

COMS 4995 W: Parallel Functional Programming

Optimizing Parallel Sorting

Jay Karp, Benjamin Flin - jlk2225, brf2117

December 23, 2020

1 Abstract

This paper presents an exploration into three main sequential sorting algorithms and their parallel counterparts. These include merge sort, quicksort, as well as bitonic sort. While bitonic sort is an inherently parallel algorithm, merge sort and quicksort on their own are not conventionally parallelized. In this paper we attempt to demonstrate both the potential performance gains of parallelized sorting in Haskell, as well as the cleaner, easier to write and simpler code required to write these parallelizations in a purely functional programming language. Finally, we present our own hybrid sorting algorithm which is built upon the best features from the different sorts which we experimented with.

2 Introduction and Problem Statement

Sorting is one of the most fundamental concepts in computer science. It is one of the first topics many introductory classes teach to new computer science students, and yet at the same time it is so foundational that many of the systems we rely on daily are built on the concept of fast, efficient sorting. This includes things such as web indexing, supercomputing, and database management.

With the possible slowing of Moore's law in the near future, it is also essential that we look at alternative methods to push these important algorithms even further. Because of this we decided to implement three main sorting algorithms using Haskell: merge sort, quicksort, and bitonic sort. Although it would probably be faster to write these sorts in a lower level language like C, we believe that the functional paradigm allows for much easier and more concise data parallelism.

After moving these three main sorts from their equivalent sequential and imperative counterparts, we further pushed the limits of sorting in Haskell to implement some of the different data parallelism models. Finally, we present 'hybrid' sort, which is composed of the best features from all of the sequential, and parallel sorts that we implemented.

3 Implementation

In this section we will discuss the specific design patterns and implementations which we used in our sorting implementations. This includes the general Haskell data structures which we used, as

well as a discussion of both the sequential and parallel algorithms which we implemented. Along with this we will provide analysis of these algorithms imperative counterparts and justifications for our design decisions.

3.1 Haskell Data Structures

In our initial testing, our implementations were based on haskell's default lists. However, we quickly realized that all of our implementations would be better suited by a different data structure. This led us towards multiple alternatives, such as regular `Data.Arrays` and `Data.Array.REPA`. Ultimately we decided on implementing all of our solutions using Haskell's `Data.Vector`, as it provided us with a boxed data structure that has both mutable and immutable variants as well as $O(1)$ indexing.

The other main data structure that we used during our testing is the `instance NFData Dumb` data structure. This data structure simply contained an int, but used a comparison function with a time cost proportional to its value. This allowed us to test the granularity of our algorithms to some extent, as we could simulate a large load without many elements in our unsorted sequence.

3.2 Sequential

After deciding on our use of vectors, our first step was to convert the pseudocode of sequential imperative code into sequential functional code in haskell. To do this we followed the following pseudocode for Merge Sort, Quicksort, and Bitonic Sort.

<pre> MergeSort(arr, left, right) if (left > right) do end mid = (left+right)/2 MergeSort(arr, left, mid) MergeSort(arr, mid+1, right) Merge(arr, left, mid, right) end </pre>	<pre> QuickSort(arr, low, high) if (low < high) do pi = partition(arr, low, high) QuickSort(arr, low, pi - 1) QuickSort(arr, pi + 1, high) end end </pre>	<pre> BitonicSort(arr, low, cnt, dir) if (cnt > 1) do k = cnt / 2 BitonicSort(a, low, k, 1) BitonicSort(a, low+k, k, 0) BitonicMerge(a,low, cnt, dir) end end </pre>
---	--	---

While this step was not extremely challenging, it did come with its own problems. The first problem that we ran into was efficient memory use. Our initial implementations made use of Haskell's `Data.Vector.slice`, which creates a copy of a subset of an initial vector. We very quickly realized that this was creating an excessive number of vector copies and using an extreme amount of memory. This was substantially hindering performance and pushed us to further look into `Data.Vector.Mutable`. This exposed multiple extremely useful functions for in place data modification, specifically,

```

(1) read  :: PrimMonad m => MVector (PrimState m) a -> Int    -> m a
(2) write :: PrimMonad m => MVector (PrimState m) a -> Int    -> a    -> m ()

```

```
(3) modify :: PrimMonad m => MVector (PrimState m) a -> (a -> a) -> Int -> m ()
(4) swap   :: PrimMonad m => MVector (PrimState m) a -> Int     -> Int -> m ()
```

This provided us with the lower level immutable operations, such as $O(1)$ reading as well as writing and swapping vector values in place.

One change we made to the sequential versions of the algorithms was a small change to the default implementation of Merge Sort. While looking at the Haskell `Data.List.sort` implementation, which is a Merge Sort, we realized that they check for subsections of the list that are already sorted in ascending or descending order. If they are in ascending order, then the algorithm skips the rest of that subsection of the sort. If they are in descending order, then flip all of the elements in the vector. This saves a surprising amount of time and also makes best case sorts even faster.

The only other aspect of the sequential sorts which we had to implement was a function `fillBitonic` (`fillBitonic :: a -> V.Vector a -> V.Vector a`). We came to the realization that because of the parallel nature of the bitonic sorting network we needed to fill our input to our sequential bitonic sort with enough elements so that the size of the sorted vector is a power of 2. To do this we filled the sorted array with empty min values (such as `0` or `""`).

3.3 Parallel

Now that we had the sequential implementations finished, we could focus more specifically on the parallel implementations and optimizations. Most of our parallelism was done using `Control.Parallel.Strategies`, with the exception of `bitonicPar`, which used `Control.Monad.Par`.

3.3.1 Merge Sort

For our parallel merge sort implementation, we split our list into each into a sequence of sorted Vectors sequentially, and then recursively merge all of these Vectors by diving our sequence of Vectors in half and merging each half in parallel. For the top few levels we ran a special merge procedure called compare-exchange described below:

1. Create two tasks (sparks), each containing both lists to merge
2. Have one task compute the lower half of the merged lists, and the other compute the upper half.
3. Concatenate the resulting halves from each task into one list

This procedure was advantageous in the first few levels in our recursive merge, where there are many elements and less can be done in parallel. When merging was done deeper in the call stack, since there were less elements to merge per spark, but many more sparks, so this procedure swapped with our fast sequential one.

3.3.2 Quick Sort

Running quicksort in parallel involved use of the following procedure:

1. Split the sequence evenly among n tasks.
2. Choose a pivot and give it to each task.
3. Have each task partition its sequence into a pair of lists, one with elements less than the pivot, and the other with elements greater than the pivot.
4. Split the tasks into two halves. Have the first half of the tasks give their upper list to the second half and vice versa.
5. Concatenate the lists so we end up with n lists, with $\frac{n}{2}$ lists less than the pivot and $\frac{n}{2}$ greater than the pivot.
6. Recurse on each half in parallel and then concatenate the lists.
7. After $\log(n)$ recursions, our base case is to use sequential sort.

Running this procedure in parallel gave good results, however, some time was lost in steps 3 and 4, as well as some time lost in concatenation.

3.3.3 Bitonic Sort

Bitonic sort was the only sort where we used `Control.Monad.Par.IO` instead, since we needed to do in-place swaps in memory, otherwise much time would be lost in copying and memory management associated with Vectors. The only way to do these in-place swaps in parallel is to run it in the `IO` monad, rather than the `ST s ()` monad since `ST` has knowledge of its sequential context in `s`. Parallelizing bitonic sort was a challenge, as our naive attempt to do each swap in parallel (as one would normally do in a sorting network) resulted in fine-grain parallelism, with enormous amounts of tasks taking very little time. Instead of parallelizing each swap, we instead batched computations by parallelizing recursive calls, with depth limiting.

3.3.4 Hybrid Sort

Hybrid sort was simply a combination of sequential quicksort and the merge operation with the compare-exchange procedure. We divided the list evenly among n tasks and then ran a sequential quicksort in each task in parallel. These lists were then concatenated using the parallel merge procedure used in the parallel merge sort algorithm. This scheme gave very good results as it had low granularity and very little communication between tasks.

4 Evaluation

4.1 Settings

We ran our experiments on an *Intel Core i7 3700k 3.50 GHz processor*.

4.2 Benchmarks

We performed many different benchmarks to test our different sorting implementations, however, we settled on two to discuss in this paper. First, we ran 128 trials of sorting a random permutation of the dictionary. Second, we ran 32 trials of sorting a uniformly distributed list of integers. To compare the speed of all of the algorithms we used 2^{20} integers, however, to compare the speed up when increasing the number of cores on our parallel implementations, we ran this sort on 2^{18} integers.

4.3 Results




Figure 1




Figure 2

Our experiments show that in almost every case our parallel algorithms ran faster than our sequential implementations. The only exception is that parallel Bitonic sort was slower than its sequential counterpart when sorting integers. One other important thing to note is that our hybrid sort implementation was the fastest sorting algorithm in all of the different tests which we ran. Along with this, our implementation was 4x faster than Haskell's sequential implementation of `Data.List.sort` as well as ~25% faster than GNU sort on a random dictionary. GNU sort is completely written in C and coming anywhere near this low level api was an unexpected result.¹

4.4 Performance Analysis

4.4.1 Parallel Merge Sort

¹ We did not have access to any sort of GNU sort api, which means that these metrics are slightly unfair. GNU sort reads files in parallel and because of the lack of an api, it is difficult to time only the sort. We assume that the sort is pretty similar in time to GNU sort's implementation.




Figure 3



Figure 4: Eventlog for Parallel Merge Sort

This figure shows around a 2x speedup with 8 cores. Threadscope reveals the compare and exchange procedure happening on two cores at the end.

4.4.2 Parallel Quicksort




Figure 5




Figure 6: Eventlog for Parallel Quicksort

The figure shows around a 2x speedup with 8 cores. Threadscope reveals that the concatenation of the lists at the end limited our speedups.

4.4.3 Parallel Bitonic Sort




Figure 7




Figure 8: Eventlog for Parallel Bitonic Sort

Bitonic sort is very evenly distributed across 8 cores, however, it still remains slower than the rest. This resulted in around a 4x speedup.

4.4.4 Hybrid Sort

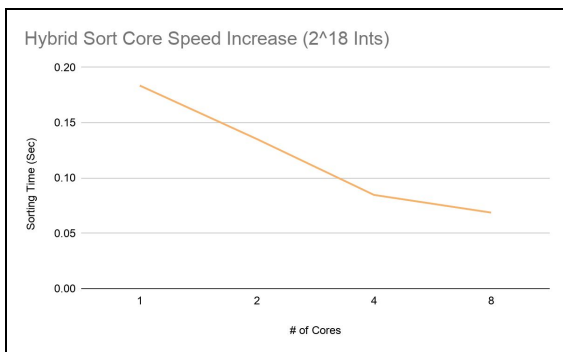


Figure 9




Figure 10: Eventlog for Hybrid Sort

The figure shows around a 3-4x speedup. Threadscope shows each quicksort being divided evenly among the 8 cores and then the compare and exchange sequence happening at the end.

5 Conclusions and Future Work

GNU sort is 4544 lines of C code. Our hybrid haskell sort implementation was ~30 lines. While we cannot give exact speeds of the two implementations, the fact these two parallel sorts were in close in speed is impressive. The ability to provide both high level API's for data parallelism and low level API's for data manipulation speaks to the robustness of parallel functional programming in Haskell. While there were challenges such as handling inter-task communication and memory mutability, our code ended up being relatively efficient, fast and clean.

One of the largest failures which we ran into during our project was the parallelization of Bitonic sort. Bitonic sort is optimized for parallel execution, because each compare and swap operation can be done at the same time. This would be extremely efficient if you had $\frac{n}{2}$ cores where $n = \text{size of list}$ but because we only had 8 cores, our program seemed to spend more time parallelizing and scheduling computations than it did evaluating them. In our future work, Bitonic sort could be greatly improved by running the algorithm using Haskell GPU support for CUDA. While this would introduce different issues such as memory synchronization, it would overcome many of the problems that Bitonic sort faces when running on the CPU.

6 Usage

The program can be simply built using `stack build` from within the root directory. It can then be run with `stack exec -- par-sort-exe [OPTIONS]`. Running without options, or with either the `-h` or `-help` flag yields the following help text

```
Usage: parsort
-s [ default | bitonicSeq | mergeSeq | quickSeq | bitonicPar |
    mergePar | quickPar | hybrid ]
-i [file]
-z [size]
Available Options:
  -s, -sort    Specifies which sort to perform
  -i, -input   [...file] Provides a pointer to a file with new-line separated
                values to sort
  -z, -size    Specifies a uniformly distributed random array of 2^[size]
                elements
  -h, -help    Shows the usage information
```

7 References

- [1] [11.4 Mergesort](#)
- [2] [Lecture 12: Parallel quicksort algorithms](#)
- [3] [Bitonic Sort: Overview](#)

A Code List

Main.hs

```
1 module Main where
2
3 import Control.DeepSeq (NFData)
4 import Data.List (sort)
5 import Data.String (fromString)
6 import qualified Data.Vector as V
```



```

7 import          Lib                (bitonicPar, bitonicSeq, hybridPar,
8                                     mergePar, mergeSeq, quickPar, quickSeq,
9                                     readLines, shuffle, time, timeIO)
10 import         System.Console.GetOpt
11 import         System.Environment (getArgs)
12 import         System.Exit        (die)
13
14 data Arg
15     = Sort String                -- -s
16     | Input String              -- -i
17     | Size String                -- -z
18     | Help                       -- --help
19     deriving (Eq, Ord, Show)
20
21 options :: [OptDescr Arg]
22 options = [
23     Option ['s'] ["sort"] (ReqArg Sort "") "default | bitonicSeq | mergeSeq |
quickSeq | bitonicPar | mergePar | quickPar | hybrid"
24     , Option ['i'] ["input"] (ReqArg Input "") "File path"
25     , Option ['z'] ["size"] (ReqArg Size "") "2^z size of array to be sorted"
26     , Option ['h'] ["help"] (NoArg Help) "Print this help message"
27     ]
28
29
30 runFromArgs :: [String] -> IO ()
31 runFromArgs args = case opt of
32     Help:_ -> die usage
33     (Sort s):(Input f):_ -> do
34         v <- readLines f
35         if s == "bitonicSeq" || s == "bitonicPar" then
36             runBitonic s (fromString "") v
37         else
38             runSort s v
39     (Sort s):(Size z):_ -> do
40         let n = read z :: Int
41             v <- shuffle $ V.enumFromN (1 :: Int) (2^n)
42             if s == "bitonicSeq" || s == "bitonicPar" then
43                 runBitonic s 0 v
44             else
45                 runSort s v
46     _ -> die usage
47     where
48         (opt,_,_) = getOpt Permute options args
49
50 usage :: String
51 usage = "Usage: parsort -s [default | bitonicSeq | mergeSeq | quickSeq |
bitonicPar | mergePar | quickPar | hybrid] -i [file] -z [size]"

```

```

52
53 runSort :: (NFData a, Ord a) => String -> V.Vector a -> IO ()
54 runSort "default"      v = time "Default Sort" (sort $ V.toList v)
55 runSort "quickSeq"    v = time "Sequential Quicksort" (quickSeq v)
56 runSort "mergeSeq"    v = time "Sequential Merge Sort" (mergeSeq v)
57 runSort "hybrid"      v = time "Parallel Hybrid Sort" (hybridPar v)
58 runSort "quickPar"    v = time "Parallel Quick Sort" (quickPar v)
59 runSort "mergePar"    v = time "Parallel Merge Sort" (mergePar v)
60 runSort _ _           = die usage
61
62 runBitonic :: (NFData a, Ord a) => String -> a -> V.Vector a -> IO ()
63 runBitonic "bitonicSeq" a v = time "Sequential Bitonic Sort" (bitonicSeq a v)
64 runBitonic "bitonicPar" a v = timeIO "Parallel Bitonic Sort" (bitonicPar a v)
65 runBitonic _ _ _       = die usage
66
67 main :: IO ()
68 main = getArgs >>= runFromArgs

```

Lib.hs

```

1 module Lib (
2     module Sequential,
3     module Parallel,
4     module Utils,
5 ) where
6
7 import           Parallel
8 import           Sequential
9 import           Utils

```

Utils.hs

```

1 {-# LANGUAGE DeriveGeneric #-}
2 module Utils
3     (
4     fillBitonic,
5     readLines,
6     shuffle,
7     Dumb,
8     time,
9     timeIO
10    )
11 where
12
13 import           Control.DeepSeq      (NFData, force)

```

```

14 import          Control.Monad          (forM_)
15 import qualified Data.ByteString        as B
16 import          Data.Time.Clock        (diffUTCTime, getCurrentTime)
17 import          Data.Vector            ((!))
18 import qualified Data.Vector            as V
19 import qualified Data.Vector.Mutable    as M
20 import          GHC.Generics            (Generic)
21 import          System.IO                (IOMode (ReadMode), hIsEOF, withFile)
22 import          System.Random            (randomRIO)
23
24 shuffle :: V.Vector a -> IO (V.Vector a)
25 shuffle v = do
26     let n = V.length v - 1
27         js <- V.forM (V.enumFromTo 0 n) $ \i -> randomRIO (i, n)
28         return $ V.create $ do
29             o <- V.thaw v
30             forM_ [1..n] $ \i -> M.swap o i (js!i)
31         return o
32
33
34 fillBitonic :: a -> V.Vector a -> V.Vector a
35 fillBitonic a v = V.create $ do
36     o <- V.thaw v
37     let l = V.length v
38         n = 2 ^ (ceiling (logBase 2 (fromIntegral l) :: Double) :: Int) - 1
39         o' <- M.grow o n
40         p <- M.replicate n a
41         M.copy (M.slice l n o') p
42     return o'
43
44 readLines :: String -> IO (V.Vector B.ByteString)
45 readLines filename = withFile filename ReadMode ((V.fromList <$>) . getLines)
46     where
47         getLines handle = do
48             eof <- hIsEOF handle
49             if eof then
50                 return []
51             else
52                 (:) <$> B.hGetLine handle <*> getLines handle
53
54 newtype Dumb = Dumb Integer deriving (Generic, Show)
55
56 instance Eq Dumb where
57     (Dumb 0) == (Dumb 0) = True
58     (Dumb _) == (Dumb 0) = False
59     (Dumb 0) == (Dumb _) = False
60     (Dumb x) == (Dumb y) = Dumb (x-1) == Dumb (y-1)

```

```

61
62 instance Ord Dumb where
63     (Dumb x) <= (Dumb 0) = x <= 0
64     (Dumb 0) <= (Dumb y) = y >= 0
65     (Dumb x) <= (Dumb y) = Dumb (x-1) <= Dumb (y-1)
66
67 instance Num Dumb where
68     (+) (Dumb a) (Dumb b) = Dumb $ a + b
69     (*) (Dumb a) (Dumb b) = Dumb $ a * b
70     abs (Dumb a) = Dumb $ abs a
71     fromInteger = Dumb
72     negate (Dumb a) = Dumb $ -a
73     signum (Dumb a) = Dumb $ signum a
74
75 instance NFData Dumb
76
77 time :: (NFData a) => String -> a -> IO ()
78 time msg a = do
79     start <- getCurrentTime
80     let a' = force a
81     end <- a' `seq` getCurrentTime
82     putStrLn $ msg ++ ": " ++ show (diffUTCTime end start)
83
84 timeIO :: NFData a => [Char] -> IO a -> IO ()
85 timeIO msg io = do
86     start <- getCurrentTime
87     a <- io
88     end <- force a `seq` getCurrentTime
89     putStrLn $ msg ++ ": " ++ show (diffUTCTime end start)

```

Sequential.hs

```

1 module Sequential (bitonicSeq, mergeSeq, quickSeq) where
2
3 import Control.Monad      (when)
4 import Data.Vector       ((!))
5 import qualified Data.Vector as V
6 import qualified Data.Vector.Mutable as M
7 import Uutils            (fillBitonic)
8
9 bitonicSeq :: Ord a => a -> V.Vector a -> V.Vector a
10 bitonicSeq = (bitonic .) . fillBitonic
11
12 bitonic :: Ord a => V.Vector a -> V.Vector a
13 bitonic v = V.create $ do

```

```

14 o <- V.thaw v
15 bitonicSort' o 0 (V.length v) True
16 return o
17 where
18   bitonicSort' o low cnt dir =
19     when (cnt > 1) $ do
20       let k = cnt `div` 2
21           bitonicSort' o low k True
22           bitonicSort' o (low + k) k False
23           bitonicMerge o low cnt dir
24   bitonicMerge o low cnt dir =
25     when (cnt > 1) $ do
26       let k = cnt `div` 2
27           loopSwap o low low k dir
28           bitonicMerge o low k dir
29           bitonicMerge o (low+k) k dir
30   loopSwap o low i k dir =
31     when (i < low + k) $ do
32       compareAndSwap o i (i+k) dir
33       loopSwap o low (i+1) k dir
34   compareAndSwap o i j dir = do
35     oi <- M.read o i
36     oj <- M.read o j
37     when (dir == (oi > oj)) $ M.swap o i j
38
39 mergeSeq :: Ord a => V.Vector a -> V.Vector a
40 mergeSeq = merge . runs
41
42 runs :: Ord a => V.Vector a -> V.Vector (V.Vector a)
43 runs x = V.create $ do
44   o <- M.new (V.length x)
45   runs' 1 x 0 o
46   where
47     runs' i v k o
48       | i < V.length v =
49         if v!(i-1) <= v!i then
50           asc (i-1) i k o
51         else
52           dsc (i-1) i k o
53       | otherwise = return $ M.slice 0 k o
54   asc s i k o =
55     if i < V.length x && x!(i-1) <= x!i then
56       asc s (i+1) k o
57     else do
58       M.write o k (V.slice s (i-s) x)
59       runs' (i+1) x (k+1) o
60   dsc s i k o =

```

```

61         if i < V.length x && x!(i-1) > x!i then
62             dsc s (i+1) k o
63         else do
64             M.write o k (V.reverse $ V.slice s (i-s) x)
65             runs' (i+1) x (k+1) o
66
67 merge :: Ord a => V.Vector (V.Vector a) -> V.Vector a
68 merge v = (!0) $ V.create $ do
69     o <- V.thaw v
70     mergeAll (V.length v) o
71     return $ M.slice 0 1 o
72     where
73         mergeAll k o
74             | k == 1 = return ()
75             | otherwise = do
76                 k' <- mergePairs 0 k o
77                 mergeAll k' o
78         mergePairs i k o
79             | i < k - 1 = do
80                 oi <- M.read o i
81                 oip1 <- M.read o (i+1)
82                 M.write o (i `div` 2) (merge2 oi oip1)
83                 mergePairs (i+2) k o
84             | i == k - 1 = do
85                 oi <- M.read o i
86                 M.write o (i `div` 2) oi
87                 return $ k `div` 2 + 1
88             | otherwise = return $ k `div` 2
89
90 merge2 :: Ord a => V.Vector a -> V.Vector a -> V.Vector a
91 merge2 a b = V.create $ do
92     v <- M.new (V.length a + V.length b)
93     a' <- V.thaw a
94     b' <- V.thaw b
95     go a' b' 0 0 v
96     return v
97     where go a' b' i j v
98         | i < V.length a && j < V.length b = do
99             ai <- M.unsafeRead a' i
100            bj <- M.unsafeRead b' j
101            if ai <= bj then do
102                M.unsafeWrite v (i+j) ai
103                go a' b' (i+1) j v
104            else do
105                M.unsafeWrite v (i+j) bj
106                go a' b' i (j+1) v
107            | i < V.length a = do

```

```

108         ai <- M.unsafeRead a' i
109         M.unsafeWrite v (i+j) ai
110         go a' b' (i+1) j v
111     | j < V.length b = do
112         bj <- M.unsafeRead b' j
113         M.unsafeWrite v (i+j) bj
114         go a' b' i (j+1) v
115     | otherwise = return ()
116
117
118 quickSeq :: Ord a => V.Vector a -> V.Vector a
119 quickSeq x = V.create $ do
120     x' <- V.thaw x
121     quickSort' x' 0 (V.length x - 1)
122     return x'
123     where
124     quickSort' v low high
125         | low < high = do
126             i <- partition v low high
127             quickSort' v low (i - 1)
128             quickSort' v (i + 1) high
129         | otherwise = return ()
130     partition v low high = do
131         i <- go (low - 1) low
132         M.swap v (i+1) high
133         return $ i + 1
134         where
135         go i j
136             | j < high = do
137                 vj <- M.read v j
138                 pivot <- M.read v high
139                 if vj < pivot then do
140                     M.swap v (i+1) j
141                     go (i+1) (j+1)
142                 else
143                     go i (j+1)
144             | otherwise = return i

```

Parallel.hs

```

1 module Parallel (bitonicPar, mergePar, hybridPar, quickPar) where
2 import Control.DeepSeq          (force)
3 import Control.Monad            (when)
4 import Control.Monad.IO.Class
5 import Control.Monad.Par.Class
6 import Control.Monad.Par.IO

```

```

7 import      Control.Parallel.Strategies
8 import      Data.List.Split          (chunksOf)
9 import      Data.Vector              ((!))
10 import qualified Data.Vector         as V
11 import qualified Data.Vector.Mutable as M
12 import qualified Data.Vector.Split   as S
13 import      Sequential              (quickSeq)
14 import      Utils                   (fillBitonic)
15
16 bitonicPar :: (NFData a, Ord a) => a -> V.Vector a -> IO (V.Vector a)
17 bitonicPar = (bitonic .) . fillBitonic
18
19 bitonic :: Ord a => V.Vector a -> IO (V.Vector a)
20 bitonic v = do
21   o <- V.thaw v
22   runParIO $ bitonicSort' o 0 (V.length v) True (0 :: Integer)
23   V.freeze o
24   where
25     bitonicSort' o low cnt dir l =
26     when (cnt > 1) $ do
27       let k = cnt `div` 2
28           if l < 7 then do
29             a <- spawn $ bitonicSort' o low k True (l+1)
30             b <- spawn $ bitonicSort' o (low + k) k False (l+1)
31             get a
32             get b
33           else do
34             bitonicSort' o low k True (l+1)
35             bitonicSort' o (low + k) k False (l+1)
36             bitonicMerge o low cnt dir l
37     bitonicMerge o low cnt dir l =
38     when (cnt > 1) $ do
39       let k = cnt `div` 2
40           loopSwap o low low k dir
41           if l < 7 then do
42             a <- spawn $ bitonicMerge o low k dir (l+1)
43             b <- spawn $ bitonicMerge o (low+k) k dir (l+1)
44             get a
45             get b
46           else do
47             bitonicMerge o low k dir (l+1)
48             bitonicMerge o (low+k) k dir (l+1)
49     loopSwap o low i k dir =
50     when (i < low + k) $ do
51       loopSwap o low (i+1) k dir
52       liftIO $ compareAndSwap o i (i+k) dir
53     compareAndSwap o i j dir = do

```



```

54     oi <- M.read o i
55     oj <- M.read o j
56     when (dir == (oi > oj)) $ M.swap o i j
57
58 hybridPar :: (NFData a, Ord a) => V.Vector a -> V.Vector a
59 hybridPar v = merge $ V.fromList $ parMap rdeepseq quickSeq chunks
60     where
61     n = V.length v
62     chunks = S.chunksOf (n `div` 32) v
63
64 quickPar :: (NFData a, Ord a) => V.Vector a -> V.Vector a
65 quickPar x = runEval $ quickPar' chunks
66     where
67     quickPar' :: (NFData a, Ord a) => [V.Vector a] -> Eval (V.Vector a)
68     quickPar' [] = return V.empty
69     quickPar' [v] = rpar (quickSeq v)
70     quickPar' (v:vs) = do
71         let p = V.head v
72             vs' <- parList rdeepseq (V.partition (<p) <$> (V.tail v:vs))
73             lower <- parList rdeepseq (fst <$> vs')
74             upper <- parList rdeepseq (snd <$> vs')
75             lower' <- parList rdeepseq (filter (not . null) $ V.concat <$>
chunksOf 2 lower)
76             upper' <- parList rdeepseq (filter (not . null) $ V.concat <$>
chunksOf 2 upper)
77             lower'' <- quickPar' lower'
78             upper'' <- quickPar' upper'
79             rpar ((lower'' `V.snoc` p) V.++ upper'')
80     n = V.length x
81     chunks = S.chunksOf (n `div` 32) x
82
83 mergePar :: (NFData a, Ord a) => V.Vector a -> V.Vector a
84 mergePar = merge . runs
85
86 runs :: Ord a => V.Vector a -> V.Vector (V.Vector a)
87 runs x = V.create $ do
88     o <- M.new (V.length x)
89     runs' 1 x 0 o
90     where
91     runs' i v k o
92     | i < V.length v =
93         if v!(i-1) <= v!i then
94             asc (i-1) i k o
95         else
96             dsc (i-1) i k o
97     | otherwise = return $ M.slice 0 k o
98     asc s i k o =

```

```

99     if i < V.length x && x!(i-1) <= x!i then
100         asc s (i+1) k o
101     else do
102         M.write o k (V.slice s (i-s) x)
103         runs' (i+1) x (k+1) o
104     dsc s i k o =
105     if i < V.length x && x!(i-1) > x!i then
106         dsc s (i+1) k o
107     else do
108         M.write o k (V.reverse $ V.slice s (i-s) x)
109         runs' (i+1) x (k+1) o
110
111
112 merge :: (NFData a, Ord a) => V.Vector (V.Vector a) -> V.Vector a
113 merge x = runEval (merge' (0::Integer) x)
114     where
115     merge' l v
116         | n > 1 = do
117             a' <- merge' (l+1) a >>= if l < 15 then rpar else rseq
118             b' <- merge' (l+1) b >>= if l < 15 then rpar else rseq
119             if l < 1 then merge2Par a' b' else return $ merge2 a' b'
120         | otherwise = return $ v!0
121     where
122     n = V.length v
123     a = V.slice 0 (n `div` 2) v
124     b = V.slice (n `div` 2) (n - n `div` 2) v
125
126 merge2 :: Ord a => V.Vector a -> V.Vector a -> V.Vector a
127 merge2 a b = V.create $ do
128     v <- M.new (V.length a + V.length b)
129     a' <- V.thaw a
130     b' <- V.thaw b
131     go a' b' 0 0 v
132     return v
133     where go a' b' i j v
134         | i < V.length a && j < V.length b = do
135             ai <- M.unsafeRead a' i
136             bj <- M.unsafeRead b' j
137             if ai <= bj then do
138                 M.unsafewrite v (i+j) ai
139                 go a' b' (i+1) j v
140             else do
141                 M.unsafewrite v (i+j) bj
142                 go a' b' i (j+1) v
143         | i < V.length a = do
144             ai <- M.unsafeRead a' i
145             M.unsafewrite v (i+j) ai

```

```

146         go a' b' (i+1) j v
147     | j < V.length b = do
148         bj <- M.unsafeRead b' j
149         M.unsafeWrite v (i+j) bj
150         go a' b' i (j+1) v
151     | otherwise = return ()
152
153 merge2Par :: (NFData a, Ord a) => V.Vector a -> V.Vector a -> Eval (V.Vector a)
154 merge2Par a b = do
155     l <- rpar (force lower)
156     u <- rpar (force upper)
157     return (l V.++ u)
158     where
159         n = V.length a + V.length b
160         h = n `div` 2
161         third (_,_,x) = x
162         lower = third <$> V.postscanl' accumLower (0,0,undefined) (V.enumFromN
163 (0::Integer) h)
164         accumLower (i, j, _) _
165         | i < V.length a && j < V.length b =
166             if a!i <= b!j then
167                 (i+1, j, a!i)
168             else
169                 (i, j+1, b!j)
170         | i < V.length a = (i+1, j, a!i)
171         | otherwise = (i, j+1, b!j)
172         upper = V.reverse $ third <$> V.postscanl' accumUpper (V.length a -
173 1,V.length b - 1,undefined) (V.enumFromN (0::Integer) (n-h))
174         accumUpper (i, j, _) _
175         | i > 0 && j > 0 =
176             if a!i >= b!j then
177                 (i-1, j, a!i)
178             else
179                 (i, j-1, b!j)
180         | i > 0 = (i-1, j, a!i)
181         | otherwise = (i, j-1, b!j)

```