

Purely Functional Real-Time Deques

Research Seminar Logic & Learning

Fabian Schaub
Supervisor: René Thiemann

16 February 2023

Abstract

Dequeues are a very useful and versatile data structure. In imperative programming they can be implemented as doubly linked lists. This yields a deque with constant worst case complexity for the operations *push*, *pop*, *inject* and *eject*. Also constant time catenation of two different doubly linked lists is possible. However, in purely functional programming we cannot make use of such an implementation. By only using *algebraic data types* the task of a deque with constant worst case complexities for the mentioned operations becomes non-trivial. In this report we describe and present an implementation of a purely functional real-time deque as defined by *Kaplan and Tarjan et al.* [5].

1 Introduction

A deque (short for **d**ouble **e**nded **q**ueue) is a data structure representing a sequence of elements. It supports at least four operations: `push`, `pop`, `inject` and `eject`. `push` and `pop` add and remove an element at the front of the deque, respectively. Likewise `inject` and `eject` add and remove an element at the back of the deque, respectively. We can go on and define catenation of deques in terms of `pop` and `inject`, or `eject` and `push`.

An example use-case for a deque is a string-builder. String-builders are commonly used to build up large strings from sub-strings because strings are often immutable in programming languages. A natural choice for an implementation of a string-builder is a deque over strings. Relative to simply catenating strings, using such a string-builder and extracting the resulting string in the end is very efficient, provided the underlying deque operations are efficient.

From this example we can see that an implementation of a deque with constant time complexity is of great interest. Moreover, in this report we are interested in a deque data structure with constant worst case complexity that is purely functional.

In the paper *purely functional real-time deques with catenation* [5] three data structures are presented: *real-time deques*, *real-time steques with catenation* and *real-time deques with catenation*. We consider *real-time deques* in this report and give a short outlook to the other two in the end.

In the following we reference related work and present the background of our work. Going on, we present the data structure from [5] and our implementation of it. In the end we show some results and finally we give an outlook to further work.

2 Related Work

Several purely functional implementations of deques achieved constant amortized complexity, which can be found in papers by Hood [2], Gajewska [1] and Hoogerwoord [3]. The implementation presented by *Kaplan and Tarjan et al.* in *Purely functional, real-time deques with catenation* [5] extends the previous implementations and achieves constant worst case complexity. They only present a textual explanation of the data structure and algorithm however. Our contribution is a pure implementation of data structure and operations in *Haskell* [4] using *algebraic data types*.

3 Background

As previously mentioned, we are in a purely functional setting. This means that our data structure has to be persistent [6], i.e. a modification of the data structure must not destroy the previous version. We consider *algebraic data types* (ADTs) as they are persistent and native to our implementation language.




Figure 1: A non-optimized deque representing the sequence 1 to 15. The top level deque has buffers over singleton elements and its child buffers over pairs of elements. the bottom of the deque is denoted as \emptyset .

The deque supports the operations `push`, `pop`, `inject` and `eject`. For these we need constant time complexity. Consider an implementation as a pair of lists $[2, 1, 3]$. This implementation typically uses a list splitting and reversal scheme. Amortized complexity analysis indeed gives us constant time complexity for all deque operations. Yet, in the worst case we have to perform list splitting and reversing which takes linear time. The deque implemented in this work achieves real-time performance, meaning all operations have a worst case complexity of $O(1)$.

4 A Purely Functional Real-Time Deque

Kaplan et al. state two key ideas to implement a purely functional real-time deque. First is to use the idea from the pair of lists mentioned in Section 3. Second is not to wait to do a list splitting and reversal, but rather to move elements inside the deque at every operation. This way the balancing of the deque, which is the essence of the list splitting and reversal scheme, is spread over all deque operations. More precisely, after each deque operation we balance the deque in a constant number of steps. As long as we retain a balanced deque we can perform all deque operations on it efficiently.

This balanced state is subsumed by *Kaplan's* definition of *regularity* and *semiregularity*. Informally, a *regular* deque is balanced and a *semiregular* deque is slightly imbalanced, such that it can be balanced in a constant number of steps.

In the following subsections we present the data structure, ideas and algorithms from [5] in Subsections 4.1, 4.2 and 4.3 and then show parts of our implementation in Section 4.4. The full implementation can be found in the Appendix.

4.1 Data Structure

As in the *pair-of-lists* approach mentioned in Section 3 the deque consists of two such lists called *buffers*. The front buffer is called *prefix* and the tail buffer is called *suffix*. Buffers are constrained to hold at most 5 elements. Since buffers themselves are deques implemented by lists this is needed to achieve real-time performance. More precisely, all deque operations on buffers are real-time since each buffer has an upper bound on its

size.

The *real-time deque* is then a triple of prefix, suffix and a child deque over pairs of elements (Figure 1). We call the deque in this chain of descendants levels, where level $i + 1$ is the descendant of level i . The idea behind the pairing of elements is a recursive slowdown when moving data inside the deque structure. Consider the levels i and $i + 1$. For every move operation on level i we move n elements. Subsequently, for every move operation on level $i + 1$ we move $2n$ elements. This means for x deque operations we have to perform up to $\frac{x}{2^{i+1}}$ move operations for any level i to balance the structure.

4.2 Regularity

For an efficient procedure to move data inside the deque structure, we first define a coloring of the buffers and the deque. The colors are *red*, *green* and *yellow* and we order them as $red < yellow < green$. The coloring indicates how many deque operations can be performed on an object. In the worst case, we cannot apply a deque operation on a *red* object. On a *yellow* object we apply at least one deque operation and on a *green* object we can apply at least two deque operations.

Buffers are then colored

- *red* if they contain 0 or 5 elements,
- *yellow* if they contain 1 or 4 elements or
- *green* if they contain 2 or 3 elements.

The number of operations applicable to a deque is the number of operations applicable to the „worse“ buffer. An exception occurs on the bottommost deque. For this deque we can move things between prefix and suffix without violating the element order. This means that if one of the buffers is empty the number of operations applicable to the bottommost deque is the number of operations applicable to the non-empty buffer. In conclusion, the coloring of a deque is defined as:

$$\text{color } d = \begin{cases} \text{color (prefix } d) & \text{if child } d \text{ and suffix } d \text{ are empty} \\ \text{color (suffix } d) & \text{if child } d \text{ and prefix } d \text{ are empty} \\ \min(\text{color (prefix } d), \text{color (suffix } d)) & \text{otherwise} \end{cases}$$

An example of a deque coloring can be seen in Figure 2. Note that Level 1 in this example would not be red if its suffix was empty, since it is the bottommost deque.

A deque is *semi-regular* if for all levels i and j , if level i and j are red then there exists a level k that is green and between i and j . More formally, let $child^m(d)$ denote level n




Figure 2: An example of a colored deque. Level 0 is green since both its buffers are green. Level 1 is red because its suffix is full.

in deque d . Semi-regularity is then defined as:

$$\begin{aligned}
 & d \text{ is semi-regular} \\
 & \quad \updownarrow \\
 & \forall i, j. i < j \rightarrow \text{child}^i(d) \text{ and } \text{child}^j(d) \text{ are red} \rightarrow \exists k. i < k < j \wedge \text{child}^k(d) \text{ is green}
 \end{aligned}$$

The idea behind semi-regularity is that any red deque can be transformed into a green one if we can move elements down or up in the deque structure. This is only the case if the descendant deque is not red. Furthermore, if there is no green descendant deque before the next red deque, such a transformation could lead to non-semi-regularity. This is why a green deque that interrupts a red chain is needed.

A deque is then *regular* if it is semi-regular and the first non-yellow deque from the top is green.

Example 4.1. Consider the color sequence of deque:

- *green* \rightarrow *yellow* \rightarrow *green* is regular.
- *yellow* \rightarrow *yellow* \rightarrow *yellow* is regular.
- *green* \rightarrow *yellow* \rightarrow *red* is regular.
- *red* \rightarrow *green* is semi-regular. It is not regular because it starts with a red deque.
- *red* \rightarrow *yellow* \rightarrow *red* is not semi-regular because there is no green deque between the red deques.

Example 4.2. The deques from Figure 2, 1 and 3 are all regular.

A deque being regular means that we can always perform at least one deque operation on the top level of the deque. Furthermore, the topmost red deque is eventually transformed into a green one before we need to perform an operation on it. Consider for example a regular deque on which we push elements. Assume a topmost red level i that has a prefix buffer with 5 elements. i is not violating the regularity constraint as long as there is a




Figure 3: **Left:** The optimized deque data structure. The left sub tree are the immediate yellow descendants, the right child node the non-yellow descendant.
Right: The inorder traversal over the optimized data structure.

green level above it. As we push elements onto the deque level i eventually becomes the topmost non-yellow deque. At this point i violates the regularity constraint and elements are moved down from its prefix to make i green. Since i has now again free capacity, any regularization that is needed above i (which may need to modify level i) is possible. The details of the regularization procedure is explained in Subsection 4.3.

An important observation is that a deque operation on a regular deque results in a regular or a semi-regular deque. Subsequently, if we regularize after each deque operation we can retain a regular deque.

At this point we see that we need to traverse to the topmost non-yellow level for our regularization. A deque structure as depicted in Figure 1 poses a problem for this. We would need to traverse the chain of descendants which has a complexity of $O(\log(n))$. Luckily, a structural optimization mitigates this problem. As seen in Figure 3 we can use a tree structure where the left sub-tree is the chain of immediate yellow descendants and the right child is the next non-yellow deque. Using this structure, the top-most non-yellow deque is the root or its right child. Moving on, we will disregard this optimization as it does not influence the principle of the regularization procedure.

4.3 Regularization

As previously mentioned we need to restore regularity of a semi-regular deque in constantly many steps after each deque operation. This is done by changing the top most red deque to a green deque. With the optimization of Figure 3 we can find this red deque within

one or two steps.

Let level i be the topmost red deque and level $i + 1$ its child. Furthermore, let P_i and S_i denote the prefix and the suffix of level i , respectively. Likewise, P_{i+1} and S_{i+1} denote the prefix and suffix of level $i + 1$. We distinguish three cases based on how many elements are in each of these buffers. In the following we will call buffers with one or zero elements *underflowing* and buffers with four or five elements *overflowing*.

Two-Buffer-Case ($|P_{i+1}| + |S_{i+1}| \geq 2$): this case occurs if level $i + 1$ contains at least two elements so we can move elements down or up to regularize. If $i + 1$ is the bottommost level we move elements between P_{i+1} and S_{i+1} , such that both are non-empty. Next we move elements down if P_i or S_i or both are overflowing. In this case level $i + 1$ contains at least two elements, enough to move elements up if P_i , S_i or both are underflowing. Finally, if level $i + 1$ is empty we delete it.

One-Buffer-Case ($|P_{i+1}| + |S_{i+1}| \leq 1 \wedge (|P_i| \geq 2 \vee |S_i| \geq 2)$): this case occurs solely at the bottom of the deque. It occurs if level $i + 1$ contains at most one element and at least one of P_i and S_i contains two or more elements. We handle this case similar as the Two-Buffer-Case, but we use only P_{i+1} on level $i + 1$. We move the element on level $i + 1$ to P_{i+1} if it exists. Next we move elements down to P_{i+1} from overflowing buffers on level i . Then we move elements up from P_{i+1} to underflowing buffers on level i . Finally, we delete level $i + 1$ if it is empty.

No-Buffer-Case ($|P_{i+1}| + |S_{i+1}| \leq 1 \wedge |P_i| \leq 1 \wedge |S_i| \leq 1$): this case also occurs solely at the bottom of the deque. In contrast to the One-Buffer-Case, here P_i and S_i are both underflowing. This means we have at most three level- i -elements in total. We move all the elements to P_i and delete level $i + 1$.

After this procedure level i is green and regularity is restored.

Example 4.3. Consider the deque d in Figure 2. The coloring is $color(d) = green$ and $color(child(d)) = red$, so d is regular. Assume we inject an element 16 into d . Then the top suffix gets yellow which makes the top level yellow. Now the topmost non-yellow level is level 1, which is red. Subsequently $inject(d, 16)$ is semi-regular. According to the regularization procedure we are in the One-Buffer-Case. Regularization now moves two elements from the suffix of level 1 to the prefix of level 1, which results in two green buffers and a green level 1. In the end we obtain $d' = regularize(inject(d, 16))$. Its coloring is $color(d') = yellow$ and $color(child(d')) = green$, which means d' is regular.

4.4 Implementation

There are several data types which can be used as a deque. For this reason `Deque` is a type class in our implementation. Other deque operations which are not class functions like `append` and `fromList` are implemented on top of these defined class functions.

```

-- | a deque as a class.
class Deque (a :: * -> *) where
  push :: b -> a b -> a b
  pop  :: a b -> (b, a b)
  inject :: b -> a b -> a b
  eject :: a b -> (b, a b)
  null  :: a b -> Bool
  empty :: a b

```

As the basic container of elements we use lists to realize buffers. A buffer is an instance of the `Deque` type class since it can act as one and indeed is used as one in this implementation. Note that buffers may contain at most five elements. This detail is not captured in the buffer type but asserted at runtime.

```

type Buffer = []
instance Deque Buffer where
  ...

```

Our real-time deque `RTDeque` is implemented as a record type. Due to the structural optimization (Figure 3) we have to make a slight modification to the data type. Since we do not know how many immediate yellow descendants there are, we also do not know the degree of pairing in the non-yellow branch. In Haskell we have no means to encode this relationship, so we have to default to arbitrary nesting. For this reason we use buffers of balanced binary trees and do not encode the pairing in the type. This means the type `RTDeque a` does not contain information about the degree of pairing, nor does its constructors contain information about the correctness of the pairing.

Also note that we have an additional constructor `Nil` which is the empty bottom of a deque. In contrast to a deque with empty buffers `Nil` does not denote an underflowing deque but is really the bottom element.

```

data RTDeque a
  = Nil
  | RTDeque
  { _prefix :: Buffer (BTree a)
  , _suffix :: Buffer (BTree a)
  , _yellows :: RTDeque a      -- ^the yellow descendants.
                                -- these must not contain nonyellow branches.
  , _nonyellows :: RTDeque a  -- ^the nonyellow descendants.
  }

```

The notion of colors is captured in our implementation as the data type `Color` and the type class `Colored`. `Buffer` and `RTDeque` are instances of the `Colored` class.

```

data Color = Red | Yellow | Green
deriving (Ord, Eq, Show)

```



```

class Colored (a :: *) where
  color :: a -> Color

instance Colored Buffer where
  ...
instance Colored RTDeque where
  ...

```

The function `restoreRegularity` uses this data definitions to implement the regularization procedure. In the source code are also checks of *regularity* and *semi-regularity* implemented. These are only for completeness and not used in the deque operations.

`restoreRegularity` is implemented as defined in the paper and is not presented at this point because it is quite long and implements the procedure shown in Subsection 4.3 or initially presented in [5]. Interested readers can find it in the Appendix. An important observation is that the regularization procedure does not contain recursive functions. The only exceptions are the functions on buffers. As we mentioned earlier buffers contain at most five elements, so these are indeed $O(1)$. Additionally, when we look at the regularization cases it becomes clear that the original description relies heavily on sequential thinking. For sake of consistency we implemented these in the same way but a non-sequential-style implementation may provide a simpler regularization function.

Finally, the instance of `Deque RTDeque` is derived by performing an operation on the appropriate top level buffer and regularizing the resulting deque.

5 Results

A measurement of `toList` and `fromList` (Figure 4) gives us the expected result. Since the two operations are implemented using `eject` and `inject`, respectively, we perform n deque operations with a list of length n . This means we expect a linear increase in runtime, which we can see in the measurement. We can also see a somewhat discrete step size in the increase in runtime. These „levels“ may correspond to the performed regularization case.

In Figure 5 we see the performance of `inject` and `eject`. These measurements match the measurements from Figure 4. The high runtime of the small deques is an effect due to our measurement setup. Included in these run times is the program start, which has more effect on smaller deque sizes. We are confident, that a repeated measurement with a more accurate setup would give us a runtimes like we see with larger deques. Unfortunately, we couldn't do these measurements for time reasons.

6 Conclusion

We were able to implement the deque presented by *Kaplan et al.* in Haskell and confirm the real-time behavior experimentally. The biggest challenge was to bring the purely




Figure 4: runtime of `toList` . `fromList` with lists of length zero to 10000.




Figure 5: runtime of `inject` and `eject` on deques of length zero to 10000.

textual explanation to a level that allowed us to implement the data structure and regularization procedure. Subsequently, our goal was not to provide a production-ready high performance but rather a conceptual and concise implementation. Though we have a working implementation there are some things which can be further simplified, for example the regularization cases and the handling of the deque shape.

In the process of bringing the data structure to code, we also implemented a monadic version that counts the cost of each deque operation. Since we saw later on that we can actually implement everything without recursion, the question of complexity became trivial and we moved to the now implemented pure version.

7 Further Work

In the paper this report is based on [5] there are two more data structures: *purely functional real-time steques with catenation* and *purely functional real-time deques with catenation*. The overall goal is to implement and formalize all three data structures. This would be a step to make them available to interactive theorem provers and dependently typed languages such as *Agda*, *Coq*, *Isabelle/HOL* and *Lean*, as well as simply typed functional languages like *Haskell* and *OCaml*. Moreover, dependent typing may give us the opportunity to index on the nested pairs in the deque datatype and refine our current solution with arbitrary binary trees. To which degree this is possible is an open question for us.

Moving away from pure functional programming to imperative programming the third data structure, the deque with real-time catenation, may also be of use. It solves the problem of self-catenation since it is persistent. In implementations like doubly linked lists self catenation may introduce non-termination if not handled properly because it leads to circular traversal. Although, one must also mention that despite being real-time these purely functional implementations are not as straight forward and performant as imperative ones. This is a downside of using deques based on persistent data structures in imperative programming. The gain of safety and the ability to self-catenation however could be useful in some domains.

8 Acknowledgements

I thank my supervisor Professor Dr. René Thiemann for helping me in conducting the project and producing the resulting material.

References

- [1] Hania Gajewska and Robert E Tarjan. Deques with heap order. *Information Processing Letters*, 22(4):197–200, 1986.
- [2] Robert T Hood. The efficient implementation of very-high-level programming language constructs. Technical report, Cornell University, 1982.

- [3] Rob R Hoogerwoord. Functional pearls a symmetric set of efficient list operations. *Journal of Functional Programming*, 2(4):505–513, 1992.
- [4] Simon Peyton Jones. *Haskell 98 language and libraries: the revised report*. Cambridge University Press, 2003.
- [5] Haim Kaplan and Robert E Tarjan. Purely functional, real-time deques with catenation. *Journal of the ACM (JACM)*, 46(5):577–603, 1999.
- [6] Neil Ivor Sarnak. *Persistent data structures*. New York University, 1986.

Appendix

Source Code

```
1
2 {-# LANGUAGE KindSignatures #-}
3 {-# LANGUAGE NoImplicitPrelude #-}
4 {-# LANGUAGE TypeSynonymInstances #-}
5
6
7
8
9 module Deque
10 where
11
12
13 import Control.Monad ((>=>), (<=<), (=<<), join)
14 import Control.Monad.Identity
15 import Control.Arrow (first, second, (>>>))
16 import Data.Bifunctor (bimap)
17 import Data.Functor ((<&>))
18 import Data.Foldable (foldl')
19 import Data.Function ((&))
20
21 import Prelude hiding (null)
22 import qualified Prelude
23
24
25 {-|
26 Implementation of
27   "Purely Functional, Real-Time Deques with Catenation, KAPLAN and
28   ↪ TARJAN et. al.,
29   Journal of the ACM, Vol. 46, No. 5, September 1999, pp. 577-603."
30
31 This module implements real-time deques without catenation.
32 we ensure the real-time complexity by restricting ourselves mainly to
33   ↪ non-recursive functions.
34 If we need to use recursive functions, we add a justification why this
35   ↪ function is still constant time or why it has no effect
36 on the behavior of the real-time deque.
37 -|
38
39 -- * Deques
```

```

38
39 -- / a deque as a class.
40 class Deque (a :: * -> *) where
41   push :: b -> a b -> a b
42   pop  :: a b -> (b, a b)
43   inject :: b -> a b -> a b
44   eject :: a b -> (b, a b)
45   null  :: a b -> Bool
46   empty :: a b
47   append :: a b -> a b -> a b
48   append xs ys = foldl' (flip inject) ys $ toList xs
49   len :: a b -> Int
50   len = length . toList
51
52
53
54 -- ** Deque packing and unpacking
55 --
56 -- note that fromList and toList are not scope of the realtime
57 -- ↪ behaviour.
58
59 toList :: (Deque a) => a b -> [b]
60 toList = go []
61   where
62     go acc d
63       | null d      = acc
64       | otherwise = uncurry go $ first (:acc) $ eject d
65
66 fromList :: Deque a => [b] -> a b
67 fromList xs = do
68   foldl' (flip inject) empty xs
69
70 -- * nonempty, leaf only binary trees
71
72 data BTree a = Leaf !a | Node (BTree a) (BTree a)
73
74 instance Show a => Show (BTree a) where
75   show (Node x y) = "(" ++ show x ++ ", " ++ show y ++ ")" -- ~this
76   ↪ method is not used in the deque
77   show (Leaf v) = show v
78
79 -- ** convenience functions

```

```

79
80 combine :: BTree a -> BTree a -> BTree a
81 combine = Node
82
83 split :: BTree a -> (BTree a, BTree a)
84 split (Leaf _) = error "split: cannot split leaf"
85 split (Node l r) = (l, r)
86
87 unleaf :: BTree a -> a
88 unleaf (Node _ _) = error "unleaf: cannot unpack node"
89 unleaf (Leaf x) = x
90
91
92 -- * Colors
93
94 -- | Colors as defined in the paper
95 data Color = Red | Yellow | Green
96   deriving (Ord, Eq, Show)
97
98 -- | Class for colored things
99 class Colored (a :: *) where
100   color :: a -> Color
101
102 -- * Buffers
103
104 -- | we use lists as Buffers. note that they must fit at most 5
105   ↪ elements.
106 -- subsequently, if we use non-constant linear functions in this context
107   ↪ they become constant time.
108 type Buffer = []
109
110 instance Deque Buffer where
111   push = (:)
112   pop (x:xs) = (x, xs)
113   inject x xs = xs ++ [x]
114   eject xs = (last xs, take (length xs - 1) xs)
115   null = Prelude.null
116   empty = []
117
118 instance Colored (Buffer a) where
119   color buf = case len buf of
120     0 -> Red

```

```

120     1 -> Yellow
121     2 -> Green
122     3 -> Green
123     4 -> Yellow
124     5 -> Red
125     _ -> error "color: buffer overflow"
126
127
128 -- * Real time Deques
129
130 -- | the real time deque.
131 data RTDeque a
132   = Nil
133   | RTDeque
134     { _prefix :: Buffer (BTree a)
135     , _suffix :: Buffer (BTree a)
136     , _yellows :: RTDeque a      -- ^the yellow descendants. these must
137     ↪ not contain nonyellow branches.
138     , _nonyellows :: RTDeque a  -- ^the nonyellow descendants.
139     }
140   deriving (Show)
141
142 -- | an empty instance of a RTDeque
143 emptyRTDeque :: RTDeque a
144 emptyRTDeque = RTDeque
145   { _prefix = empty
146   , _suffix = empty
147   , _yellows = Nil
148   , _nonyellows = Nil
149   }
150
151 -- ** guarded getters, modifiers and setters
152
153 prefix, suffix :: RTDeque a -> Buffer (BTree a)
154 prefix Nil = empty
155 prefix d   = _prefix d
156 suffix Nil = empty
157 suffix d   = _suffix d
158
159 yellows, nonyellows :: RTDeque a -> RTDeque a
160 yellows Nil         = Nil
161 yellows d           = _yellows d

```



```

162 nonyellows Nil = Nil
163 nonyellows d  = _nonyellows d
164
165 onPrefix, onSuffix :: (Buffer (BTree a) -> Buffer (BTree a)) -> RTDeque
    ↪ a -> RTDeque a
166 onPrefix f Nil = emptyRTDeque { _prefix = f empty }
167 onPrefix f deq = deq { _prefix = f (_prefix deq) }
168 onSuffix f Nil = emptyRTDeque { _suffix = f empty }
169 onSuffix f deq = deq { _suffix = f (_suffix deq) }
170 setPrefix, setSuffix :: Buffer (BTree a) -> RTDeque a -> RTDeque a
171 setPrefix = onPrefix . const
172 setSuffix = onSuffix . const
173
174
175 onNonyellows, onYellows :: (RTDeque a -> RTDeque a) -> RTDeque a ->
    ↪ RTDeque a
176 onNonyellows f Nil = emptyRTDeque { _nonyellows = f emptyRTDeque }
177 onNonyellows f deq = deq { _nonyellows = f $ _nonyellows deq }
178 onYellows f Nil = emptyRTDeque { _yellows = f emptyRTDeque }
179 onYellows f deq = deq { _yellows = f $ _yellows deq }
180
181 setNonyellows, setYellows :: RTDeque a -> RTDeque a -> RTDeque a
182 setNonyellows = onNonyellows . const
183 setYellows    = onYellows . const
184
185 -- ** auxiliary functions
186
187 -- | check if a RTDeque is a bottom element (i.e does not contain any
    ↪ values).
188 -- due to our regularization procedure we know that a deque is bottom if
    ↪ the buffers are empty.
189 bottom :: RTDeque a -> Bool
190 bottom Nil = True
191 bottom deq = all null [prefix deq, suffix deq]
192
193
194 -- | replace RTDeque with Nil if it is the bottom of the deque
195 truncate :: RTDeque a -> RTDeque a
196 truncate deq
197   | bottom deq = Nil
198   | otherwise  = deq
199
200

```

```

201
202 -- / apply a function on the descendant of the RTDeque
203 withNext :: (RTDeque a -> b) -> RTDeque a -> b
204 withNext fun deq
205   | bottom (yellows deq) = fun $ nonyellows deq
206   | otherwise            = fun $ yellows deq
207
208 instance Colored (RTDeque a) where
209   color deq =
210     let
211       bot = withNext bottom deq
212       pnull = null $ prefix deq
213       snull = null $ suffix deq
214     in
215       case (bot, pnull, snull) of
216         (True, True, _ ) -> color $ suffix deq
217         (True, _ , True) -> color $ prefix deq
218         (_ , _ , _ ) -> min (color $ prefix deq) (color $ suffix
↪ deq)
219
220
221 -- * Regularity
222 --
223 -- The regularity checks are implemented for completeness. they are not
↪ used in the actual RTDeque since regularity and semiregularity is
224 -- invariant to the functions. This also means that the checks have no
↪ effect on the complexity.
225
226
227 -- / semiregularity check
228 -- note that we also check if the partition (yellows/nonyellows) is
↪ correct (is an error case)
229 semiregular :: RTDeque a -> Bool
230 semiregular deque = bottom deque
231   || ( semiregular (nonyellows deque)
232       && allyellow (yellows deque)
233       && ( Red /= color deque || greenBeforeRed (nonyellows deque)
234         )
235     )
236 where
237   greenBeforeRed d = case (bottom d, color d) of
238     (False, Red ) -> False

```

```

239     (_ , Yellow) -> error "semiregular: yellows deque in the
↳ nonyellows stack"
240     (_ , _ ) -> True
241     allyellow d = bottom d
242     || ( color d == Yellow
243         && bottom (nonyellows d)
244         && allyellow (yellows d)
245         )
246     || error "semiregular: nonyellows deque in yellows stack"
247
248
249 -- / regularity check
250 regular :: RTDeque a -> Bool
251 regular deque
252 | bottom deque           = True
253 | not (semiregular deque) = False
254 | otherwise              = case (color deque, bottom (nonyellows
↳ deque), color (nonyellows deque)) of
255   (Green, _ , _ ) -> True
256   (Red , _ , _ ) -> False
257   (_ , True, _ ) -> True
258   (_ , _ , Green) -> True
259   (_ , _ , _ ) -> False
260
261
262 -- * Restoring Regularity
263
264
265 -- / restore a semiregular deque to a regular deque
266 --
267 -- this function is constant time, since it is not recursive and does
↳ not use any recursive functions.
268 restoreRegularity :: RTDeque a -> RTDeque a
269 restoreRegularity deque
270 | bottom deque           = deque
271 | color deque == Green   = deque
272 | color deque == Red     =
↳ withNext (restore deque) deque
273 | color deque == Yellow && color (nonyellows deque) == Green = deque
274 | color deque == Yellow && color (nonyellows deque) == Red   =
↳ onNonyellows (\ny -> withNext (restore ny) ny) deque
275 | otherwise              = error
↳ "restoreRegularity: nonexhaustive matching"

```

```

276 where
277   restore deque child =
278     let
279       p1 = len $ prefix deque
280       s1 = len $ suffix deque
281       p2 = len $ prefix child
282       s2 = len $ suffix child
283       two_buffer_case_cond = p2 + s2 >= 2
284       one_buffer_case_cond = p2 + s2 <= 1 && (p1 >= 2 || s1 >= 2)
285       no_buffer_case_cond  = p2 + s2 <= 1 && (p1 <= 1 && s1 <= 1)
286       case_fun = case (two_buffer_case_cond, one_buffer_case_cond,
↪ no_buffer_case_cond) of
287         (True, _ , _ ) -> twoBufferCase
288         (_ , True, _ ) -> oneBufferCase
289         (_ , _ , True) -> noBufferCase
290         (_ , _ , _ ) -> error "restoreRegularity: no case
↪ applicable"
291     in
292       combineDeqs $ case_fun buffers
293     where
294       -- the buffers to modify
295       buffers = (prefix deque, suffix deque, prefix child, suffix
↪ child)
296       -- moving elements between buffers
297
↪ p1_to_p2,p2_to_p1,p2_to_s2,p2_to_s1,s2_to_p2,s2_to_s1,s1_to_s2,s1_to_p2
↪ :: (Buffer (BTree a), Buffer (BTree a), Buffer (BTree a), Buffer
↪ (BTree a)) -> (Buffer (BTree a), Buffer (BTree a), Buffer (BTree a),
↪ Buffer (BTree a))
298       p1_to_p2 (p1, s1, p2, s2) = let (y,(x,p1')) = second eject $
↪ eject p1 ; p2' = push (combine x y) p2 in (p1', s1 , p2' , s2 )
299       p2_to_p1 (p1, s1, p2, s2) = let ((x,y),p2') = first split $ pop
↪ p2 ; p1' = inject y $ inject x p1 in (p1', s1 , p2' , s2 )
300       p2_to_s2 (p1, s1, p2, s2) = let (x,p2') = eject p2
↪ ; s2' = push x s2 in (p1 , s1 , p2' , s2')
301       p2_to_s1 (p1, s1, p2, s2) = let ((x,y),p2') = first split $
↪ eject p2 ; s1' = push x $ push y s1 in (p1 , s1' , p2' , s2 )
302       s2_to_p2 (p1, s1, p2, s2) = let (x,s2') = pop s2
↪ ; p2' = inject x p2 in (p1 , s1 , p2' , s2')
303       s2_to_s1 (p1, s1, p2, s2) = let ((x,y),s2') = first split $
↪ eject s2 ; s1' = push x $ push y s1 in (p1 , s1' , p2 , s2')
304       s1_to_s2 (p1, s1, p2, s2) = let (x,(y,s1')) = second pop $ pop
↪ s1 ; s2' = inject (combine x y) s2 in (p1 , s1' , p2 , s2')

```

```

305     s1_to_p2 (p1, s1, p2, s2) = let (x,(y,s1')) = second pop $ pop
↪ s1     ; p2' = inject (combine x y) p2 in (p1 , s1' , p2' , s2 )
306     -- combine deque and child with the new buffers.
307     -- the parent deque has changed from red to green.
308     -- we need to rotate the tree if child changed from yellow to
↪ nonyellow or vice versa.
309     combineDeqs (p1, s1, p2, s2) =
310         let
311             child' = truncate $ setPrefix p2 $ setSuffix s2 child
312             deque' = truncate $ setPrefix p1 $ setSuffix s1 deque
313         in
314             case (color child, color child') of
315                 (Yellow, Yellow) -> setYellows child' deque'
316                 (Yellow, _      ) -> let child'' = truncate $ setNonyellows
↪ (nonyellows deque') child'
317                                     in setYellows Nil $ setNonyellows
↪ child'' deque'
318                 (_      , Yellow) -> setYellows child' $ setNonyellows
↪ (nonyellows child') deque'
319                 (_      , _      ) -> setNonyellows child' deque'
320     -- two buffer case: p2 + s2 >= 2
321     twoBufferCase =
322         let
323             balance_lower                (p1, s1, p2, s2)
324             | len p2 == 0                = s2_to_p2 (p1, s1, p2, s2)
325             | len s2 == 0                = p2_to_s2 (p1, s1, p2, s2)
326             | otherwise                  = (p1, s1, p2, s2)
327             prop_prefix_down            (p1, s1, p2, s2)
328             | len p1 >= 4                = p1_to_p2 (p1, s1, p2, s2)
329             | otherwise                  = (p1, s1, p2, s2)
330             prop_prefix_up              (p1, s1, p2, s2)
331             | len p1 <= 1                = p2_to_p1 (p1, s1, p2, s2)
332             | otherwise                  = (p1, s1, p2, s2)
333             prop_suffix_down            (p1, s1, p2, s2)
334             | len s1 >= 4                = s1_to_s2 (p1, s1, p2, s2)
335             | otherwise                  = (p1, s1, p2, s2)
336             prop_suffix_up              (p1, s1, p2, s2)
337             | len s1 <= 1                = s2_to_s1 (p1, s1, p2, s2)
338             | otherwise                  = (p1, s1, p2, s2)
339         in
340             balance_lower
341             >>> prop_prefix_down
342             >>> prop_suffix_down

```

```

343     >>> prop_prefix_up
344     >>> prop_suffix_up
345     -- one buffer case: p2 + s2 <= 1 &&& (p1 >= 2 || s1 >= 2)
346     oneBufferCase =
347     let
348         move_lower_to_prefix      (p1, s1, p2, s2)
349         | len s2 == 1 = s2_to_p2 (p1, s1, p2, s2)
350         | otherwise   =          (p1, s1, p2, s2)
351     prop_prefix_down      (p1, s1, p2, s2)
352         | len p1 >= 4 = p1_to_p2 (p1, s1, p2, s2)
353         | otherwise   =          (p1, s1, p2, s2)
354     prop_prefix_up      (p1, s1, p2, s2)
355         | len p1 <= 1 = p2_to_p1 (p1, s1, p2, s2)
356         | otherwise   =          (p1, s1, p2, s2)
357     prop_suffix_down    (p1, s1, p2, s2)
358         | len s1 >= 4 = s1_to_p2 (p1, s1, p2, s2)
359         | otherwise   =          (p1, s1, p2, s2)
360     prop_suffix_up      (p1, s1, p2, s2)
361         | len s1 <= 1 = p2_to_s1 (p1, s1, p2, s2)
362         | otherwise   =          (p1, s1, p2, s2)
363     in
364         move_lower_to_prefix
365         >>> prop_prefix_down
366         >>> prop_suffix_down
367         >>> prop_prefix_up
368         >>> prop_suffix_up
369     -- no buffer case: p2 + s2 <= 1 &&& (p1 <= 1 &&& s1 <= 1)
370     noBufferCase =
371     let
372         move_s2_to_p2      (p1, s1, p2, s2)
373         | len s2 == 1 = s2_to_p2 (p1, s1, p2, s2)
374         | otherwise   =          (p1, s1, p2, s2)
375     prop_prefix_up      (p1, s1, p2, s2)
376         | len p2 == 1 = p2_to_p1 (p1, s1, p2, s2)
377         | otherwise   =          (p1, s1, p2, s2)
378     in
379         move_s2_to_p2 >>> prop_prefix_up
380
381
382
383     -- / the deque instance for Real Time deque
384     instance Deque RTDeque where
385         push x = restoreRegularity . onPrefix (push (Leaf x))

```

```

386   pop deq =
387     let
388       pnull = null (prefix deq) -- if the prefix is empty we pop the
↪   suffix
389       buf = if pnull then suffix deq else prefix deq
390       (x, buf') = pop buf
391       deq' = restoreRegularity $ (if pnull then setSuffix else
↪   setPrefix) buf' deq
392       in (unleaf x, deq')
393   inject x = restoreRegularity . onSuffix (inject (Leaf x))
394   eject deq =
395     let
396       snull = null $ suffix deq -- if the suffix is empty we eject the
↪   prefix
397       buf = if snull then prefix deq else suffix deq
398       (x, buf') = eject buf
399       deq' = restoreRegularity $ (if snull then setPrefix else
↪   setSuffix) buf' deq
400       in (unleaf x, deq')
401   null = bottom
402   empty = Nil

```