (\Gamma, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_{j+l} \Theta$ (7), by Lemma 22
$\forall l \leqslant k, \Theta.\ (\Delta, x \mapsto \mathsf{lazy}_F\ v') \sqsubseteq_l \Theta \Rightarrow F\ v \in \mathsf{E}_{k-j-l, \Theta}[\![B + (A' \twoheadrightarrow_F B)]\!]$ (8), by (4) and (7)
$((\Delta, x \mapsto \mathsf{lazy}_F\ v', z \mapsto \mathsf{lazy}_F\ v''),\ z) \in \mathcal{V}_{k-j}[\![A' \twoheadrightarrow_F B]\!]$ (4), by definition with (3) and

Lemma 32. (*Short-cutting to memo preserves value denotation.*)
$((\Gamma,\ x \mapsto \mathsf{indirect}\ y),\ v) \in \mathcal{V}_k[\![A]\!]$ and $y \mapsto \mathsf{memo}\ w \in \Gamma$ $((\Gamma,\ x \mapsto \mathsf{memo}\ w),\ v) \in \mathcal{V}_{k-1}[\![A]\!]$
Proof. By induction on k. If $k = 0$, then obvious. Else: induction on A.

Case $A = 1$: Obvious

Case $A = A' + B$ or $A = A' \times B$: Follows directly from the inner inductive hypothesis.

Case $A = \mu\alpha.\ A'$: Follows directly from the outer inductive hypothesis.

Case $A = A' \twoheadrightarrow_F B, z \neq x$. If $z \mapsto \mathsf{memo}\ w' \in (\Gamma,\ x \mapsto \mathsf{indirect}\ y)$, then the claim follows from the inner inductive hypothesis. If $z \mapsto \mathsf{indirect}\ w' \in (\Gamma,\ x \mapsto \mathsf{indirect}\ y)$, then the claim follows from the outer inductive hypothesis. If $z \mapsto \mathsf{lazy}_F\ w' \in (\Gamma,\ x \mapsto \mathsf{indirect}\ y)$, then the claim follows from the inner inductive hypothesis and applying CUTM to $\Delta \sqsubseteq \Theta$.

Case $A = A' \twoheadrightarrow_F B, z = x$.
$((\Gamma,\ x \mapsto \mathsf{indirect}\ y),\ x) \in \mathcal{V}_k[\![A]\!]$ (1), by assumption
$(\Gamma,\ y) \in \mathcal{V}_{k-1}[\![A' \twoheadrightarrow_F B]\!]$ (2), by (1)
$\Gamma = \Gamma_1,\ y \mapsto \mathsf{memo}\ w$ (3), by assumption
$(\Gamma_1,\ w) \in \mathcal{V}_{k-1}[\![B]\!]$ (4), by (2) and (3)
$(\Gamma,\ w) \in \mathcal{V}_{k-1}[\![B]\!]$ (5), by Lemma 30
$((\Gamma,\ x \mapsto \mathsf{memo}\ w),\ x) \in \mathcal{V}_{k-1}[\![A' \twoheadrightarrow_F B]\!]$ (6), by (5)

Lemma 33. (*Indirecting to memo preserves value denotation.*)
$((\Gamma,\ x \mapsto \mathsf{memo}\ w),\ v) \in \mathcal{V}_k[\![A]\!]$ and $y \mapsto \mathsf{memo}\ w \in \Gamma$ $((\Gamma,\ x \mapsto \mathsf{indirect}\ y),\ v) \in \mathcal{V}_{k-1}[\![A]\!]$
Proof. By induction on k. If $k = 0$, then obvious. Else: induction on A.
Case $A = 1$: Obvious

Case $A = A' + B$ or $A = A' \times B$: Follows directly from the inner inductive hypothesis.

Case $A = \mu\alpha.\ A'$: Follows directly from the outer inductive hypothesis.

Case $A = A' \twoheadrightarrow_F B, z \neq x$. If $z \mapsto \mathsf{memo}\ w' \in (\Gamma,\ x \mapsto \mathsf{memo}\ w)$, then the claim follows from the inner inductive hypothesis. If $z \mapsto \mathsf{indirect}\ w' \in (\Gamma,\ x \mapsto \mathsf{memo}\ w)$, then the claim follows from the outer inductive hypothesis. If $z \mapsto \mathsf{lazy}_F\ w' \in (\Gamma,\ x \mapsto \mathsf{memo}\ w)$, then the claim follows from the inner inductive hypothesis and applying INDM to $\Delta \sqsubseteq \Theta$.

Case $A = A' \twoheadrightarrow_F B, z = x$.
$((\Gamma,\ x \mapsto \mathsf{memo}\ w),\ x) \in \mathcal{V}_k[\![A]\!]$ (1), by assumption
$(\Gamma,\ w) \in \mathcal{V}_k[\![B]\!]$ (2), by (1)
$\Gamma = \Gamma_1,\ y \mapsto \mathsf{memo}\ w$ (3), by assumption
$(\Gamma,\ y) \in \mathcal{V}_k[\![A' \twoheadrightarrow_F B]\!]$ (4), by (2) and (3)
$(\Gamma,\ y) \in \mathcal{V}_j[\![A' \twoheadrightarrow_F B]\!]$ (5), for $j < k + 1$, by Lemma 29
$((\Gamma,\ x \mapsto \mathsf{indirect}\ y),\ x) \in \mathcal{V}_{k+1}[\![A' \twoheadrightarrow_F B]\!]$ (6), by (5)
$((\Gamma,\ x \mapsto \mathsf{indirect}\ y),\ x) \in \mathcal{V}_{k-1}[\![A' \twoheadrightarrow_F B]\!]$ (7), by (6) and Lemma 29

Lemma 34. (*Short-cutting to indirect preserves value denotation.*)
$((\Gamma,\ x\ \mapsto \mathsf{indirect}\ y),\ v)\ \in \mathcal{V}_k[\![\ A\ ]\!]$ and $y\ \mapsto \mathsf{indirect}\ w\ \in \Gamma$ $((\Gamma,\ x\ \mapsto \mathsf{indirect}\ w),\ v)\ \in \mathcal{V}_{k\ -\ 1}[\![\ A\ ]\!]$
Proof. By induction on k. If $k\ =\ 0$, then obvious. Else: induction on A.

Case $A\ =\ 1$: Obvious

Case $A\ =\ A'\ +\ B$ or $A\ =\ A'\ \times B$: Follows directly from the inner inductive hypothesis.

Case $A\ =\ \mu\alpha.\ A'$: Follows directly from the outer inductive hypothesis.

Case $A\ =\ A'\ \nrightarrow_F\ B, z\ \neq x$. If $z\ \mapsto \mathsf{memo}\ w'\ \in (\Gamma,\ x\ \mapsto \mathsf{indirect}\ y)$, then the claim follows from the inner inductive hypothesis. If $z\ \mapsto \mathsf{indirect}\ w'\ \in (\Gamma,\ x\ \mapsto \mathsf{indirect}\ y)$, then the claim follows from the outer inductive hypothesis. If $z\ \mapsto \mathsf{lazy}_F\ w'\ \in (\Gamma,\ x\ \mapsto \mathsf{indirect}\ y)$, then the claim follows from the inner inductive hypothesis and applying $\textsc{cuti}$ to $\Delta\ \sqsubseteq \Theta$.

Case $A\ =\ A'\ \nrightarrow_F\ B, z\ =\ x$.

$((\Gamma,\ x\ \mapsto \mathsf{indirect}\ y),\ x)\ \in \mathcal{V}_k[\![\ A\ ]\!]$	(1), by assumption
$(\Gamma,\ y)\ \in \mathcal{V}_{k\ -\ 1}[\![\ A'\ \nrightarrow_F\ B\ ]\!]$	(2), by (1)
$\Gamma\ =\ \Gamma_1,\ y\ \mapsto \mathsf{indirect}\ w$	(3), by assumption
$(\Gamma_1,\ w)\ \in \mathcal{V}_{k\ -\ 2}[\![\ A'\ \nrightarrow_F\ B\ ]\!]$	(4), by (2) and (3)
$(\Gamma,\ w)\ \in \mathcal{V}_{k\ -\ 2}[\![\ A'\ \nrightarrow_F\ B\ ]\!]$	(5), by Lemma 30
$((\Gamma,\ x\ \mapsto \mathsf{indirect}\ w),\ x)\ \in \mathcal{V}_{k\ -\ 1}[\![\ A'\ \nrightarrow_F\ B\ ]\!]$	(6), by (5)

Lemma 35. (*Indirecting to indirect preserves value denotation.*)
$((\Gamma,\ x\ \mapsto \mathsf{indirect}\ w),\ v)\ \in \mathcal{V}_k[\![\ A\ ]\!]$ and $y\ \mapsto \mathsf{indirect}\ w\ \in \Gamma$ $((\Gamma,\ x\ \mapsto \mathsf{indirect}\ y),\ v)\ \in \mathcal{V}_{k\ -\ 1}[\![\ A\ ]\!]$
Proof. By induction on k. If $k\ =\ 0$, then obvious. Else: induction on A.

Case $A\ =\ 1$: Obvious

Case $A\ =\ A'\ +\ B$ or $A\ =\ A'\ \times B$: Follows directly from the inner inductive hypothesis.

Case $A\ =\ \mu\alpha.\ A'$: Follows directly from the outer inductive hypothesis.

Case $A\ =\ A'\ \nrightarrow_F\ B,\ z\ \neq x$. If $z\ \mapsto \mathsf{memo}\ w'\ \in (\Gamma,\ x\ \mapsto \mathsf{indirect}\ w)$, then the claim follows from the inner inductive hypothesis. If $z\ \mapsto \mathsf{indirect}\ w'\ \in (\Gamma,\ x\ \mapsto \mathsf{indirect}\ w)$, then the claim follows from the outer inductive hypothesis. If $z\ \mapsto \mathsf{lazy}_F\ w'\ \in (\Gamma,\ x\ \mapsto \mathsf{indirect}\ w)$, then the claim follows from the inner inductive hypothesis and applying $\textsc{indi}$ to $\Delta\ \sqsubseteq \Theta$.

Case $A\ =\ A'\ \nrightarrow_F\ B, z\ =\ x$.

$((\Gamma,\ x\ \mapsto \mathsf{indirect}\ w),\ x)\ \in \mathcal{V}_k[\![\ A\ ]\!]$	(1), by assumption
$(\Gamma,\ w)\ \in \mathcal{V}_j[\![\ A'\ \nrightarrow_F\ B\ ]\!]$	(2), for $j\ <\ k$, by (1)
$\Gamma\ =\ \Gamma_1,\ y\ \mapsto \mathsf{indirect}\ w$	(3), by assumption
$(\Gamma,\ y)\ \in \mathcal{V}_k[\![\ A'\ \nrightarrow_F\ B\ ]\!]$	(4), by (2) and (3)
$((\Gamma,\ x\ \mapsto \mathsf{indirect}\ y),\ x)\ \in \mathcal{V}_{k\ +\ 1}[\![\ A'\ \nrightarrow_F\ B\ ]\!]$	(5), by (4)
$((\Gamma,\ x\ \mapsto \mathsf{indirect}\ y),\ x)\ \in \mathcal{V}_{k\ -\ 1}[\![\ A'\ \nrightarrow_F\ B\ ]\!]$	(6), by (5) and Lemma 29

Lemma 36. (*Store extension preserves types.*)
If $(\Delta,\ v)\ \in \mathcal{V}_k[\![\ A\ ]\!]$ and $\Delta\ \sqsubseteq_j\ \Theta$, then $(\Theta,\ v)\ \in \mathcal{V}_{k\ -\ j}[\![\ A\ ]\!]$.
Proof. By induction on j and then case-split on $\Delta\ \sqsubseteq_j\ \Theta$.

Case $\textsc{cutm}$:

$\Gamma,\ x\ \mapsto \mathsf{indirect}\ y\ \sqsubseteq_1\ \Gamma,\ x\ \mapsto \mathsf{memo}\ w$	(1), by assumption
$y\ \mapsto \mathsf{memo}\ w\ \in \Gamma$	(2), by (1)
$((\Gamma,\ x\ \mapsto \mathsf{indirect}\ y),\ v)\ \in \mathcal{V}_k[\![\ A\ ]\!]$	(3), by assumption
$((\Gamma,\ x\ \mapsto \mathsf{memo}\ w),\ v)\ \in \mathcal{V}_{k\ -\ 1}[\![\ A\ ]\!]$	(4), by Lemma 32

Case $\textsc{cuti}$:

$\Gamma,\ x \mapsto \mathsf{indirect}\ y \sqsubseteq_1 \Gamma,\ x \mapsto \mathsf{indirect}\ z$ (1), by assumption
$y \mapsto \mathsf{indirect}\ z \in \Gamma$ (2), by (1)
$((\Gamma,\ x \mapsto \mathsf{indirect}\ y),\ v) \in \mathcal{V}_k[\![A]\!]$ (3), by assumption
$((\Gamma,\ x \mapsto \mathsf{indirect}\ z),\ v) \in \mathcal{V}_{k-1}[\![A]\!]$ (4), by Lemma 34

Case INDM:
$\Gamma,\ x \mapsto \mathsf{memo}\ w \sqsubseteq_1 \Gamma,\ x \mapsto \mathsf{indirect}\ y$ (1), by assumption
$y \mapsto \mathsf{memo}\ w \in \Gamma$ (2), by (1)
$((\Gamma,\ x \mapsto \mathsf{memo}\ w),\ v) \in \mathcal{V}_k[\![A]\!]$ (3), by assumption
$((\Gamma,\ x \mapsto \mathsf{indirect}\ y),\ v) \in \mathcal{V}_{k-1}[\![A]\!]$ (4), by Lemma 33

Case INDI:
$\Gamma,\ x \mapsto \mathsf{indirect}\ z \sqsubseteq_1 \Gamma,\ x \mapsto \mathsf{indirect}\ y$ (1), by assumption
$y \mapsto \mathsf{indirect}\ z \in \Gamma$ (2), by (1)
$((\Gamma,\ x \mapsto \mathsf{indirect}\ z),\ v) \in \mathcal{V}_k[\![A]\!]$ (3), by assumption
$((\Gamma,\ x \mapsto \mathsf{indirect}\ y),\ v) \in \mathcal{V}_{k-1}[\![A]\!]$ (4), by Lemma 35

C.2.4 Referential Transparency.

Lemma 37. (*Evaluation in store extension*)
If $(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k (\Delta,\ x \mapsto \mathsf{memo}\ w)$, then there are stores $\Gamma',\ \Delta'$ such that $\Gamma' : F\ v \Downarrow_j \Delta' : w$ and $j \leqslant k$ and $(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Gamma',\ x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Delta',\ x \mapsto \mathsf{memo}\ w) \sqsubseteq (\Delta,\ x \mapsto \mathsf{memo}\ w)$.
Proof. By induction on $(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) \sqsubseteq_k (\Delta,\ x \mapsto \mathsf{memo}\ w)$.

Case CUTM: Impossible.

Case CUTI: Impossible.

Case INDM: Impossible.

Case INDI: Impossible.

Theorem 8. (*Short-cutting indirections is referentially transparent.*)
If $\Gamma : e \Downarrow_k \Delta : v$ and $\Gamma \sqsubseteq_j \Gamma'$, then $\Gamma' : e \Downarrow_l \Delta' : v$ with $\Delta \sqsubseteq_{l'} \Delta'$.
Proof. By induction on k and then case-split on the evaluation. If $k = 0$, then only VALUE applies. Then the claim follows directly. If $k > 0$:

Case UNFOLD: The claim follows directly.

Case LET: By the inductive hypothesis.

Case APP: By the inductive hypothesis.

Case SPLIT: By the inductive hypothesis.

Case CASE: By the inductive hypothesis.

Case LAZY:
$\Gamma : lv \Downarrow_1 (\Gamma,\ z \mapsto lv) : z$ (1), by LAZY
$\Gamma \sqsubseteq_j \Gamma'$ (2), by assumption
$\Gamma' : lv \Downarrow_1 (\Gamma',\ z \mapsto lv) : z$ (3), by LAZY
$(\Gamma,\ z \mapsto lv) \sqsubseteq_j (\Gamma',\ z \mapsto lv)$ (4), by (2) and Lemma 24

Case RECALL:
$\Gamma : \mathsf{step}\ z \Downarrow_1 \Gamma : w$ (1), by RECALL
$z \mapsto \mathsf{memo}\ w \in \Gamma$ (2), by RECALL
$\Gamma \sqsubseteq_j \Gamma'$ (3), by assumption
$z \mapsto \mathsf{memo}\ w \in \Gamma'$ (4), by (3) and Lemma 25
$\Gamma' : \mathsf{step}\ z \Downarrow_1 \Gamma' : w$ (5), by (4) and RECALL

Case STEP: We have $\Gamma = \Gamma_1, z \mapsto \mathsf{lazy}_F\ v$ and $\Delta = \Delta_1,\ z \mapsto \mathsf{memo}\ w$. If $\Gamma' = \Gamma_1', z \mapsto \mathsf{lazy}_F\ v$:

$\Gamma : \mathsf{step}\ z \Downarrow_{k+1} \Delta : w$	(1), by STEP
$\Gamma_1 : F\ v \Downarrow_k \Delta_1 : w$	(2), by STEP
$\Gamma_1' : F\ v \Downarrow_l \Delta_1' : w$	(3), by the inductive hypothesis
$\Delta_1 \sqsubseteq \Delta_1'$	(4), by the inductive hypothesis
$\Gamma_1', z \mapsto \mathsf{lazy}_F\ v : \mathsf{step}\ z \Downarrow_1 \Delta_1',\ z \mapsto \mathsf{memo}\ w : w$	(5), by (3) and STEP
$\Delta \sqsubseteq \Delta_1',\ z \mapsto \mathsf{memo}\ w$	(6), by (4) and Lemma 24

Else $\Gamma' = \Gamma_1', z \mapsto \mathsf{memo}\ w$:

$\Gamma : \mathsf{step}\ z \Downarrow_{k+1} \Delta : w$	(1), by STEP
$\Gamma_1 : F\ v \Downarrow_k \Delta_1 : w$	(2), by STEP
$\Gamma_1'' : F\ v \Downarrow_j \Delta_1' : w'$	(3), by Lemma 37
$\Gamma \sqsubseteq (\Gamma_1'', z \mapsto \mathsf{lazy}_F\ v) \sqsubseteq (\Delta_1', z \mapsto \mathsf{memo}\ w) \sqsubseteq \Gamma'$	(4), by Lemma 37
$\Gamma_1 \sqsubseteq \Gamma_1''$	(5), by Lemma 26
$\Gamma_1'' : F\ v \Downarrow_j \Delta_1' : w$	(6), by the inductive hypothesis
$\Delta_1 \sqsubseteq \Delta_1'$	(7), by the inductive hypothesis
$\Delta \sqsubseteq (\Delta_1', z \mapsto \mathsf{memo}\ w)$	(8), by Lemma 24
$\Delta \sqsubseteq \Gamma'$	(9), by transitivity

C.3 Soundness of Implementation

C.3.1 Implementation calculus. In this section, we give the full rules for the Implementation Calculus.

$$\frac{}{\varnothing \mid \Gamma,\ x : A \vdash x : A}\ \textsc{var} \qquad \frac{L_1 \mid \Gamma \vdash e_1 : A \quad L_2 \mid \Gamma,\ x : A \vdash e_2 : B}{L_1,\ L_2 \mid \Gamma \vdash \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2}\ \textsc{let}$$

$$\frac{\varnothing \mid \Gamma \vdash v : A_i}{\varnothing \mid \Gamma \vdash \mathsf{in}_i\ v : A_l + A_r}\ \textsc{inl/inr} \qquad \frac{\varnothing \mid \Gamma \vdash v : A_l + A_r \quad L \mid \Gamma,\ x : A_i \vdash e_i : C}{L \mid \Gamma \vdash \mathsf{case}\ v\ \{\ \mathsf{inl}\ x \to e_l;\ \mathsf{inr}\ y \to e_r\ \} : C}\ \textsc{case}$$

$$\frac{\varnothing \mid \Gamma \vdash v : A \quad \varnothing \mid \Gamma \vdash w : B}{\varnothing \mid \Gamma \vdash (v,\ w) : A \times B}\ \textsc{pair} \qquad \frac{\varnothing \mid \Gamma \vdash v : A \times B \quad L \mid \Gamma,\ x : A,\ y : B \vdash e : C}{L \mid \Gamma \vdash \mathsf{split}\ v\ \{\ (x,\ y) \to e\ \} : C}\ \textsc{split}$$

$$\frac{\varnothing \mid \Gamma \vdash v : A[\mu\alpha.\ A/\alpha]}{\varnothing \mid \Gamma \vdash \mathsf{fold}\ v : \mu\alpha.\ A}\ \textsc{fold} \qquad \frac{\varnothing \mid \Gamma \vdash v : \mu\alpha.\ A}{\varnothing \mid \Gamma \vdash \mathsf{unfold}\ v : A[\mu\alpha.\ A/\alpha]}\ \textsc{unfold}$$

$$\frac{}{\varnothing \mid \Gamma \vdash () : 1}\ \textsc{unit} \qquad \frac{F : A \to B \in \Sigma \quad \varnothing \mid \Gamma \vdash v : A}{\varnothing \mid \Gamma \vdash F\ v : B}\ \textsc{app}$$

$$\frac{}{\Vdash \varnothing}\ \textsc{defbase} \qquad \frac{\Vdash \Sigma \quad x : A \vdash e : B}{\Vdash \Sigma, F(x) = e : A \to B}\ \textsc{deffun}$$

$$\frac{F : A \to B \in \Sigma \quad \varnothing \mid \Gamma \vdash v : A}{\varnothing \mid \Gamma \vdash \mathsf{lazy}_F\ v : A \nrightarrow_F B} \qquad \frac{F : A \to B \in \Sigma \quad \varnothing \mid \Gamma \vdash v : B}{\varnothing \mid \Gamma \vdash \mathsf{memo}\ v : A \nrightarrow_F B}$$

$$\frac{\Vdash \Sigma \quad l : A \mid x : B \vdash e : C}{\Vdash \Sigma, F(l; x) = e : A \to B \to C}\ \textsc{deflapp} \qquad \frac{F : A \to B \to C \in \Sigma \quad \varnothing \mid \Gamma \vdash v : B}{l : A \mid \Gamma \vdash F\ l\ v : C}\ \textsc{lapp}$$

$$\frac{\varnothing \mid \Gamma \vdash w : B}{l : B \mid \Gamma \vdash \text{memoize } l\ w : B}\ \textsc{memoize}$$

$$\frac{\varnothing \mid \Gamma \vdash v : A \nrightarrow_F B \quad L, l : B \mid x : A \vdash e_1 : C \quad L \mid \Gamma,\ y : B \vdash e_2 : C}{L \mid \Gamma \vdash \text{lazy match } v\ \{\ \text{lazy}_F\ l\ x \rightarrow e_1;\ \text{memo } y \rightarrow e_2\ \} : C}\ \textsc{lazymatch}$$

C.3.2 Soundness of Implementation Calculus.

Lemma 38. (*The low-level calculus implements the high-level calculus*)
If $\Gamma \vdash e : A$, then $\varnothing \mid \Gamma \vdash e : A$.

Proof. By induction on $\Gamma \vdash e : A$. Obvious for any rule except STEP. We define:

$$\begin{aligned}
\text{step } x\ &= \text{lazy match } x \\
&\quad \text{lazy}_F\ l\ v \rightarrow \text{let } w\ =\ F\ v \text{ in memoize } l\ w \\
&\quad \text{memo } y \rightarrow y
\end{aligned}$$

$F : A \rightarrow B\ \in \Sigma$	(1), by assumption
$\varnothing \mid v : A \vdash F\ v : B$	(2), by APP
$l : B \mid v : A,\ w : B \vdash \text{memoize } l\ w : B$	(3), by MEMOIZE
$l : B \mid v : B \vdash \text{let } w\ =\ F\ v \text{ in memoize } l\ w : B$	(4), by LET
$y : B \vdash y : B$	(5), by VAR
$x : A \nrightarrow_F > B \vdash x : A\ F\ B$	(6), by VAR
$\varnothing \mid x : A \nrightarrow_F B \vdash \text{lazy match } x\ \{ \dots \} : B$	(7), by LAZYMATCH
$\Vdash \text{step}(x)\ =\ \text{lazy match } x\ \{ \dots \} : (A \nrightarrow_F B) \rightarrow B$	(8), by DEFFUN

Lemma 39. (*The small-step semantics implements high-level semantics.*)
If $\Gamma : e \Downarrow_k \Delta : v$, then $\Gamma \mid e \longmapsto^* \Delta \mid v$.

Proof. By induction on k and case-split on $\Gamma : e \Downarrow_k \Delta : v$.

Case VALUE:

$\Gamma : v \Downarrow_0 \Gamma : v$	(1), by assumption
$\Gamma \mid v \longmapsto^* \Gamma \mid v$	(2), since reduction is finished

Case APP:

$\Gamma : F\,v \Downarrow_{k+1} \Delta : w$	(1), by assumption
$F(x)\ =\ e\ \in \Sigma$	(2), by (1)
$\Gamma : e[v/x] \Downarrow_k \Delta : w$	(3), by (2)
$\Gamma \mid e[v/x] \longmapsto^* \Delta \mid w$	(4), by the inductive hypothesis
$\Gamma \mid F\,v \longrightarrow \Gamma \mid e[v/x]$	(5), by (*app*)
$\Gamma \mid F\,v \longmapsto^* \Delta \mid w$	(6), by (4) and (5)

Case LAZY: If $lv\ =\ \text{lazy}_F\ v$:

$\Gamma : \text{lazy}_F\ v \Downarrow_1 (\Gamma,\ z \mapsto \text{lazy}_F\ v) : z$	(1), by assumption
$\Gamma \mid \text{lazy}_F\ v \longrightarrow (\Gamma,\ z \mapsto \text{lazy}_F\ v) \mid z$	(2), by (*lazy*)

If $lv\ =\ \text{memo } w$:

$\Gamma : \text{memo } w \Downarrow_1 (\Gamma,\ z \mapsto \text{memo } w) : z$	(1), by assumption
$\Gamma \mid \text{memo } w \longrightarrow (\Gamma,\ z \mapsto \text{memo } w) \mid z$	(2), by (*memo*)

Case LET:

$$\begin{array}{ll}
\Gamma : \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 \Downarrow_{k+j+1} \Theta : w & \text{(1), by assumption} \\
\Gamma : e_1 \Downarrow_k \Delta : v & \text{(2), by (1)} \\
\Delta : e_2[v/x] \Downarrow_j \Theta : w & \text{(3), by (1)} \\
\Gamma \mid e_1 \longmapsto^* \Delta \mid v & \text{(4), by the inductive hypothesis} \\
\Delta \mid e_2[v/x] \longmapsto^* \Theta \mid w & \text{(5), by the inductive hypothesis} \\
\Delta \mid \mathsf{let}\ x = v\ \mathsf{in}\ e_2 \longrightarrow \Delta \mid e_2[v/x] & \text{(6), by } (\textit{let}) \\
\Gamma \mid \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 \longmapsto^* \Theta \mid w & \text{(7), by (4),(5),(6)}
\end{array}$$

Case SPLIT:

$$\begin{array}{ll}
\Gamma : \mathsf{split}\ (v_1,\ v_2)\ \{\ (x,\ y) \to e\ \} \Downarrow_{k+1} \Delta : w & \text{(1), by assumption} \\
\Gamma : e[v_1/x,\ v_2/y] \Downarrow_k \Delta : w & \text{(2), by (1)} \\
\Gamma \mid e[v_1/x,\ v_2/y] \longmapsto^* \Delta \mid w & \text{(3), by the inductive hypothesis} \\
\Gamma \mid \mathsf{split}\ (v_1,\ v_2)\ \{\ (x,\ y) \to e\ \} \longrightarrow \Delta \mid e[v_1/x,\ v_2/y] & \text{(4), by } (\textit{split}) \\
\Gamma \mid \mathsf{split}\ (v_1,\ v_2)\ \{\ (x,\ y) \to e\ \} \longmapsto^* \Delta \mid w & \text{(5), by (3),(4)}
\end{array}$$

Case UNFOLD:

$$\begin{array}{ll}
\Gamma : \mathsf{unfold}\ (\mathsf{fold}\ v) \Downarrow_1 \Gamma : v & \text{(1), by assumption} \\
\Gamma \mid \mathsf{unfold}\ (\mathsf{fold}\ v) \longmapsto^* \Gamma \mid v & \text{(2), by } (\textit{unfold})
\end{array}$$

Case CASE:

$$\begin{array}{ll}
\Gamma : \mathsf{case}\ (\mathsf{in}_i\ v)\ \{\ \mathsf{in}_l\ x_l \to e_l;\ \mathsf{in}_r\ x_r \to e_r\ \} \Downarrow_{k_i\,+\,1} \Delta : w & \text{(1), by assumption} \\
\Gamma : e_i[v/x_i] \Downarrow_{k_i} \Delta : w & \text{(2), by (1)} \\
\Gamma \mid e_i[v/x_i] \longmapsto^* \Delta \mid w & \text{(3), by the inductive hypothesis} \\
\Gamma \mid \mathsf{case}\ (\mathsf{in}_i\ v)\ \{\ \mathsf{in}_l\ x_l \to e_l;\ \mathsf{in}_r\ x_r \to e_r\ \} \longrightarrow \Delta \mid e_i[v/x_i] & \text{(4), by } (\textit{case}) \\
\Gamma \mid \mathsf{case}\ (\mathsf{in}_i\ v)\ \{\ \mathsf{in}_l\ x_l \to e_l;\ \mathsf{in}_r\ x_r \to e_r\ \} \longmapsto^* \Delta \mid w & \text{(5), by (3),(4)}
\end{array}$$

Case RECALL:

$$\begin{array}{ll}
\Gamma : \mathsf{step}\ z \Downarrow_1 \Gamma : v & \text{(1), by assumpt} \\
z \mapsto \mathsf{memo}\ v \in \Gamma & \text{(2), by (1)} \\
\Gamma \mid \mathsf{step}\ z \longrightarrow \Gamma \mid \mathsf{lazy\ match}\ z\ \{\ \mathsf{lazy}_F\ l\ v \to \mathsf{let}\ w = F\ v\ \mathsf{in\ memoize}\ l\ w;\ \mathsf{memo}\ y \to y\ \} & \text{(3), by } (\textit{app}) \\
\Gamma \mid \mathsf{step}\ z \longrightarrow \Gamma \mid y[v/y] & \text{(4), by } (\textit{lazyma} \\
\Gamma \mid \mathsf{step}\ z \longmapsto^* \Gamma \mid v & \text{(5), by (4)}
\end{array}$$

Case STEP:

$$\begin{array}{l}
(\Gamma,\ x \mapsto \mathsf{lazy}_F\ v) : \mathsf{step}\ x \Downarrow_{k+1} (\Delta,\ x \mapsto \mathsf{memo}\ w) : w \\
\Gamma : F\ v \Downarrow_k \Delta : w \\
\Gamma \mid F\ v \longmapsto^* \Delta \mid w \\
(\Gamma,\ x \mapsto \mathsf{locked}) \mid F\ v \longmapsto^* (\Delta,\ x \mapsto \mathsf{locked}) \mid w \\
(\Gamma,\ z \mapsto \mathsf{lazy}_F\ v) \mid \mathsf{step}\ z \longrightarrow (\Gamma,\ z \mapsto \mathsf{lazy}_F\ v) \mid \mathsf{lazy\ match}\ z\ \{\ \mathsf{lazy}_F\ l\ v \to \mathsf{let}\ w = F\ v\ \mathsf{in\ memoize}\ l\ w; \\
(\Gamma,\ z \mapsto \mathsf{lazy}_F\ v) \mid \mathsf{step}\ z \longrightarrow (\Gamma,\ z \mapsto \mathsf{locked}) \mid \mathsf{let}\ w = F\ v\ \mathsf{in\ memoize}\ z\ w \\
(\Gamma,\ z \mapsto \mathsf{lazy}_F\ v) \mid \mathsf{step}\ z \longmapsto^* (\Gamma,\ z \mapsto \mathsf{locked}) \mid \mathsf{memoize}\ z\ w \\
(\Gamma,\ z \mapsto \mathsf{lazy}_F\ v) \mid \mathsf{step}\ z \longmapsto^* (\Gamma,\ z \mapsto \mathsf{memo}\ w) \mid w
\end{array}$$

C.3.3 Short-cutting.

Lemma 40. (*The low-level calculus implements the high-level calculus*)
If $\Gamma \vdash e : A$, then $\varnothing \mid \Gamma \vdash e : A$.

Proof. By induction on $\Gamma \vdash e : A$. Obvious for any rule except EVAL. We define:

$$\begin{array}{l}
\mathsf{eval}\ x = \mathsf{lazy\ match}\ x \\
\qquad \mathsf{lazy}_F\ l\ v \to \mathsf{let}\ w = F\ v\ \mathsf{in\ case}\ w\ \{\ \mathsf{inl}\ y \to \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \to (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \} \\
\qquad \mathsf{indirect}\ y \to \mathsf{eval}\ y \\
\qquad \mathsf{memo}\ y \to y
\end{array}$$

$F : A \to B + (A \twoheadrightarrow_F B) \in \Sigma$	(1), by assumption
$\varnothing \mid v : A \vdash F\ v : B + (A \twoheadrightarrow_F B)$	(2), by APP
$l : A \twoheadrightarrow_F B \mid y : B \vdash \mathsf{memoize}\ l\ y : B$	(3), by MEMOIZE
$\varnothing \mid z : A \twoheadrightarrow_F B \vdash \mathsf{eval}\ z : B$	(4), by APP
$l : A \twoheadrightarrow_F B \mid y : A \twoheadrightarrow_F B \vdash \mathsf{indirect}\ l\ y : B$	(5), by INDIRECT
$l : A \twoheadrightarrow_F B \mid y : A \twoheadrightarrow_F B \vdash \mathsf{let}\ z = \mathsf{indirect}\ l\ y\ \mathsf{in}\ \mathsf{eval}\ z : B$	(6), by LET
$l : A \twoheadrightarrow_F B \mid w : B + (A \twoheadrightarrow_F B) \vdash \mathsf{case}\ w\ \{ \dots \}$	(7), by CASE
$l : A \twoheadrightarrow_F B \mid v : A \vdash \mathsf{let}\ w = F\ v\ \mathsf{in}\ \mathsf{case}\ w\ \{ \dots \}$	(8), by LET
$\varnothing \mid y : A \twoheadrightarrow_F B \vdash \mathsf{eval}\ y : B$	(9), by APP
$\varnothing \mid y : B \vdash y : B$	(10), by VAR
$\varnothing \mid x : A \twoheadrightarrow_F B \vdash \mathsf{lazy\ match}\ x\ \{ \dots \} : B$	(11), by LAZYMATCH
$\Vdash \mathsf{eval}(x) = \mathsf{lazy\ match}\ x\ \{ \dots \} : (A \twoheadrightarrow_F B) \to B$	(7), by DEFFUN

C.3.4 Tail-recursive Evaluation. We use typical equational reasoning laws as given in e.g. Leijen and Lorenzen [2023]. Our translation function is:

$$[\![e]\!]_l = \mathsf{case}\ (\mathsf{memoize}\ l\ e)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}$$

which yields the calculations:

$$\begin{aligned}
[\![\mathsf{let}\ y = e_1\ \mathsf{in}\ e_2]\!]_l &= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{let}\ y = e_1\ \mathsf{in}\ e_2))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{case}\ (\mathsf{let}\ y = e_1\ \mathsf{in}\ \mathsf{memoize}\ l\ e_2)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{let}\ y = e_1\ \mathsf{in}\ \mathsf{case}\ (\mathsf{memoize}\ l\ e_2)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{let}\ y = e_1\ \mathsf{in}\ [\![e_2]\!]_l
\end{aligned}$$

$$\begin{aligned}
&[\![\mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to e_1;\ \mathsf{inr}\ y \to e_2\ \}]\!]_l\\
&\quad= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to e_1;\ \mathsf{inr}\ y \to e_2\ \}))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&\quad= \mathsf{case}\ (\mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to \mathsf{memoize}\ l\ e_1;\ \mathsf{inr}\ y \to \mathsf{memoize}\ l\ e_2\ \})\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&\quad= \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to \mathsf{case}\ (\mathsf{memoize}\ l\ e_1)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \};\ \mathsf{inr}\ y \to \mathsf{case}\ (\mathsf{memoize}\ l\ e_2)\ \{\ \mathsf{i}\\
&\quad= \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to [\![e_1]\!]_l;\ \mathsf{inr}\ y \to [\![e_2]\!]_l\ \}
\end{aligned}$$

$$\begin{aligned}
&[\![\mathsf{split}\ v\ \{\ (y,\ z) \to e\ \}]\!]_l\\
&\quad= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{split}\ v\ \{\ (y,\ z) \to e\ \}))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&\quad= \mathsf{case}\ (\mathsf{split}\ v\ \{\ (y,\ z) \to \mathsf{memoize}\ l\ e\ \})\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&\quad= \mathsf{split}\ v\ \{\ (y,\ z) \to \mathsf{case}\ (\mathsf{memoize}\ l\ e)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\ \}\\
&\quad= \mathsf{split}\ v\ \{\ (y,\ z) \to [\![e]\!]_l\ \}
\end{aligned}$$

$$\begin{aligned}
[\![\mathsf{inl}\ w]\!]_l &= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{inl}\ w))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{let}\ v = \mathsf{memoize}\ l\ (\mathsf{inl}\ w)\ \mathsf{in}\ \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{memoize}\ l\ (\mathsf{inl}\ w);\ \mathsf{case}\ (\mathsf{inl}\ w)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{memoize}\ l\ (\mathsf{inl}\ w);\ w
\end{aligned}$$

$$\begin{aligned}
[\![\mathsf{inr}\ w]\!]_l &= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{inr}\ w))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{let}\ v = \mathsf{memoize}\ l\ (\mathsf{inr}\ w)\ \mathsf{in}\ \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{memoize}\ l\ (\mathsf{inr}\ w);\ \mathsf{case}\ (\mathsf{inr}\ w)\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&= \mathsf{memoize}\ l\ (\mathsf{inr}\ w);\ \mathsf{eval}\ w
\end{aligned}$$

$$\begin{aligned}
&[\![\mathsf{inr}\ (\mathsf{fold}\ (\mathsf{lazy}_F\ v))]\!]\\
&\quad= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ (\mathsf{lazy}_F\ v))))\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&\quad= \mathsf{let}\ y = \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \mathsf{eval}\ (\mathsf{fold}\ y)\\
&\quad= \mathsf{let}\ y = \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \mathsf{lazy\ match}\ y\\
&\qquad\quad \mathsf{lazy}_F\ l\ v \to F'\ l\ v\\
&\qquad\quad \mathsf{memo}\ v \to \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \to y;\ \mathsf{inr}\ y \to \mathsf{eval}\ y\ \}\\
&\quad= \mathsf{let}\ y = \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \mathsf{lock}\ y\ \mathsf{in}\ F'\ y\ v
\end{aligned}$$

where

$\mathsf{lock}\ y\ \mathsf{in}\ e := \mathsf{lazy\ match}\ y\ \{\ \mathsf{lazy}_F\ y\ \_ \rightarrow e;\ \mathsf{memo}\ v \rightarrow \mathsf{impossible}\ \}$

C.3.5 Short-cutting. With the short-cutting primitives, we use the modified evaluation function:

$$\begin{aligned}\mathsf{eval}\ x\ =\ & \mathsf{lazy\ match}\ x\\ & \quad \mathsf{lazy}_F\ l\ v \rightarrow \mathsf{case}\ (F\ v)\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ & \quad \mathsf{indirect}\ y \rightarrow \mathsf{eval}\ y\\ & \quad \mathsf{memo}\ y \rightarrow y\end{aligned}$$

and thus the translation:

$$[\![e]\!]_l\ = \mathsf{case}\ e\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}$$

which yields the calculations:

$$\begin{aligned}[\![\mathsf{let}\ y\ =\ e_1\ \mathsf{in}\ e_2]\!]_l\ &=\ \mathsf{case}\ (\mathsf{let}\ y\ =\ e_1\ \mathsf{in}\ e_2)\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ &=\ \mathsf{let}\ y\ =\ e_1\ \mathsf{in\ case}\ e_2\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ &=\ \mathsf{let}\ y\ =\ e_1\ \mathsf{in}\ [\![e_2]\!]_l\end{aligned}$$

$$\begin{aligned}&[\![\mathsf{case}\ v\ \{\ \mathsf{inl}\ y \rightarrow e_1;\ \mathsf{inr}\ y \rightarrow e_2\ \}]\!]_l\\ &\quad=\ \mathsf{case}\ (\mathsf{case}\ v\ \{\ \mathsf{inl}\ y \rightarrow e_1;\ \mathsf{inr}\ y \rightarrow e_2\ \})\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ &\quad=\ \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{case}\ (e_1)\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \};\ \mathsf{inr}\ y \rightarrow \mathsf{case}\ (e_2\\ &\quad=\ \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \rightarrow [\![e_1]\!]_l;\ \mathsf{inr}\ y \rightarrow [\![e_2]\!]_l\ \}\end{aligned}$$

$$\begin{aligned}&[\![\mathsf{split}\ v\ \{\ (y,\ z) \rightarrow e\ \}]\!]_l\\ &\quad=\ \mathsf{case}\ (\mathsf{split}\ v\ \{\ (y,\ z) \rightarrow e\ \})\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ &\quad=\ \mathsf{split}\ v\ \{\ (y,\ z) \rightarrow \mathsf{case}\ (e)\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\ \}\\ &\quad=\ \mathsf{split}\ v\ \{\ (y,\ z) \rightarrow [\![e]\!]_l\ \}\end{aligned}$$

$$\begin{aligned}[\![\mathsf{inl}\ w]\!]_l &=\ \mathsf{case}\ (\mathsf{inl}\ w)\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ &=\ \mathsf{memoize}\ l\ w\end{aligned}$$

$$\begin{aligned}[\![\mathsf{inr}\ w]\!]_l &=\ \mathsf{case}\ (\mathsf{inr}\ w)\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ &=\ \mathsf{indirect}\ l\ w;\ \mathsf{eval}\ w\end{aligned}$$

$$\begin{aligned}&[\![\mathsf{inr}\ (\mathsf{lazy}_F\ v)]\!]\\ &\quad=\ \mathsf{case}\ (\mathsf{inr}\ (\mathsf{lazy}_F\ v))\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ &\quad=\ \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in\ case}\ (\mathsf{inr}\ y)\ \{\ \mathsf{inl}\ y \rightarrow \mathsf{memoize}\ l\ y;\ \mathsf{inr}\ y \rightarrow (\mathsf{indirect}\ l\ y;\ \mathsf{eval}\ y)\ \}\\ &\quad=\ \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in\ indirect}\ l\ y;\ \mathsf{eval}\ y\\ &\quad=\ \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in\ indirect}\ l\ y;\ \mathsf{lazy\ match}\ y\\ &\qquad\quad \mathsf{lazy}_F\ l\ v \rightarrow F'\ l\ v\\ &\qquad\quad \mathsf{memo}\ v \rightarrow \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \rightarrow y;\ \mathsf{inr}\ y \rightarrow \mathsf{eval}\ y\ \}\\ &\quad=\ \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in\ indirect}\ l\ y;\ \mathsf{lock}\ y\ \mathsf{in}\ F'\ y\ v\end{aligned}$$

After $F'\ y\ v$, y points to a chain of indirections ending in a memo. The CUTM short-cutting rule gives us the law $\mathsf{indirect}\ l\ y;\ \mathsf{memoize}\ y\ z\ =\ \mathsf{memoize}\ l\ z;\ \mathsf{memoize}\ y\ z$, while the CUTI rule yields $\mathsf{indirect}\ l\ y;\ \mathsf{indirect}\ y\ z\ =\ \mathsf{indirect}\ l\ z;\ \mathsf{indirect}\ y\ z$. Then we can replace the evaluation of $F'\ y\ v$

by $F'\ l\ v$. Finally, the cell y becomes unused altogether and we can remove it:

$$
\begin{array}{l}
\ \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in indirect } l\ y;\ \text{lock } y \text{ in } F'\ y\ v \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in indirect } l\ y;\ \text{lock } y \text{ in case } (F\ y\ v) \\
\qquad \text{inl } z \longrightarrow \text{memoize } y\ z \\
\qquad \text{inr } z \longrightarrow (\text{indirect } y\ z;\ \text{eval } z) \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in lock } y \text{ in indirect } l\ y;\ \text{case } (F\ y\ v) \\
\qquad \text{inl } z \longrightarrow \text{memoize } y\ z \\
\qquad \text{inr } z \longrightarrow (\text{indirect } y\ z;\ \text{eval } z) \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in lock } y \text{ in case } (F\ y\ v) \\
\qquad \text{inl } z \longrightarrow (\text{indirect } l\ y;\ \text{memoize } y\ z) \\
\qquad \text{inr } z \longrightarrow (\text{indirect } l\ y;\ \text{indirect } y\ z;\ \text{eval } z) \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in lock } y \text{ in case } (F\ y\ v) \\
\qquad \text{inl } z \longrightarrow (\text{memoize } l\ z;\ \text{memoize } y\ z) \\
\qquad \text{inr } z \longrightarrow (\text{indirect } l\ z;\ \text{indirect } y\ z;\ \text{eval } z) \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in lock } y \text{ in case } (F\ y\ v) \\
\qquad \text{inl } z \longrightarrow (\text{memoize } l\ z;\ \text{memoize } y\ z) \\
\qquad \text{inr } z \longrightarrow (\text{indirect } l\ z;\ \text{indirect } y\ z;\ \text{eval } z) \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in lock } y \text{ in case } (F\ y\ v) \\
\qquad \text{inl } z \longrightarrow (\text{memoize } y\ z;\ \text{memoize } l\ z) \\
\qquad \text{inr } z \longrightarrow (\text{indirect } y\ z;\ \text{indirect } l\ z;\ \text{eval } z) \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in lock } y \text{ in case } (F\ y\ v) \\
\qquad \text{inl } z \longrightarrow (\text{indirect } y\ l;\ \text{memoize } l\ z) \\
\qquad \text{inr } z \longrightarrow (\text{indirect } y\ l;\ \text{indirect } l\ z;\ \text{eval } z) \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in lock } y \text{ in indirect } y\ l;\ \text{case } (F\ y\ v) \\
\qquad \text{inl } z \longrightarrow \text{memoize } l\ z \\
\qquad \text{inr } z \longrightarrow (\text{indirect } l\ z;\ \text{eval } z) \\
= \ \text{let } y \ = \ \text{lazy}_F\ v \text{ in lock } y \text{ in indirect } y\ l;\ F'\ l\ v \\
= \ F'\ l\ v
\end{array}
$$

C.3.6 Schorr-Waite Evaluation of Lazy Constructors. We use the translation function:

$$[\![e]\!]_{z,l} = \text{case } (\text{memoize } l\ e)\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \}$$

which yields the calculations:

$$
\begin{array}{lcl}
[\![\text{let } y \ = \ e_1 \text{ in } e_2]\!]_{z,l} & = & \text{case } (\text{memoize } l\ (\text{let } y \ = \ e_1 \text{ in } e_2))\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \} \\
& = & \text{case } (\text{let } y \ = \ e_1 \text{ in memoize } l\ e_2)\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \} \\
& = & \text{let } y \ = \ e_1 \text{ in case } (\text{memoize } l\ e_2)\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \} \\
& = & \text{let } y \ = \ e_1 \text{ in } [\![e_2]\!]_{z,l}
\end{array}
$$

$$
\begin{array}{l}
[\![\text{case } v\ \{\ \text{inl } y \longrightarrow e_1;\ \text{inr } y \longrightarrow e_2\ \}]\!]_{z,l} \\
\qquad = \ \text{case } (\text{memoize } l\ (\text{case } v\ \{\ \text{inl } y \longrightarrow e_1;\ \text{inr } y \longrightarrow e_2\ \}))\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \} \\
\qquad = \ \text{case } (\text{case } v\ \{\ \text{inl } y \longrightarrow \text{memoize } l\ e_1;\ \text{inr } y \longrightarrow \text{memoize } l\ e_2\ \})\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z \\
\qquad = \ \text{case } v\ \{\ \text{inl } y \longrightarrow \text{case } (\text{memoize } l\ e_1)\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \};\ \text{inr } y \longrightarrow \text{case } (\text{memo} \\
\qquad = \ \text{case } v\ \{\ \text{inl } y \longrightarrow [\![e_1]\!]_{z,l};\ \text{inr } y \longrightarrow [\![e_2]\!]_{z,l}\ \}
\end{array}
$$

$$
\begin{array}{l}
[\![\text{split } v\ \{\ (y,\ z) \longrightarrow e\ \}]\!]_{z,l} \\
\qquad = \ \text{case } (\text{memoize } l\ (\text{split } v\ \{\ (y,\ z) \longrightarrow e\ \}))\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \} \\
\qquad = \ \text{case } (\text{split } v\ \{\ (y,\ z) \longrightarrow \text{memoize } l\ e\ \})\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \} \\
\qquad = \ \text{split } v\ \{\ (y,\ z) \longrightarrow \text{case } (\text{memoize } l\ e)\ \{\ \text{inl } y \longrightarrow \text{unroll } z\ y;\ \text{inr } y \longrightarrow \text{eval}'\ z\ y\ \}\ \} \\
\qquad = \ \text{split } v\ \{\ (y,\ z) \longrightarrow [\![e]\!]_{z,l}\ \}
\end{array}
$$

$$
\begin{aligned}
[\![\mathsf{inl}\ w]\!]_{z,l} &= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{inl}\ w))\ \{\ \mathsf{inl}\ y \longrightarrow \mathsf{unroll}\ z\ y;\ \ \mathsf{inr}\ y \longrightarrow \mathsf{eval}'\ z\ y\ \}\\
&= \mathsf{let}\ v\ =\ \mathsf{memoize}\ l\ (\mathsf{inl}\ w)\ \mathsf{in}\ \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \longrightarrow \mathsf{unroll}\ z\ y;\ \ \mathsf{inr}\ y \longrightarrow \mathsf{eval}'\ z\ y\ \}\\
&= \mathsf{memoize}\ l\ (\mathsf{inl}\ w);\ \ \mathsf{case}\ (\mathsf{inl}\ w)\ \{\ \mathsf{inl}\ y \longrightarrow \mathsf{unroll}\ z\ y;\ \ \mathsf{inr}\ y \longrightarrow \mathsf{eval}'\ z\ y\ \}\\
&= \mathsf{memoize}\ l\ (\mathsf{inl}\ w);\ \ \mathsf{unroll}\ z\ w
\end{aligned}
$$

$$
\begin{aligned}
[\![\mathsf{inr}\ w]\!]_{z,l} &= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{inr}\ w))\ \{\ \mathsf{inl}\ y \longrightarrow \mathsf{unroll}\ z\ y;\ \ \mathsf{inr}\ y \longrightarrow \mathsf{eval}'\ z\ y\ \}\\
&= \mathsf{let}\ v\ =\ \mathsf{memoize}\ l\ (\mathsf{inr}\ w)\ \mathsf{in}\ \mathsf{case}\ v\ \{\ \mathsf{inl}\ y \longrightarrow \mathsf{unroll}\ z\ y;\ \ \mathsf{inr}\ y \longrightarrow \mathsf{eval}'\ z\ y\ \}\\
&= \mathsf{memoize}\ l\ (\mathsf{inr}\ w);\ \ \mathsf{case}\ (\mathsf{inr}\ w)\ \{\ \mathsf{inl}\ y \longrightarrow \mathsf{unroll}\ z\ y;\ \ \mathsf{inr}\ y \longrightarrow \mathsf{eval}'\ z\ y\ \}\\
&= \mathsf{memoize}\ l\ (\mathsf{inr}\ w);\ \ \mathsf{eval}'\ z\ w
\end{aligned}
$$

$$
\begin{aligned}
&[\![\mathsf{inr}\ (\mathsf{fold}\ (\mathsf{lazy}_F\ v))]\!]\\
&\quad= \mathsf{case}\ (\mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ (\mathsf{lazy}_F\ v))))\ \{\ \mathsf{inl}\ y \longrightarrow \mathsf{unroll}\ z\ y;\ \ \mathsf{inr}\ y \longrightarrow \mathsf{eval}'\ z\ y\ \}\\
&\quad= \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \ \mathsf{eval}'\ z\ (\mathsf{fold}\ y)\\
&\quad= \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \ \mathsf{lazy\ match}\ (\mathsf{unfold}\ x)\\
&\qquad\quad \mathsf{lazy}_F\ l\ v \longrightarrow F'\ z\ l\ v\\
&\qquad\quad \mathsf{memo}\ y \longrightarrow \mathsf{case}\ y\ \{\ \mathsf{inl}\ y \longrightarrow \mathsf{unroll}\ z\ y;\ \ \mathsf{inr}\ y \longrightarrow \mathsf{eval}'\ z\ y\ \}\\
&\quad= \mathsf{let}\ y\ =\ \mathsf{lazy}_F\ v\ \mathsf{in}\ \mathsf{memoize}\ l\ (\mathsf{inr}\ (\mathsf{fold}\ y));\ \ \mathsf{lock}\ y\ \mathsf{in}\ F'\ z\ y\ v
\end{aligned}
$$