

A Lightweight Approach to Formal Algorithmic Complexity

Functional Pearl

ANONYMOUS AUTHOR(S)

Proof assistants are typically used to verify programs' correctness – yet reasoning about the performance of programs appears to be less straightforward. This paper explores a lightweight technique for reasoning about the algorithmic complexity of purely functional data structures and algorithms: defining a cost function by induction on a program's graph. To illustrate this approach, this paper presents the worst case and amortized analysis of several data structures, including batched queues and real time queues. Using recent work on first order laziness, this paper reasons formally about sharing in a purely functional setting, without introducing an explicit heap or resorting to separation logic. Simple induction suffices.

CCS Concepts: • **Theory of computation** → **Complexity theory and logic**; **Program specifications**; **Functional constructs**; **Data structures design and analysis**.

1 Introduction

Q: *Can proof assistants reason about complexity of programs?*

A: *Complexity theory is just math. And all math can be formalized.*

– an exchange on the Proof Assistants StackExchange forum¹

Modern proof assistants are highly effective tools for reasoning about purely functional programs. While there is a large body of work proving functional correctness [see for instance, Appel 2026], this paper ploughs a different furrow: exploring a lightweight approach to reasoning about algorithmic complexity of purely functional algorithms and data structures. The approach presented here have several advantages:

- The key technique is simple enough to explain to an undergraduate student, yet powerful enough to handle a wide range of examples – including algorithms for sorting and searching, and data structures, such as batched queues and real time queues. Unlike existing approaches, this approach does not rely on type annotations of functions [Danielsson 2008] or working in a particular logical framework [Grodin et al. 2024; Niu et al. 2022]. In fact, the study of a function's algorithmic complexity is entirely decoupled from the function's definition, enabling the post-hoc verification of non-functional properties.
- Besides analysing the worst case time complexity, we prove the amortized complexity cost of ephemeral data structures. Using dependent types, the distinction between ephemeral usage and persistent usage is made explicit.
- Using recent work on *first-order laziness* [Lorenzen et al. 2025], we show how to reason about data structures that rely on lazy evaluation and sharing. This is notoriously difficult to do, yet our proposed solution avoids introducing an explicit heap or mutation.

¹<https://proofassistants.stackexchange.com/questions/1976/can-proof-assistants-reason-about-complexity-of-programs>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICFP '26, Indiana University, Indianapolis

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM

<https://doi.org/XXXXXXX.XXXXXXX>

This paper aims to dispel the widespread misconception that the *only* properties proof assistants can reason about is their extensional behaviour. Although the development we present is in Agda, the same ideas can be readily ported to any other proof assistant using dependent types, including Lean, Idris or Rocq.

2 Code is data: structural induction on the graph of a function

To illustrate the key ideas, we begin by revisiting a classical example of functional programming: list reversal. The ‘fast’ algorithm for list reversal uses an accumulating parameter:

```
reverse-acc : (xs acc : List A) → List A
reverse-acc [] acc = acc
reverse-acc (x : xs) acc = reverse-acc xs (x : acc)

fast-reverse : List A → List A
fast-reverse xs = reverse-acc xs []
```

The ‘slow’ algorithm for list reversal repeatedly calls the append function on lists:

```
slow-reverse : List A → List A
slow-reverse [] = []
slow-reverse (x : xs) = slow-reverse xs + [ x ]
```

A simple inductive argument shows that these two algorithms produce the same result. The key property relating these two definitions reads as follows:

```
reverse-correct : (xs ys : List A) → reverse-acc xs ys ≡ reverse xs + ys
```

Yet we have no way to argue that one algorithm is faster than the other. In fact, if the ambient type theory supports *function extensionality*, both these functions are *equal*, as they produce the same output for input lists.

What to do? Even in an intensional type theory, there is no way to observe a function’s implementation. In this paper, we propose a lightweight approach: evoking the old functional programmer’s maxim that ‘code is data.’ The key idea is to make the definition of a function explicit by reifying the *graph* of the function as an inductively defined *relation*. In general, the graph of a function $f : A \rightarrow B$ is a relation $R : A \rightarrow B \rightarrow \text{Set}$, such that $R\ x\ (f\ x)$ holds. In the case of the `reverseAcc` function, this corresponds to the following data type:

```
data ReverseAcc : List A → List A → List A → Set where
  base : ReverseAcc [] ys ys
  step : ReverseAcc xs (x : ys) zs → ReverseAcc (x : xs) ys zs
```

You may want to think of this as a data type describing the *call graph* of the `reverse-acc` function; or alternatively, a Prolog programming for reversing one list onto another. Such relations provide an *intensional* inductive description of the function’s behaviour. It is straightforward to prove that this definition is sound and complete with respect to the function’s definition:

```
Reverse-acc-sound : ReverseAcc xs acc ys → reverse-acc xs acc ≡ ys
Reverse-acc-complete : reverse-acc xs acc ≡ ys → ReverseAcc xs acc ys
```

But crucially, reifying the definition of the function in this fashion lets us compute the *size* of the graph and relate this to the size of its inputs. To do so, we define a simple *cost function*, a function that counts the size of the graph of our reversal function:

99 $\text{cost-reverse-acc} : \text{ReverseAcc } xs \ ys \ zs \rightarrow \mathbb{N}$

100 $\text{cost-reverse-acc base} = 1$

101 $\text{cost-reverse-acc (step } r) = 1 + \text{cost-reverse-acc } r$

102 The accumulating reverse function takes a number of steps linear in the size of its first argument.
 103 With these definitions, however, we make this more precise by *proving* that $\text{reverse-acc } xs \ ys$
 104 requires one step for each element of xs and one final step in the base case:
 105

106 $\text{bound-reverse-acc} : (\text{rev} : \text{ReverseAcc } xs \ ys \ zs) \rightarrow \text{cost-reverse-acc rev} \equiv 1 + \text{length } xs$

107 The proof uses straightforward induction on the graph of reverse-acc . In this fashion, reasoning
 108 about the complexity of a function definition is no harder than proving its functional correctness.
 109

110 2.1 Size and *Ordnung*

111 Algorithmic complexity relates size of inputs to the number of steps necessary to produce the
 112 output. To measure things systematically, we introduce some overloaded notation that we will
 113 use throughout this paper. Typically, we are interested in measuring the ‘size’ of our inputs. We
 114 introduce the corresponding record, Size , with a single field:
 115

116 **record** $\text{Size} (A : \text{Set}) : \text{Set}$ **where**

117 **field**

118 $|_| : A \rightarrow \mathbb{N}$

119 For example, we typically consider the size of a list to be its length:

120 **instance**

121 $\text{size-List} : \text{Size} (\text{List } A)$

122 $\text{size-List} = \text{record } \{ |_| = \text{length} \}$

123 We will use Agda’s *instance arguments* [Devriese and Piessens 2011] to find the necessary Size
 124 record, when needed. Now this notion of Size only works for types in Set . The graph of a function
 125 $f : A \rightarrow B$, however, will have type $A \rightarrow B \rightarrow \text{Set}$. For such families of types, we introduce a
 126 separate notion of size for indexed families of types:
 127
 128

129 **record** $\text{Size}_2 (P : A \rightarrow B \rightarrow \text{Set}) : \text{Set}$ **where**

130 **field**

131 $|_||_2 : \forall \{x : A\} \{y : B\} \rightarrow P \ x \ y \rightarrow \mathbb{N}$

132 Similar definitions exist for relations on more than two arguments. In the remainder of this paper,
 133 we will omit the subscripts, as these can readily be inferred.
 134

135 Typically, we are not interested in the exact number of steps required, but rather want an upper
 136 bound. We introduce a (simplified) version of the typical ‘big O’ notation:

137 $_ \in \text{O} _ : \{ \{ \text{Size } A \} \} \rightarrow (G : A \rightarrow B \rightarrow \text{Set}) \rightarrow \{ \{ \text{Size } G \} \} \rightarrow (\mathbb{N} \rightarrow \mathbb{N}) \rightarrow \text{Set}$

138 $G \in \text{O } f = \exists [c_0] \ \exists [c] \ (g : G \ x \ y) \rightarrow |x| \equiv n \rightarrow |g| \leq c_0 + c * f(n)$

139 We say that the relation G is in the complexity class f when we can find constants c_0 and c ,
 140 such that the cost of the graph $G \ x \ y$ is less than $c_0 + c * f(n)$, for any input x of size n . This
 141 simplified definition will serve our purposes for this paper, but a much more careful treatment of
 142 ‘big O notation’ in type theory can be found in Guéneau’s PhD thesis [2019].
 143

144 Using this notation, we show that the accumulating reversal function is linear in its first argument:

145 $\text{rev-acc-linear} : \forall (ys : \text{List } A) \rightarrow (\lambda \ xs \rightarrow \text{ReverseAcc } xs \ ys) \in \text{O} (\lambda \ n \rightarrow n)$

146

147

2.2 Quadratic complexity

What about our slow reversal function? Can we also prove that it is quadratic? To do so, we once again reify the slow-reverse function, together with the append function it uses, by defining the graphs of both functions:

```
data Append : List A → List A → List A → Set where
  base : Append [] ys ys
  step : Append xs ys zs → Append (x : xs) ys (x : zs)
```

```
data Slow-reverse : List A → List A → Set where
  base : Slow-reverse [] []
  step : Slow-reverse xs ys → Append ys [ x ] zs → Slow-reverse (x : xs) zs
```

Next, we define the obvious cost functions. Since slow-reverse calls append, we add the associated cost to that of the recursive call:

```
append-cost : Append xs ys zs → ℕ
slow-reverse-cost : Slow-reverse xs ys → ℕ
slow-reverse-cost base = 1
slow-reverse-cost (step rev app) = 1 + slow-reverse-cost rev + append-cost app
```

Unsurprisingly, we show that the slow-reverse functions takes a quadratic number of steps:

```
slow-reverse-quadratic : Slow-reverse ∈ O (λ n → n * n)
```

The proof proceeds by induction on the graph of the slow-reverse function. The proof requires a few auxiliary properties of addition and multiplication, but is otherwise straightforward.

2.3 Limitations and caveats

Before covering more interesting examples, we want to reflect briefly on the general approach. Firstly, the function’s graph and the associated cost are a crucial part of the problem specification. It is all too easy provide ‘incorrect’ definitions. Here is an alternative definition of the graph of our slow-reverse function:

```
data Slow-reverse-wrong : List A → List A → Set where
  wrong : Slow-reverse-wrong xs (slow-reverse xs)
```

The data type, while it also describes the graph of slow-reverse, is not very useful. Crucially, we expect the graphs to follow the same recursive structure as the functions they represent. Provided the type indices are built from constructors and variables, this will be the case.

Many proof assistants, including Rocq and Lean, provide automation for generating a function’s induction principles from its definition [Breitner 2024; Sozeau 2010; Sozeau and Mangin 2019]; generating the corresponding inductively defined graph is very similar. In Agda, we could also use the reflection mechanism to generate these definitions automatically [van der Walt and Swierstra 2012]. There is some manual overhead (and room for error) in writing these definitions by hand – but doing so is instructive at first.

Similarly, the cost functions that calculate the expected number of steps are part of the specification. Consider the following definition:

```
wrong-cost : Slow-reverse-wrong xs ys → ℕ
wrong-cost _ = 1
```

197 The wrong-cost function defines a cost function for our slow-reverse function, albeit not a very
 198 useful one! It assumes that every slow reverse takes constant time, which is not typically what one
 199 would expect. One could, once again, use metaprogramming or generic programming techniques
 200 to generate these cost functions to avoid this issue [Backhouse et al. 1998; Escot and Cockx 2022;
 201 Jansson and Jeuring 1997; Ko et al. 2022].

202 This drawback, however, is not unique to our approach: many textbooks on algorithms leave
 203 underspecified what operations actually count towards a program’s execution time. Just as any other
 204 formally verified proof, the value of the proof inherently depends on the specification. Nonetheless,
 205 there are situations where having some flexibility in defining a suitable cost functions is useful – as
 206 we will see next.

207

208 2.4 Insertion sort

209 Consider *insertion sort*. The insert function inserts a new element into a sorted list:

210 $\text{insert} : \mathbb{N} \rightarrow \text{List } \mathbb{N} \rightarrow \text{List } \mathbb{N}$
 211 $\text{insert } y [] = [y]$
 212 $\text{insert } y (x : xs) = \text{if } y \leq^b x \text{ then } y : x : xs \text{ else } x : \text{insert } y xs$

214 As this function has two results when the input list is non-empty, there are two constructors in the
 215 corresponding graph:

216

217 **data** $\text{Insert } (y : \mathbb{N}) : \text{List } \mathbb{N} \rightarrow \text{List } \mathbb{N} \rightarrow \text{Set}$ **where**

218 $\text{base} : \text{Insert } y [] [y]$
 219 $\text{stop} : y \leq x \rightarrow \text{Insert } y (x : xs) (y : x : xs)$
 220 $\text{step} : y > x \rightarrow \text{Insert } y xs zs \rightarrow \text{Insert } y (x : xs) (x : zs)$

221 The obvious cost function is appears equally straightforward:

222

223 $\text{cost-insert} : \text{Insert } y xs zs \rightarrow \mathbb{N}$
 224 $\text{cost-insert base} = 1$
 225 $\text{cost-insert (stop } x \leq y) = 1 + \text{cost-}\leq x \leq y$

226 **where**

227 $\text{cost-}\leq : x \leq y \rightarrow \mathbb{N}$
 228 $\text{cost-}\leq _ = 1$

229 $\text{cost-insert (step } y > x \text{ rec)} = \text{cost-}\rightarrow y > x + \text{succ (cost-insert rec)}$

230 **where**

231 $\text{cost-}\rightarrow : y > x \rightarrow \mathbb{N}$
 232 $\text{cost-}\rightarrow _ = 1$

233

234 Yet we have made an important design decision here: we have chosen to consider the cost
 235 of comparing two natural numbers to be constant. These seems to be an obvious choice – but if
 236 these natural numbers are implemented naively using Peano numerals, this cost is *not* constant: it
 237 takes longer to compare bigger numbers. Given that we expect to compile this code to (arbitrary
 238 precision) machine integers, this seems a reasonable choice. This illustrates how having some
 239 choice in defining the cost function is important and how this forms part of our specification.

240 With these definitions in place, it is easy to prove that the insert function is linear in the length
 241 of its input list:

242

243 $\text{insert-linear} : (y : \mathbb{N}) \rightarrow \text{Insert } y \in O(\lambda n \rightarrow n)$

244 As a result, repeatedly calling insert leads to a quadratic time sorting function.

245

246 2.5 Logs come from trees

247 Our next example explores binary search trees. There are many different ways to enforce the
 248 desired invariants of our data structure: intrinsically or extrinsically? Do we require the trees to be
 249 balanced? Or are we only interested in the ordering of elements? The purpose of this example is to
 250 demonstrate how to incorporate additional invariants extrinsically in the cost analysis. As a result,
 251 we choose a simply typed definition of binary trees:

```
252 data Tree : Set where  

  253   leaf  : Tree  

  254   node  : Tree → ℕ → Tree → Tree
```

256 The find function looks up its argument key k in a given tree:

```
257 find : ℕ → Tree → Bool  

  258 find k leaf      = false  

  259 find k (node l x r) = if  $k \equiv^b x$  then true  

  260                   else (if  $k <^b x$  then find k l else find k r)
```

262 What is the complexity of this find function? To answer this question, we need a notion of size for
 263 our trees, the graph of the find function, and the corresponding cost function:

```
264 size-tree : Tree → ℕ  

  265 size-tree leaf      = 1  

  266 size-tree (node l x r) = 1 + size-tree l + size-tree r
```

```
268 data Find (k : ℕ) : Tree → Bool → Set
```

```
269 cost-find : Find k t b → ℕ
```

271 We have left out the definitions of the graph and cost function, as we hope the pattern is familiar by
 272 now. Using these definitions, we prove that the *worst case complexity* is linear in the size of the tree:

```
273 worst-case : (k : ℕ) → Find k ∈O (λ n → n)
```

275 Oftentimes, however, we do not search through an arbitrary tree, but rather one that is suitably
 276 balanced. We could add indices to our Tree type to ensure that the tree is balanced – but as we
 277 mentioned previously, we instead choose to add this assumption to the *graph* of the find function.
 278 To simplify our analysis, we consider *perfect* trees, where the left and right subtree have the exact
 279 same size:

```
280 Balanced : Tree → Tree → Set  

  281 Balanced l r = size-tree l ≡ size-tree r
```

283 It is straightforward to relax this constraint, bounding the difference in size between the two subtrees,
 284 as is the case for AVL trees [Adelson-Velskii and Landis 1962] and Braun trees [Hoogerwoord 1992a;
 285 Rem and Braun 1983], for instance.

286 We can add this additional assumption to the graph of the find function:

```
287 data Find-Balanced (k : ℕ) : Tree → Bool → Set where  

  288   base : Find-Balanced k leaf false  

  289   stop  : Find-Balanced k (node l k r) true  

  290   left  : Balanced l r → k < x → Find-Balanced k l b → Find-Balanced k (node l x r) b  

  291   right : Balanced l r → k > x → Find-Balanced k r b → Find-Balanced k (node l x r) b  

  292 cost-find-balanced : Find-Balanced k t b → ℕ
```

294

Here the two possible recursive calls, corresponding to the left and right constructors, introduce an explicit assumption that the two subtrees are have equal size. Using this assumption, we can prove the expected logarithmic complexity:

worst-case-balanced : $(k : \mathbb{N}) \rightarrow \text{Find-Balanced } k \in O(\lambda n \rightarrow \log_2 n)$

Of course, this analysis hinges on the assumption that the tree is perfectly balanced. To prove that this graph accurately reifies the find function on perfectly balanced trees, we prove the following lemma, establishing that the revised graph of our find function accurately reflects the search through a perfect tree:

Perfect : Tree \rightarrow Set

Perfect leaf = \top

Perfect (node l _ r) = Balanced l r \times Perfect l \times Perfect r

find-perfect : $(k : \mathbb{N}) \rightarrow (t : \text{Tree}) \rightarrow \text{Perfect } t \rightarrow \text{Find-Balanced } k \ t \ (\text{find } k \ t)$

We could have achieved a similar result in a much simpler fashion: defining an inductive family of perfectly balanced trees. This choice would have simplified the analysis enormously – the purpose of the example is to show how to incorporate extrinsic assumptions in our cost analysis. Alternatively, we could also pass the Perfect assumption as an argument to the cost-find function, as we will do towards the end of this paper. As is often the case, there are many different ways to formulate and prove the same result.

3 Batched queues: amortized cost

All our analyses so far establish upper bounds on the *worst case* time complexity of our functions. In this section we explore a classic *amortized* analysis of a queue implementation using two lists. The code in Figure 1 is a variation of the batched queues from Okasaki’s classic textbook [1998]. Each queue consists of two lists, the front and the rear. Elements are enqueued onto the rear, but removed from the front. Once the front runs out of elements, the elements in the rear are reversed to form the new front.

We enforce one invariant intrinsically: when the front of the queue is empty, so is the rear. By doing so, we know that when the deq operation encounters an empty front queue, there is no need to inspect (or reverse) the rear. Note that both the enq and deq operations are not recursive. The balance function, however, that restores our invariant, may take non-constant time, if the rear needs reversing.

To say anything about the cost of these queue operations, we define their graphs:

data Balance : List A \times List A \rightarrow Q A \rightarrow Set **where**

stop : Balance (x : xs , ys) (queue (x : xs) ys (λ ()))

work : Reverse xs ys \rightarrow Balance ([], xs) (queue ys []) (const tt)

data Enq : Q A \rightarrow Q A \rightarrow Set **where**

enq-q : Balance (front q , x : rear q) q₁ \rightarrow Enq q q₁

data Deq : Q A \rightarrow Q A \rightarrow Set **where**

deq-: : $\forall \{ \text{inv} \} \rightarrow$ Balance (xs , ys) q \rightarrow Deq (queue (x : xs) ys inv) q

Note that Balance may call the (linear) list reversal function defined in the previous section. The obvious cost functions count the size of these graphs:

```

344
345
346 record Q (A : Set) : Set where
347   constructor queue
348   field
349     front : List A
350     rear  : List A
351     inv   : So (null front) → So (null rear)
352
353 open Q
354 empty : Q A
355 empty = queue [] [] id
356 balance : List A × List A → Q A
357 balance ([] , rear) = queue (reverse rear) [] (const tt)
358 balance (front @ (_ : _) , rear) = queue front rear λ ()
359
360 enq : A → Q A → Q A
361 enq x q = balance (front q , x : rear q)
362
363 deq : Q A → Maybe (A × Q A)
364 deq (queue [] r i) = nothing
365 deq (queue (x : xx) rear i) = just (x , balance (xx , rear))
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392

```

Fig. 1. Implementing queues using two lists

```

370 cost-enq    : Enq q0 q1 → ℕ
371 cost-deq    : Deq q0 q1 → ℕ
372 cost-balance : Balance (xs,ys) q → ℕ
373

```

As balance may require reversing the rear list, the worst case time complexity of balance is linear in the length of the rear list; correspondingly, the worst case complexity of both enq and deq is linear:

```

377 balance-linear : Balance ∈ O (λ n → n)
378

```

As an aside, we've omitted the Size instance for pairs of lists; in this particular example, we define the size of the pair to be the length of the second component, i.e., the length of the rear list.

3.1 Sums of costs

The key idea behind an *amortized* analysis is to find an upper bound on the total cost of a *series of operations*. If the expensive operations are sufficiently rare, these should not have a modest effect on the total cost of the series. We will reproduce the amortized analysis of batched queues in Agda.

Our analysis shifts from the cost of a single operation to series of operations. What are valid series? We restrict ourselves to two queue operations, enqueueing and dequeuing:

```

389 data Op (q0 : Q A) : Q A → Set where
390   enq-op : Enq q0 q1 → Op q0 q1
391   deq-op : Deq q0 q1 → Op q0 q1
392

```

393 The data type $\text{Op } q_0 \ q_1$ is either an enq or deq operation, taking the old queue q_0 to a new queue
 394 q_1 . Even though dequeuing may fail (when the queue is empty), the *graph* of Deq only considers
 395 the successful case – hence never have to worry about the possibility of failure, even if the deq
 396 *function* may not return a result. Next, we chain together a list of queue operations, one after the
 397 other:

398 **data** $\text{Ops } (q_0 : \text{Q } A) : \text{Q } A \rightarrow \text{Set}$ **where**

399 $\text{stop} : \text{Ops } q_0 \ q_0$

400 $\text{step} : \text{Op } q_0 \ q_1 \rightarrow \text{Ops } q_1 \ q_2 \rightarrow \text{Ops } q_0 \ q_2$

402 Note that the types ensure that each subsequent operation acts on the result of the previous one.

403 The total cost of such a sequence of queue operations is the sum of the cost of the individual
 404 operations:

405 $\text{cost-op} : \text{Op } q_0 \ q_1 \rightarrow \mathbb{N}$

406 $\text{cost-op } (\text{enq-op } e) = \text{cost-enq } e$

407 $\text{cost-op } (\text{deq-op } d) = \text{cost-deq } d$

408 $\text{cost-ops} : \text{Ops } q_0 \ q_1 \rightarrow \mathbb{N}$

409 $\text{cost-ops } (\text{step op ops}) = \text{cost-op } \text{op} + \text{cost-ops } \text{ops}$

410 $\text{cost-ops } \text{stop} = \text{zero}$

411
 412 Our aim is to find an upper bound on the total cost of a series of queue operations. It is not
 413 immediately obvious how to do so: any upper bound must predict how often rebalancing occurs
 414 and what the cost of each expensive rebalancing operation is.
 415

416 3.2 Amortized cost

417 The key idea behind amortized complexity is due to [Tarjan \[1985\]](#), although we closely follow
 418 the presentation by [Hoogerwoord \[1992b\]](#). Rather than reasoning about the cost of the sum of
 419 operations, we define a notion of credit:
 420

421 $\text{credit} : \text{Q } A \rightarrow \mathbb{N}$

422 For now, we leave underspecified what assumptions we make about this credit function. Note
 423 that credit is a function on queues – not on operations. We define the amortized cost of a single
 424 operation as its cost, together with its change in credit:
 425

426 $\text{amortized-cost-op} : \text{Op } q_0 \ q_1 \rightarrow \mathbb{N}$

427 $\text{amortized-cost-op } \{q_0\} \{q_1\} \text{op} = \text{cost-op } \text{op} + \text{credit } q_1 \div \text{credit } q_0$

428 The amortized cost of a series of operations simply computes the sum of amortized costs:
 429

430 $\text{amortized-cost-ops} : \text{Ops } q_0 \ q_1 \rightarrow \mathbb{N}$

431 $\text{amortized-cost-ops } \text{stop} = \text{zero}$

432 $\text{amortized-cost-ops } (\text{step op ops}) = \text{amortized-cost-op } \text{op} + \text{amortized-cost-ops } \text{ops}$

433 Why is this definition useful? When we sum the amortized cost, adjacent credits cancel out! To
 434 make this precise, we prove the following equality:
 435

436 $\text{amortized-property} : (\text{ops} : \text{Ops } q_0 \ q_1) \rightarrow$

437 $\text{credit } q_0 + \text{amortized-cost-ops } \text{ops} \equiv \text{cost-ops } \text{ops} + \text{credit } q_1$

438
 439 If we choose the credit of our empty queue to be zero, then we know that the $\text{amortized-cost-ops}$
 440 gives an upper bound for total cost. The inductive case of the corresponding proof is in [Figure 2](#),
 441

442 credit q_0 + amortized-cost-ops (step op ops)
 443 \equiv (definition of amortized-cost-ops)
 444 credit q_0 + (amortized-cost-op op + amortized-cost-ops ops)
 445 \equiv (definition of amortized-cost)
 446 credit q_0 + ((cost-op op + credit q_1 \div credit q_0) + amortized-cost-ops ops)
 447 \equiv (associativity)
 448 (credit q_0 + (cost-op op + credit q_1 \div credit q_0)) + amortized-cost-ops ops
 449 \equiv (monus-addition associativity)
 450 (credit q_0 + (cost-op op + credit q_1) \div credit q_0) + amortized-cost-ops ops
 451 \equiv (monus-addition cancellation)
 452 (cost-op op + credit q_1) + amortized-cost-ops ops
 453 \equiv (associativity)
 454 cost-op op + (credit q_1 + amortized-cost-ops ops)
 455 \equiv (induction hypothesis)
 456 cost-op op + (cost-ops ops + credit q_2)
 457 \equiv (associativity)
 458 (cost-op op + cost-ops ops) + credit q_2
 459 \equiv (definition of cost-ops)
 460 cost-ops (step op ops) + credit q_2
 461
 462
 463
 464
 465
 466
 467
 468
 469
 470
 471
 472
 473
 474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485
 486
 487
 488
 489
 490

Fig. 2. An equational proof relating amortized and worst case cost

where $op : Op\ q_0\ q_1$ and $ops : Ops\ q_1\ q_2$; the Agda proof terms have been replaced with a more legible description.

This proof does hinge on several lemmas. All costs are positive natural numbers. The definition of amortized cost uses a ‘monus’ operation, i.e., cut-off subtraction on natural numbers. Reasoning about the monus is rather subtle, as it does not satisfy the familiar properties of subtraction on integers. The proof requires that monus cancels addition and that monus and addition are associative:

$$m+n\div m\equiv n : \forall m\ n \rightarrow (m+n)\div m\equiv n$$

$$+-\div\text{-assoc} : \forall m \rightarrow o \leq n \rightarrow (m+n)\div o\equiv m+(n\div o)$$

To use the second lemma, we need to ensure the amortized cost both operations is positive:

$$\text{amortized-cost-positive} : (op : Op\ q_0\ q_1) \rightarrow \text{credit } q_0 \leq \text{cost-op } op + \text{credit } q_1$$

This side condition guarantees that the amortized cost never ‘overspends’. Each operation uses no more credits than are available; credits are never lost.

3.3 Credits and queues

All that remains to be done, is to choose the notion of credit for our batched queues. Okasaki [1998] argues that the *length of the rear* queue gives a suitable notion of credit:

$$\text{credit} : QA \rightarrow \mathbb{N}$$

$$\text{credit } q = \text{length } (\text{rear } q)$$

491 How do our queue operations change the credits available? Each enqueue operation takes constant
 492 time, but increases the accumulated credit as it adds an element to the rear. When rebalancing
 493 occurs, typically after a dequeue operation, we use the accumulated credits to pay for the expensive
 494 list reversal.

495 Why is amortization important? Instead of reasoning about the entire sequence of operations,
 496 we know that the amortized cost of the series gives an upper bound on the total cost. We therefore
 497 prove an upper bound on the amortized cost of each single operation:

498 amortized-bound : (op : Op q₀ q₁) → amortized-cost-op op ≤ 3
 499

500 Why three? Consider an enqueue operation that does not rebalance: 1) it takes a step to enter the
 501 body of enq; 2) it calls balance; 3) it allocates one new credit as the resulting rear list is longer. It is
 502 easy to construct a similar argument for the other queue operations. The most interesting case is
 503 where rebalancing occurs – but by construction, we will have accumulated enough credits to pay
 504 for the expensive reversal operation.

505 If we know that a single operation has a constant amortized cost, the overall cost of any series of
 506 operations is bounded accordingly:

507 amortized-bound-ops : (ops : Ops q₀ q₁) → amortized-cost-ops ops ≤ 3 * | ops |
 508

509 Using the key amortized-property from the figure above, together with the fact that the empty
 510 queue has zero credit, we show that the *total costs* are bounded:

511 amortized-constant : (ops : Ops empty q) → cost-ops ops ≤ 3 * | ops |
 512

512 Hence we conclude that enqueue, dequeue and rebalancing operations are all *amortized constant*
 513 *time operations*.

514 4 Real time queues: laziness and sharing

516 There is one important point we have glossed over in the previous section. Recall the definition of
 517 a series of operations from the previous section:

518 **data** Ops (q₀ : Q A) : Q A → Set **where**
 519

520 stop : Ops q₀ q₀

521 step : Op q₀ q₁ → Ops q₁ q₂ → Ops q₀ q₂

522 Here the queue operations are applied one after the other; each new operation is applied to the
 523 result of the previous one. As a result, the expensive rebalancing operations are rare and do not
 524 dominate the overall cost. Unfortunately, not all programs that work with queues satisfy this pattern
 525 of usage. Consider the following (admittedly contrived) example:

526 f : List A → List A

527 f ys = **let** q = queue [0] [1000000, ..., 2, 1] λ () **in**

528 map (λ y → ... deq q ...) ys
 529

530 This function constructs a queue, q, that will require more than a million steps for the next deq
 531 operation. In this example, this rebalancing happens over and over again as the map function is
 532 executed. The amortized analysis from the previous section hinges on rebalancing a queue only
 533 once – and cannot be used to limit the cost of the function f above. Here the queue q is said to be
 534 used *persistently* rather than *ephemerally*.

535 To address this, Okasaki gives several implementations of queues with *worst case constant bounds*,
 536 including real time queues [Okasaki 1998, §7.2]. Just like batched queues, real time queues are
 537 implemented using a pair of lists; the key differences are in their invariants and rebalancing
 538 operations. Where batched queues perform one expensive rebalance operation every so often; real
 539

time queues perform a part of the expensive rebalancing after every queue operation: the total amount of work is the same, but the work is scheduled differently. The implementation of real time queues, however, makes essential use of *sharing* and *lazy evaluation*.

So far, we have implicitly assumed a strict evaluation order: the cost functions compute the required time to fully evaluate a function call, assuming the arguments are already fully evaluated. How should we account for thunked computations, laziness and sharing? Recently, Pottier et al. [2024] have shown how to use separation logic and the ‘debit style’ reasoning to prove amortized bounds for lazy data structures. Yet we will not succumb to the siren’s call of separation logic – pure functions and simple induction shall suffice.

4.1 First order laziness

Many *implementations* of laziness use a mutable reference to store thunks on the heap; upon evaluation, a thunk is replaced with its value. Rather than model such mutable state in Agda, however, we employ the recent work on *first order laziness* for the Koka programming language [Lorenzen et al. 2025]. In this work, a fixed number of lazy computations may be suspended, in the form of so-called *lazy constructors*. Although Agda lacks the language support that Koka provides for lazy constructors, we will illustrate their key ideas and how to model these in Agda, before returning to the implementation of real time queues.

Consider the following data type declaration:

```
data Stream (A : Set) : Set where
  []   : Stream A
  _:_  : A → Stream A → Stream A
  app  : Stream A → Stream A → Stream A
  rev  : Stream A → Stream A → Stream A
```

Besides the familiar constructors for lists, we have included two additional constructors: `app` appends one stream onto another; `rev` reverses its first argument onto its second. To distinguish between regular and lazy constructors, we need to define which constructors we consider to be in *weak head normal form*:

```
whnf? : Stream A → Set
whnf? []           = ⊤
whnf? (_ : _)     = ⊤
whnf? (app _ _)   = ⊥
whnf? (rev _ _)   = ⊥
```

The familiar list constructors are in weak head normal form. The `app` and `rev` constructors, on the other hand, are *lazy constructors* that require further evaluation. To specify how these constructors are evaluated, we define a single reduction step on streams that are not yet in weak head normal form:

```
data _↪_ : Stream A → Stream A → Set where
  app↪      : xs ↪ xs' → app xs ys ↪ app xs' ys
  rev↪      : xs ↪ xs' → rev xs ys ↪ rev xs' ys
  app↪-base : app [] ys ↪ ys
  app↪-step : app (x : xs) ys ↪ (x : app xs ys)
  rev↪-base : rev [] ys ↪ ys
  rev↪-step : rev (x : xs) ys ↪ rev xs (x : ys)
```

589 We can also define the function implementing this single step that is the identity on streams in
 590 weak head normal form:

591 `step : Stream A → Stream A`

593 In what should be a familiar pattern, the small step relation is the graph of the step function.
 594 Repeatedly executing a single evaluation step, yields a weak head normal form:

595 **data** \rightsquigarrow^* : Stream A → Stream A → Set **where**

596 `step- \rightsquigarrow` : $xs \rightsquigarrow ys \rightarrow ys \rightsquigarrow^* zs \rightarrow xs \rightsquigarrow^* zs$

597 `stop-` : $(x : xs) \rightsquigarrow^* (x : xs)$

598 `stop-[]` : $[] \rightsquigarrow^* []$

600 We can also define the corresponding function, that we will call `seq`. To pattern match on the result
 601 of `seq`, we include an implicit proof argument that the result is indeed in weak head normal form:

602 **data** WHNF (A : Set) : Set **where**

603 `<_>` : $(xs : Stream A) \rightarrow \{whnf? xs\} \rightarrow WHNF A$

604 `seq` : $(xs : Stream A) \rightarrow WHNF A$

606 In this fashion, programming with streams appears quite similar to programming with lists: when
 607 pattern matching on the result of a call to `seq`, we never observe the lazy constructors. For example,
 608 to map over the stream (and thereby force any suspended lazy constructors), we write:

609 `map` : $(f : A \rightarrow B) \rightarrow Stream A \rightarrow List B$

610 `map f xs with seq xs`

611 `... | <[]>` = `[]`

612 `... | <x : xs>` = `f x : map f xs`

614 Here Agda uses the (implicit) proof argument to rule out the cases for `app` and `rev`; these will be
 615 forced by the call to `seq`.

616 After a call to `seq`, a stream may still contain (possibly nested) lazy constructors. If a stream is
 617 fully evaluated, i.e., it does not contain any lazy constructors at all, we shall refer to it as a *list*:

618 `isList?` : Stream A → Set

620 Finally, we define a simple cost function for our `seq` function: counting the number of evaluation
 621 steps necessary to produce a weak head normal form.

622 `cost-steps` : $xs \rightsquigarrow^* ys \rightarrow \mathbb{N}$

623 `cost-steps (step- \rightsquigarrow _ steps)` = `1 + cost-steps steps`

624 `cost-steps stop-` = `1`

625 `cost-steps stop-[]` = `1`

627 Using this cost function, we prove simple properties about the number of steps required to force
 628 evaluation of a lazy constructors to their weak head normal form:

629 `app-cost` : $(steps : app xs ys \rightsquigarrow^* zs) \rightarrow isList? xs \rightarrow isList? ys \rightarrow$

630 `cost-steps steps` ≤ 2

631 `rev-cost` : $(steps : rev xs ys \rightsquigarrow^* zs) \rightarrow isList? xs \rightarrow isList? ys \rightarrow$

632 `cost-steps steps` $\leq 2 + length xs$

634 Unsurprisingly, evaluating an `app` constructor to weak head normal form requires constant time;
 635 evaluating a `rev` constructor is linear in its first argument (provided the arguments to both con-
 636 structors are already fully evaluated).

637

4.2 Real time queues

Now that we have this model of laziness in place, we turn our attention to implementing real time queues. The implementation is given in Figure 3. Just as we saw for batched queues, we maintain a front and a rear stream: new elements are added to the rear; elements are removed from the front. Unlike batched queues, however, we enforce a different invariant, relating the front to a third stream: the *schedule*. What is this schedule? Okasaki writes:

We can think of s (the schedule) in two ways, either as a suffix of f [the front] or as a pointer to the first unevaluated suspension in f .

The invariant that we enforce intrinsically, $\text{sched} \sqsubseteq \text{front}$, in our definition of real time queues will make this statement formal: the schedule is a suffix of the front; the prefix of the front that is not in the schedule is fully evaluated. Furthermore, the schedule stream itself is either:

- a fully evaluated list. Initially, for example, the schedule is the empty list.
- or a suspended computation of the form $\text{app } xs \ (\text{rev } ys \ zs)$ for lists xs , ys , and zs , where the length of ys is equal to the length of $xs + 1$.

If you look ahead to Figure 4, you can see that the rebalancing functions inspect the schedule, forcing its evaluation. By also evaluating one step of the list reversal in the call to the auxiliary function `step-rev`, we increase the cost of each queue operation – but doing so ensures that by the time we run out of elements in xs , the reversal of ys will be almost finished. As we never have to compute an unbounded linear time reversal operation, the worst case cost of each operation remains constant. Once the suspended `app` computation is finished, the remainder of the schedule is a (reversed) list that is fully evaluated – as in the first bullet point above.

Specifying these invariant effectively is one of the key challenges when it comes to reasoning about real time queues. We will return to the definition of rebalancing real time queues shortly, but first we formally define the notion of suffix – together with three other invariants that real time queues satisfy.

4.3 Contexts and suffixes

Just as there are many equivalent definitions of the ‘less than or equals to’ order on natural numbers, there are many ways to specify that one list is a suffix of another. We have chosen a definition that relies on the notion of *one hole contexts*:

data Ctx (A : Set) : Set **where**

\square : Ctx A

$_:_$: A \rightarrow Ctx A \rightarrow Ctx A

A one hole context for our data type of streams is simply a list of values, ending in a hole \square . We require contexts to be fully evaluated, that is, they do not contain lazy constructors. This definition generalises naturally to other data types, with or without lazy constructors: the derivative of the underlying functor gives rise to the corresponding context [Hinze et al. 2004; McBride 2001].

There are two operations on contexts that we will employ: filling in the hole with a stream; appending one context after another. Both operations are easy to define:

$_[_]$: Ctx A \rightarrow Stream A \rightarrow Stream A

$\square [xs] = xs$

$(x : c) [xs] = x : (c [xs])$

```

687
688
689 record Q (A : Set) : Set where
690   constructor queue
691   field
692     front : Stream A
693     rear  : Stream A
694     sched : Stream A
695     inv   : sched  $\sqsubseteq$  front
696
697   empty : Q A
698   empty = queue [] [] []  $\sqsubseteq$ -refl
699
700   enq : Q A  $\rightarrow$  A  $\rightarrow$  Q A
701   enq q y = balance-enq (front q) (y : rear q) (sched q) (inv q)
702
703   deq : Q A  $\rightarrow$  Maybe (A  $\times$  Q A)
704   deq q with seq (front q)
705   ... |  $\langle$  []  $\rangle$  = nothing
706   ... |  $\langle$  x : xx  $\rangle$  = just (x , balance-deq xx (rear q) (sched q) (inv q))
707
708
709
710

```

Fig. 3. Real time queues – definition and operations

```

711  $\_ + \_ : \text{Ctx } A \rightarrow \text{Ctx } A \rightarrow \text{Ctx } A$ 
712  $\square + c_2 = c_2$ 
713  $(x : c_1) + c_2 = x : (c_1 + c_2)$ 
714

```

Minamide [1998] has proposed that these operations on contexts may be made efficient by using an additional pointer that gives immediate access to the end of the context. The Koka programming language implements this idea [Lorenzen et al. 2024] – hidden behind a purely functional interface – but guarantees that both operations are executed by means of constant time pointer manipulation.

Using this definition of context, we define the suffix relation between streams as follows:

```

720 data  $\_ \sqsubseteq \_ (xs : \text{Stream } A) : \text{Stream } A \rightarrow \text{Set} where
721   suffix : (c : Ctx A)  $\rightarrow$  xs  $\sqsubseteq$  c [ xs ]
722$ 
```

In words, $xs \sqsubseteq ys$ holds precisely when we can find a context c such that ys is equal to $c [xs]$. It is straightforward to prove that this relation is both transitive and reflexive. With this definition, we enforce that the prefix of the front is fully evaluated: the prefix is described by a context without lazy constructors. Furthermore, as we shall see, this definition readily captures our intuition behind the sharing of evaluation: if the context c witnesses $xs \sqsubseteq ys$ and $xs \rightsquigarrow xs'$ then $c [xs]$ becomes $c [xs']$.

4.4 Invariants and rebalancing

To complete our definition of real time queues, we need to define the balance functions that restores the property that the schedule is suffix of the front stream – but also forces enough evaluation to ensure that enq and deq remain constant time. To understand how the balance functions work, however, we need to identify several additional invariants that real time queues satisfy. The code

```

736 balance-enq : (front rear sched : Stream A) → (sched ⊆ front) → Q A
737
738 balance-enq .(c [ sched ]) rear sched (suffix c) with seq sched
739 ... | ⟨ [] ⟩ = let f = app (c [ [] ]) (rev rear [])
740               in queue f [] f ⊆-refl
741
742 ... | ⟨ x : xx ⟩ = let sched' = step-rev xx in
743                   let c' = c + (x : []) in
744                       queue (c' [ sched' ]) rear sched' (suffix c')
745
746 step-rev : Stream A → Stream A
747 step-rev (app xs (rev ys zs)) = app xs (step (rev ys zs))
748 step-rev xs = xs
749

```

Fig. 4. Balancing real time queues

to rebalance after an enq operation is given in Figure 4. Before studying this further, however, we identify the invariants that we will prove *extrinsically* about our real time queues.

Invariants. We call a property on queues an *invariant* when it holds for the empty queue and is preserved by each queue operation:

Invariant : (P : Q A → Set) → Set

Invariant P = P empty × (∀ q q' → Op q q' → P q → P q')

The real time queues satisfy three invariants that we now cover one by one.

Our first simple invariant states that the rear of the queue is always a fully evaluated list:

RearInv : Q A → Set

RearInv q = isList? (rear q)

This is obviously true for the empty queue. The enq operation extends the rear list, but does not introduce lazy constructors; the balancing and dequeue operations do not change the rear list. It is straightforward to prove that this invariant holds.

A second invariant relates the lengths of the streams that constitute our queues. Okasaki suggests the following invariant:

LengthInv : Q A → Set

LengthInv q = length (rear q) + length (sched q) ≡ length (front q)

In our implementation, however, we are a bit more precise. The intrinsic invariant, stating that $\text{sched} \subseteq \text{front}$, already tells us something about the length of the front queue, namely it must be equal to the length of the evaluated prefix plus the length of the schedule. A more precise invariant therefore removes the length of the schedule from both sides of this equality:

LengthInv : Q A → Set

LengthInv q = length (rear q) ≡ length-ctx (inv q)

length-ctx : Ctx A → ℕ

length-ctx c = length (c [[]])

Put differently: the evaluated prefix of the front queue has the same length as the rear. Why does this hold? It is clearly true of the empty queue. After each enq operation, the rear is too long – but evaluating the schedule, as done by the call to seq in the balance-enq function, extends the evaluated prefix. After each successful deq operation, the front is too short – but again, evaluating the schedule extends the evaluated prefix, reestablishing the desired invariant.

Both these properties, however, are only ever used to prove the key invariant that we need to establish the constant time behaviour of real time queues. We previously claimed that the schedule was always fully evaluated or of the form `app xs (rev ys zs)`. To formalise this in Agda, we introduce the following predicate, characterising valid schedules:

```
data Schedule : Stream A → Set where
  done   : (isList? xs) → Schedule xs
  thunk  : (xs ys zs : Stream A) → isList? xs → isList? ys → isList? zs →
    succ (length xs) ≡ length ys → Schedule (app xs (rev ys zs))
```

This property is the key to proving that all queue operations are worst case constant time.

Rebalancing. With the specification of the key invariants out of the way, we finally return to our rebalancing operations. After each deq or enq operation, the length invariant is broken: either the front is too short or the rear is too long. To restore this – and the other invariants – there are two rebalancing functions, one of which is given in Figure 4, that we explain now in more detail.

First of all, note that pattern matching on the proof that `sched ⊆ front` tells us *exactly* what the *value* of the front stream is, namely `c [sched]`. The pattern we see for the front queue is *forced*. As the body of the rebalance function evaluates the schedule to weak head normal form, as done by the call `seq sched`, we need to ensure this evaluation is also shared with the front queue. We distinguish two possible cases, one for each possible weak head normal form of the schedule:

- If the schedule is empty, we create a new suspension appending the evaluated prefix of the front to the reversed rear. As a result of the enqueue operation, the length invariant no longer holds: the rear is one longer than the front – but this is exactly what is necessary to ensure the schedule is well-formed: the thunk constructor reestablishes the schedule invariant. Of course, we still need to check that all the streams involved (`c [[]]`, rear and `[]`) are fully evaluated – but the only interesting case is for the rear stream, which we proved separately in the `RearInv` above.
- If the schedule is non-empty, its head is accumulated in the evaluated prefix of the front, extending the context – thereby reestablishing the desired length invariant relating this prefix to the length of the rear. Furthermore, if the schedule contains a suspended rev constructor, one step of this reversal is also executed. This ensures that the new schedule continues to satisfy the Schedule predicate above – where the first argument to `app` is only one shorter than the first argument to `rev`.

While we not gone into detail how to complete the formal proofs that these invariants hold, these are largely straightforward. Most of the proofs need to distinguish the two different cases for rebalancing, keeping track of where evaluation occurs and is shared – but these require very few auxiliary definitions or lemmata.

What about rebalancing after an deq operation? The function has almost exactly the same behaviour. To understand why there are two different balancing functions, it is important to understand how enq and deq break the suffix invariant in different ways. After a *dequeue* operation, the front stream has become too short: we need to take the tail of *both* the front and schedule. We therefore drop the first element of the evaluated prefix of the front:

834 $\text{tail} : \text{Ctx } A \rightarrow \text{Ctx } A$

835 $\text{tail } \square = \square$

836 $\text{tail } (x : c) = c$

837 Besides this modest change, the rebalancing function is identical; the proofs reestablishing the
838 invariants are no more complicated. The complete code for rebalancing can be found in Appendix A.

839 Note that the presentation of real time queues is slightly different from the one given by Okasaki
840 [1998]. The original presentation uses one lazy operation, `rotate`, that takes three arguments
841 – but behaves identically to the combination of `append` and `reverse` used here. One advantage
842 of our definition is that it potentially offers more memory reuse [Lorenzen et al. 2023]. As the
843 lazy constructors are evaluated, the corresponding memory locations may be reused, rather than
844 allocating fresh cons cells.
845

846 4.5 Cost analysis

847 To prove the worst case constant time bounds for queue operations, we rely on one key insight: it
848 takes a constant amount of time to force a valid schedule to weak head normal form:

849 $\text{schedule-constant} : \exists [c] (\text{Schedule } s \rightarrow (\text{steps} : s \rightsquigarrow^* s') \rightarrow \text{cost-steps steps} \leq c)$

850 As a result, forcing the front queue (as is done by the `deq` operation), also takes constant time.
851 Furthermore, both rebalancing functions are not recursive – but do require forcing the schedule.

852 There is one rather subtle point we want to discuss explicitly: our pure specification restores
853 the invariant, $\text{sched} \sqsubseteq \text{front}$, by extending the context after the schedule has been evaluated to
854 weak head normal form. The cost functions consider this operation to have zero cost. We make
855 the assumption that the `let` binding, introduced when rebalancing encounters an empty schedule,
856 creates a *shared* thunk for the front and schedule. Any subsequent evaluation of the schedule is
857 automatically shared with the front. As a result, the cost of ‘updating’ the front after evaluating the
858 schedule is zero; this update is usually left implicit – as the front and the schedule share memory
859 locations – in traditional semantics of lazy evaluation with an explicit heap [Launchbury 1993].
860 For the moment, however, we defer a more thorough discussion about suitable *cost models* to the
861 discussion in the next section.
862

863 Using the lemma above, we prove that each operation takes at most constant time:

864 $\text{real-time-constant} : \exists [c] (\text{Schedule } (\text{sched } q) \rightarrow (\text{op} : \text{Op } q) \rightarrow \text{cost-op op} \leq c)$

865 The proof itself is rather dull – the challenge is identifying and establishing the invariants that
866 make it so.
867

868 This cost analysis does have one interesting aspect: we establish a constant time upper bound
869 on queues *provided the schedule is valid*. Without this assumption, evaluating the schedule to weak
870 head normal form could take an arbitrary amount of time, for instance, when the schedule contains
871 many nested `rev` and `app` constructors. We justify this assumption by proving this invariant holds
872 for every queue arises from a sequence of `enq` and `deq` operations on the empty queue. Any library
873 should therefore take care to provide only these operations as the queue interface – or this analysis
874 will become unsound.
875

876 5 Discussion

877 *Proof assistants.* There is a great deal of related work in the intersection of cost analysis, functional
878 languages, and proof assistants. The recent textbook by Nipkow [2025] gives an overview of purely
879 functional data structures and algorithms, together with their implementation, specification and
880 verification using the interactive proof assistant Isabelle. This book not only gives the precise
881 specification and correctness proofs, but also covers the worst case and amortized complexity
882

883 analysis. Rather than reify the graph of a function explicitly and define the cost functions inductively,
884 the authors use a metaprogram to generate the cost functions from a function definition directly.
885 Nipkow, however, avoids reasoning about lazy evaluation as ‘reasoning about running time under
886 lazy evaluation is nontrivial’ [2025, page 248]. Yet Okasaki writes: ‘functional languages have
887 the curious property that *all* data structures are automatically persistent’ and ‘lazy evaluation
888 can mediate the conflict between amortization and persistence’. Laziness, despite its apparent
889 complexity, is crucial for the efficient implementation of many purely functional data structures.

890 In spirit, our work is most closely related to the Bove-Capretta method [2005]. Bove and Capretta
891 have shown how to reify a function’s call graph to decouple its termination proof from its definition;
892 any function can be made structurally recursive in this fashion. In our work, we reify the entire
893 graph of a function to define a suitable cost function. Yet there are some differences. As our analysis
894 of the slow reverse function shows, it is important to have access to *all* auxiliary function calls
895 – such as the calls to `append` – and not just recursive structure of the function definition under
896 inspection. Furthermore, by including the function’s result in the reified data type, we check that the
897 corresponding data type is sound and complete with respect to its function definition. Where both
898 the Bove-Capretta predicates and function graphs can be generated automatically from a function’s
899 definition, our cost analyses do not require modification to the function’s original definition.

900 There are several related approaches to incorporating the cost analysis of a function together
901 with its definition in a proof assistant. Danielsson [2008] has shown how track the necessary
902 evaluation steps for lazy programs in a function’s type signature in Agda. More recently, the
903 work on cost-aware logical frameworks is also accompanied by an Agda implementation [Grodin
904 et al. 2024; Niu et al. 2022]. Provided users write their function in these libraries, they provide
905 an accurate upper bound on the function’s cost. There is a trade off here: using the lightweight
906 techniques suggested in this paper, cost analysis is decoupled from a function’s definition and can
907 be performed *post hoc*, just as other forms of verification. Reasoning about complexity is not always
908 straightforward: the analysis of real time queues hinges on a complex invariant (the validity of the
909 schedule) and requires a non-trivial argument establish the constant time bounds – these reasoning
910 steps are hard to infer automatically. On the other hand, the approach taken in this paper requires
911 more manual work (and meta-theoretic justification) that these frameworks provide out of the box.
912 Further afield, there are numerous type and effect systems that track a functional program’s cost or
913 resource usage [Hoffmann et al. 2012; Hoffmann and Hofmann 2010; Hoffmann and Jost 2022; Jost
914 et al. 2010; Rajani et al. 2021].

915
916 *Cost models.* What makes a good cost function? We have side stepped the meta-theoretic discus-
917 sion up until this point – arguing that the graphs and cost functions are part of the specification of
918 the cost analysis. To substantiate any claim, however, these cost functions must respect a *cost model*,
919 relating the cost function to the program’s semantics. Fortunately, there is decades of research on
920 this topic [Danner et al. 2015; Gurr 1990; Le Métayer 1988; Rosendahl 1989; Sands 1995; Shultis
921 1984; Van Stone 2003; Wegbreit 1975]. For the first order data structures and functions that we
922 have covered in the first half of this paper, counting function calls – and thereby the size of the
923 graph of a function – suffices [Sands 1990]. The syntactic restriction, requiring all type indices in
924 the function’s graph to be constructors or variables, ensures that all costs of intermediate results
925 are taken into account. Variable access, memory allocation, garbage collection, cache locality, and
926 timing of primitive functions all matter for performance – but do not affect the complexity analysis
927 overall.

928 What about laziness? Can we substantiate the claim that real time queues can be implemented
929 using sharing in the style suggested here? One way to do so would rewrite the real time queues
930 given here with an explicit heap, tracking thunks and aliasing. There is already existing work on
931

formalizing constructor contexts in Iris [Jung et al. 2018] by Lorenzen et al. [2024]. Constructing such a cost model is still necessary to justify the validity of the cost function used in our analysis, but for a pearl such as this, however, we appeal to the reader’s intuition: intrinsically enforcing sharing guarantees an efficient implementation using aliasing is possible.

Amortized analysis and lazy evaluation. Where we have argued above that laziness is essential for efficient purely functional data structures, we have only considered the *worst case* analyses and not *amortized* analyses. Okasaki [1998] has shown how to use debit based reasoning can be used to establish amortized bounds on persistent data structures. The amortized analyses of lazy data structures is notoriously difficult. The natural semantics of laziness given by Launchbury [1993] use a heap to store thunks. Using separation logic, recent work has shown how to record additional information in the ghost state associated with a program [Mével et al. 2019; Pilkiewicz and Pottier 2011; Pottier et al. 2024], making it possible to formalize the informal arguments given by Okasaki.

Alternative semantics for lazy evaluation may make amortized analyses easier. Bjerner and Holmström [1989] give an untyped demand-driven analysis of first order lazy functional programs. This has formed the basis of recent work, formalized in Rocq, on the cost of evaluating lazy functional programs, including an amortized analysis of lazy queues [Xia et al. 2024]. Modelling lazy evaluation in Rocq, however, requires a careful treatment of approximations and inserting explicit ‘ticks’ to count evaluation steps. Alternatively, the recent work on clairvoyant semantics [Hackett and Hutton 2019] shows how call by need can be viewed as a form of non-deterministic call by value. This work, in turn, inspired a ‘memorist semantics’ [Li et al. 2026] that tracks both the cost and usage information by means of an explicit monad.

Lorenzen [2025] has developed a testing library for checking amortized bounds of lazy data structures. Their work suggests using *credit passing style*, where credits are only ever placed on lazy constructors, may still be sound, even for data structures that are used persistently. The key property of amortized analysis in Figure 2 establishes an *equality*; the associated side condition guarantees that operations never overspend. In the persistent setting, however, we may want to adapt this argument: proving an *inequality* instead and allowing credits to be discarded, for instance, when a shared thunk has been evaluated further than you expect. Proving the soundness of such an amortized argument would offer a new avenue towards reasoning about lazy data structures that are used persistently.

Conclusion. In a purely functional language, reasoning relies on elementary methods: functions are defined by structural induction; proofs follow suit. This paper explores how simple induction – albeit on the function’s graph rather than its inputs – offers a convenient way to reason about a function’s computational complexity in a proof assistant with dependent types. In this style, we have written worst case analyses, amortized analyses of ephemeral data structures, and worst case analyses of persistent data structures, offering a practical approach to machine-checked complexity analyses all within a purely functional setting.

References

- M Adelson-Velskii and Evgenii Mikhailovich Landis. 1962. *An algorithm for the organization of information*. Technical Report.
- Andrew W. Appel. 2026. *Verified Functional Algorithms*. Software Foundations, Vol. 3. Electronic textbook. Version 1.6.0, <http://softwarefoundations.cis.upenn.edu>.
- Roland Backhouse, Patrik Jansson, Johan Jeuring, and Lambert Meertens. 1998. Generic programming. In *International School on Advanced Functional Programming*. Springer, 28–115.
- Bror Bjerner and Sören Holmström. 1989. A composition approach to time analysis of first order lazy functional programs. In *Proceedings of the fourth International Conference on Functional Programming Languages and Computer Architecture*. 157–165.

- 981 Ana Bove and Venanzio Capretta. 2005. Modelling general recursion in type theory. *Mathematical Structures in Computer*
982 *Science* 15, 4 (2005), 671–708.
- 983 Joachim Breitner. 2024. Functional Induction theorems. (2024). Published online at [https://lean-lang.org/blog/2024-5-17-](https://lean-lang.org/blog/2024-5-17-functional-induction)
984 [functional-induction](https://lean-lang.org/blog/2024-5-17-functional-induction).
- 985 Nils Anders Danielsson. 2008. Lightweight semiformal time complexity analysis for purely functional data structures. In *Proceedings of the 35th Annual ACM SIGPLAN Symposium on Principles of Programming Languages (POPL '08)*. 133–144.
986 <https://doi.org/10.1145/1328438.1328457>
- 987 Norman Danner, Daniel R Licata, and Ramyaa Ramyaa. 2015. Denotational cost semantics for functional languages with
988 inductive types. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming*. 140–151.
- 989 Dominique Devriese and Frank Piessens. 2011. On the bright side of type classes: instance arguments in Agda. In *Proceedings*
990 *of the 16th ACM SIGPLAN International Conference on Functional Programming (ICFP '11)*. 143–155. [https://doi.org/10.](https://doi.org/10.1145/2034773.2034796)
991 [1145/2034773.2034796](https://doi.org/10.1145/2034773.2034796)
- 992 Lucas Escot and Jesper Cockx. 2022. Practical generic programming over a universe of native datatypes. *Proceedings of the*
993 *ACM on Programming Languages* 6, ICFP (2022), 625–649.
- 994 Harrison Grodin, Yue Niu, Jonathan Sterling, and Robert Harper. 2024. Decalf: A directed, effectful cost-aware logical
995 framework. *Proceedings of the ACM on Programming Languages* 8, POPL (2024), 273–301.
- 996 Armaël Guéneau. 2019. *Mechanized Verification of the Correctness and Asymptotic Complexity of Programs*. Ph.D. Dissertation.
997 Université de Paris.
- 998 Douglas J Gurr. 1990. *Semantic frameworks for complexity*. Ph.D. Dissertation. The University of Edinburgh.
- 999 Jennifer Hackett and Graham Hutton. 2019. Call-by-need is clairvoyant call-by-value. *Proceedings of the ACM on Programming*
1000 *Languages* 3, ICFP (2019), 1–23.
- 1001 Ralf Hinze, Johan Jeuring, and Andres Löf. 2004. Type-indexed data types. *Science of Computer Programming* 51, 1-2 (2004),
1002 117–151.
- 1003 Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. 2012. Multivariate amortized resource analysis. *ACM Transactions on*
1004 *Programming Languages and Systems (TOPLAS)* 34, 3 (2012), 1–62.
- 1005 Jan Hoffmann and Martin Hofmann. 2010. Amortized resource analysis with polynomial potential: A static inference of
1006 polynomial bounds for functional programs. In *European Symposium on Programming*. Springer, 287–306.
- 1007 Jan Hoffmann and Steffen Jost. 2022. Two decades of automatic amortized resource analysis. *Mathematical Structures in*
1008 *Computer Science* 32, 6 (2022), 729–759.
- 1009 Rob R Hoogerwoord. 1992a. A logarithmic implementation of flexible arrays. In *International Conference on Mathematics of*
1010 *Program Construction*. Springer, 191–207.
- 1011 Rob R Hoogerwoord. 1992b. A symmetric set of efficient list operations (Functional pearl). *Journal of Functional Programming*
1012 2, 4 (1992), 505–513.
- 1013 Patrik Jansson and Johan Jeuring. 1997. PolyP—a polytypic programming language extension. In *Proceedings of the 24th*
1014 *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 470–482.
- 1015 Steffen Jost, Kevin Hammond, Hans-Wolfgang Loidl, and Martin Hofmann. 2010. Static determination of quantitative
1016 resource usage for higher-order programs. In *Proceedings of the 37th annual ACM SIGPLAN-SIGACT Symposium on*
1017 *Principles of Programming Languages (POPL)*. 223–236.
- 1018 Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Aleš Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground
1019 up: A modular foundation for higher-order concurrent separation logic. *Journal of Functional Programming* 28 (2018),
1020 e20.
- 1021 Hsiang-Shang Ko, Liang-Ting Chen, and Tzu-Chi Lin. 2022. Datatype-generic programming meets elaborator reflection.
1022 *Proceedings of the ACM on Programming Languages* 6, ICFP (2022), 225–253.
- 1023 John Launchbury. 1993. A natural semantics for lazy evaluation. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium*
1024 *on Principles of Programming Languages (POPL)*. 144–154.
- 1025 Daniel Le Métayer. 1988. ACE: An automatic complexity evaluator. *ACM Transactions on Programming Languages and*
1026 *Systems (TOPLAS)* 10, 2 (1988), 248–266.
- 1027 Xing Li, Yao Li, Peter Schachte, and Christine Rizkallah. 2026. The Memorist Tale: Every Thunk Every Cost All At Once.
1028 (2026). Accepted for publication at ESOP.
- 1029 Anton Lorenzen. 2025. Lightweight Testing of Persistent Amortized Time Complexity in the Credit Monad. In *Proceedings*
1030 *of the 18th ACM SIGPLAN International Haskell Symposium (Haskell '25)*. Association for Computing Machinery, New
1031 York, NY, USA, 80–93. <https://doi.org/10.1145/3759164.3759351>
- 1032 Anton Lorenzen, Daan Leijen, and Wouter Swierstra. 2023. FP²: Fully in-Place Functional Programming. *Proceedings of the*
1033 *ACM on Programming Languages* 7, ICFP (2023), 275–304.
- 1034 Anton Lorenzen, Daan Leijen, Wouter Swierstra, and Sam Lindley. 2024. The Functional Essence of Imperative Binary
1035 Search Trees. *Proceedings of the ACM on Programming Languages PLDI* (2024). <https://doi.org/10.1145/3656398>

- 1030 Anton Lorenzen, Daan Leijen, Wouter Swierstra, and Sam Lindley. 2025. First-Order Laziness. *Proc. ACM Program. Lang.* 9,
 1031 ICFP, Article 261 (Aug. 2025), 29 pages. <https://doi.org/10.1145/3747530>
- 1032 Conor McBride. 2001. The derivative of a regular type is its type of one-hole contexts. *Unpublished manuscript* (2001),
 1033 74–88.
- 1034 Glen Mével, Jacques-Henri Jourdan, and François Pottier. 2019. Time credits and time receipts in Iris. In *European Symposium
 1035 on Programming*. Springer, 3–29.
- 1036 Yasuhiko Minamide. 1998. A functional representation of data structures with a hole. In *Proceedings of the 25th ACM
 1037 SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 75–84.
- 1038 Tobias Nipkow. 2025. *Functional Data Structures and Algorithms: A Proof Assistant Approach*. ACM.
- 1039 Yue Niu, Jonathan Sterling, Harrison Grodin, and Robert Harper. 2022. A cost-aware logical framework. *Proceedings of the
 1040 ACM on Programming Languages* 6, POPL (2022), 1–31.
- 1041 Chris Okasaki. 1998. *Purely Functional Data Structures*. Cambridge University Press.
- 1042 Alexandre Pilkiewicz and François Pottier. 2011. The essence of monotonic state. In *Proceedings of the 7th ACM SIGPLAN
 1043 Workshop on Types in Language Design and Implementation*. 73–86.
- 1044 François Pottier, Armaël Guéneau, Jacques-Henri Jourdan, and Glen Mével. 2024. Thunks and Debits in Separation Logic
 1045 with Time Credits. *Proc. ACM Program. Lang.* 8, POPL (2024). <https://doi.org/10.1145/3632892>
- 1046 Vineet Rajani, Marco Gaboardi, Deepak Garg, and Jan Hoffmann. 2021. A unifying type-theory for higher-order (amortized)
 1047 cost analysis. *Proc. ACM Program. Lang.* 5, POPL (2021). <https://doi.org/10.1145/3434308>
- 1048 Martin Rem and Wim Braun. 1983. *A logarithmic implementation of flexible arrays*. Technical Report Memorandum MR83/4.
 1049 Eindhoven University of Technology.
- 1050 Mads Rosendahl. 1989. Automatic complexity analysis. In *Proceedings of the fourth international conference on Functional
 1051 programming languages and computer architecture*. 144–156.
- 1052 David Sands. 1990. *Calculi for time analysis of functional programs*. Ph.D. Dissertation. Imperial College London, UK.
- 1053 David Sands. 1995. A naive time analysis and its theory of cost equivalence. *Journal of Logic and Computation* 5, 4 (1995),
 1054 495–541.
- 1055 Jon Shultis. 1984. *Complexity of higher-order programs*. Technical Report. Colorado Univ., Boulder (USA). Dept. of Computer
 1056 Science.
- 1057 Matthieu Sozeau. 2010. Equations: A dependent pattern-matching compiler. In *International Conference on Interactive
 1058 Theorem Proving*. Springer, 419–434.
- 1059 Matthieu Sozeau and Cyprien Mangin. 2019. Equations reloaded: High-level dependently-typed functional programming
 1060 and proving in Coq. *Proceedings of the ACM on Programming Languages* 3, ICFP (2019), 1–29.
- 1061 Robert Endre Tarjan. 1985. Amortized computational complexity. *SIAM Journal on Algebraic Discrete Methods* 6, 2 (1985),
 1062 306–318.
- 1063 Paul van der Walt and Wouter Swierstra. 2012. Engineering Proof by Reflection in Agda. In *Implementation and Application of
 1064 Functional Languages - 24th International Symposium, IFL 2012, Oxford, UK, August 30 - September 1, 2012, Revised Selected
 1065 Papers (Lecture Notes in Computer Science)*, Ralf Hinze (Ed.), Vol. 8241. Springer, 157–173. https://doi.org/10.1007/978-3-642-41582-1_10
- 1066 Kathryn Van Stone. 2003. *A denotational approach to measuring complexity in functional programs*. Ph.D. Dissertation.
 1067 Carnegie Mellon.
- 1068 Ben Wegbreit. 1975. Mechanical program analysis. *Commun. ACM* 18, 9 (1975), 528–539.
- 1069 Li-yao Xia, Laura Israel, Maite Kramarz, Nicholas Coltharp, Koen Claessen, Stephanie Weirich, and Yao Li. 2024. Story of Your
 1070 Lazy Function’s Life: A Bidirectional Demand Semantics for Mechanized Cost Analysis of Lazy Programs. *Proceedings of
 1071 the ACM on Programming Languages* 8, ICFP (2024), 30–63.

```

1079 A Implementation of real time queues
1080 record Q (A : Set) : Set where
1081   constructor queue
1082   field
1083     front : Stream A
1084     rear  : Stream A
1085     sched : Stream A
1086     inv   : sched  $\sqsubseteq$  front
1088 open Q
1089 empty : Q A
1090 empty = queue [] [] []  $\sqsubseteq$ -refl
1091
1092 balance-enq : (front rear sched : Stream A)  $\rightarrow$  (sched  $\sqsubseteq$  front)  $\rightarrow$  Q A
1093 balance-enq .(c [ sched ]) rear sched (suffix c) with seq sched
1094 ... |  $\langle [] \rangle$  = let f = app (c [ [] ]) (rev rear [])
1095               in queue f [] f  $\sqsubseteq$ -refl
1096
1097 ... |  $\langle x : xx \rangle$  = let sched' = step-rev xx in
1098                   let c' = c + (x : []) in
1099                   queue (c' [ sched' ]) rear sched' (suffix c')
1100
1101 tail : Ctx A  $\rightarrow$  Ctx A
1102 tail [] = []
1103 tail (x : c) = c
1104
1105 balance-deq : (xx front rear sched : Stream A)  $\rightarrow$  (sched  $\sqsubseteq$  front)  $\rightarrow$  Q A
1106 balance-deq xx .(c [ sched ]) rear sched (suffix c) with seq sched
1107 ... |  $\langle [] \rangle$  = let f = app xx (rev rear [])
1108               in queue f [] f  $\sqsubseteq$ -refl
1109
1110 ... |  $\langle x : xx \rangle$  = let sched' = step-rev xx in
1111                   let c' = tail (c + (x : [])) in
1112                   queue (c' [ sched' ]) rear sched' (suffix c')
1113
1114 enq : Q A  $\rightarrow$  A  $\rightarrow$  Q A
1115 enq q y = balance-enq (front q) (y : rear q) (sched q) (inv q)
1116
1117 deq : Q A  $\rightarrow$  Maybe (A  $\times$  Q A)
1118 deq q with seq (front q)
1119 ... |  $\langle [] \rangle$  = nothing
1120
1121 ... |  $\langle x : xx \rangle$  = just (x , balance-deq xx (front q) (rear q) (sched q) (inv q))
1122
1123
1124
1125
1126
1127

```