

Project Report: Parallel Convex Hull

Andrei Coman, ac4808@columbia.edu

December 18, 2022

1 Introduction

This report presents a parallel Haskell implementation for the Convex Hull problem. More formally, given N points in the xy -plane, we want to find their convex hull - the smallest convex set containing all N points, or, alternatively, the intersection of all convex sets containing all N points.

An efficient solution to this problem is called Graham's Scan and is presented in Section 2. Then, the remaining sections present the parallelization of this algorithm along with several performance measures of the accompanying Haskell implementation. The complete code listing and usage are presented in the Annex.

2 Sequential Algorithm

Graham's Scan considers the convex-hull of the N points as composed of an upper-hull and a lower-hull. If we let L denote the leftmost point of this convex hull and R denote the rightmost point, the upper-hull contains all the vertices which lie above the $[LR]$ segment, whereas the lower-hull contains the ones below. To compute the upper-hull, the algorithm first sorts the points in increasing order of their x -coordinate. Then, it iterates through the points, maintaining at each step a "convex" stack - where, by convex, we mean that any three adjacent points on the stack form a clockwise turn. Whenever a point is inserted, the stack is popped as much as needed to preserve the convexity property (see Fig. 1 and Fig. 2). The computation of the lower-hull is symmetric, but a concavity property is now enforced. Once the upper and lower hull are computed, the convex hull is their union. This algorithm has complexity $O(N \log N)$ for the sort and $O(N)$ for the following sweeps. A Haskell implementation of this idea can be found in Annex B, under `src/Lib.hs`.




Figure 1: Stack Before Insertion (Convexity Violated)




Figure 2: Stack Following Insertion (Convexity Restored)

3 Parallel Algorithm

3.1 Convex Hull Merging

To parallelize this algorithm, we can adopt the following divide and conquer approach: instead of computing the hull of the entire dataset directly, we split the dataset into two halves, compute their individual hulls in parallel, then merge them to obtain the overall hull. While merging two arbitrary possibly intersecting convex polygons (in particular, convex hulls) can be performed efficiently using more involved data structures, we can choose a split that minimizes the complexity of this operation.

In particular, let us choose some arbitrary vertical line in the xy-plane and partition the points into two subsets according to the side of the separating line they fall on. Then, the convex hulls corresponding to these two subsets will be linearly separated by the aforementioned line. Thus, we only need to handle the problem of merging two linearly-separated convex hulls. This operation can be performed efficiently as follows.

Consider, as in the sequential algorithm, that a convex hull is the union of an upper-hull and a lower-hull. Then, the convex hull resulting from the merge is determined by the common upper-tangent of the two upper-hulls and the common lower-tangent of the two lower-hulls (Fig. 3, Fig. 4).

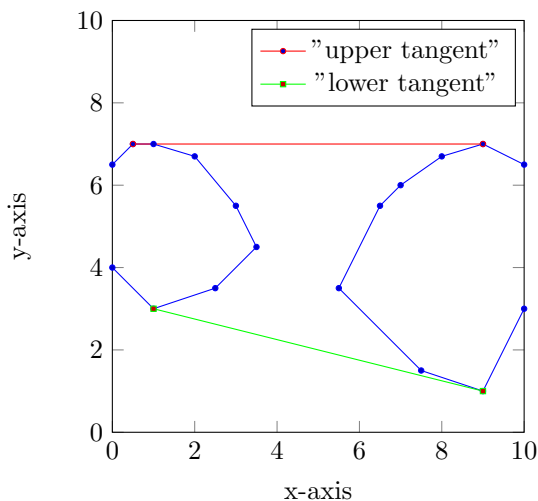


Figure 3: Upper/Lower Tangents




Figure 4: Merged Convex Hulls

Now, these two common tangents can be found by the following "iterative" algorithm. We will look at the process for finding the upper-tangent (the one for the lower-tangent is symmetric). Consider the segment between the left hull's rightmost point and the right hull's leftmost point as the first candidate for the upper-tangent. If shifting this segment to the next point on the left would create a concave angle (Fig. 5), shift the tangent to the left and continue the process recursively. If shifting the segment to the next point on the right would create a concave angle (Fig. 6), shift the tangent to the right and continue the process recursively. Otherwise, the current candidate is the common upper-tangent of the two hulls.

3.2 Parallelism Strategies

Having this merging algorithm for linearly-separated convex hulls, I considered several possible strategies for parallelizing the overall algorithm. The most obvious strategy is to first sort the points by their x-coordinate, then split them into several contiguous chunks and compute the convexHull of each chunk in parallel using parList. Since the points had been sorted before the split, the resulting convex-hulls are necessarily separated by some vertical lines and, so, the algorithm discussed above can be applied. The merging can either be applied as a fold over the resulting hulls or in a divide-and-conquer manner. However, since the convex hull has a logarithmic size in the number of points, this choice does not have a big impact.




Figure 5: Left-shifting a Tangent




Figure 6: Right-shifting a Tangent

Parallelization using parList

```

1 convexHullHelper :: [Point] -> Hull
2 convexHullHelper xs =
3   let chunks = chunksOf chunkSize xs in
4   let hulls = map (force . convexHullNaive) chunks 'using' parList rseq in
5   List.foldl1 mergeHulls hulls

```

However, profiling this solution reveals that, when using several threads, the runtime becomes dominated by the sorting algorithm. This can be improved by implementing a simple parallel quicksort which partitions the array using some pivot, but then proceeds to sort the two partitions in parallel before merging them. To reduce the number of unnecessary (overflowed/fizzled/GC'd) sparks, we can also introduce a threshold for parallelism, after which the algorithm is just executed in series.

Parallel Naive QuickSort

```

1 quicksort :: (NFData a, Ord a) => [a] -> [a]
2 quicksort [] = []
3 quicksort (x : xs) =
4   let l = quicksort [y | y <- xs, y <= x] in
5   let r = quicksort [y | y <- xs, y > x] in
6   if length xs > chunkSize then
7     let parRes = do l' <- rpar (force l)
8                   r' <- rpar (force r)
9                   _ <- rseq l'
10                  _ <- rseq r'
11                 return (l', r') in
12     let (l', r') = runEval $ parRes in
13     l' ++ [x] ++ r'
14   else l ++ [x] ++ r

```

These two ideas already lead to good parallel performance. However, the runtime can still be improved by leveraging the similarity between the sorting algorithm and the divide-and-conquer approach for splitting and merging the convex hulls. More specifically, we can modify the parallel computation of the convex hull can be modified to follow the same structure as the sort: split the dataset into three partitions using some arbitrary pivot: elements to the left of the pivot, the pivot, elements to the right of the pivot. Then, depending on whether some threshold is exceeded, either compute the convex hull using the naive sequential algorithm or recursively compute the convex hull for all partitions recursively. Then, merge the left hull with the pivot and the right hull.

```

1 {- recursively partition the input array; call naive
2   - convex hull algorithm on leaves; merge left and right
3   - hulls at internal nodes
4   -}
5 convexHullHelper :: [Point] -> Hull
6 convexHullHelper xs =
7   if length xs <= chunkSize
8   then convexHullNaive $ quicksort xs
9   else -- partition list of points into left/right halves
10      let pivot = head xs in
11          let (ls, _, rs) = partition xs pivot [] [] [] in
12
13              -- compute the convex hulls of the two halves in parallel
14              let lhull = convexHullHelper ls in
15                  let rhull = convexHullHelper rs in
16                      let parRes = do lhull' <- rpar (force lhull)
17                                      rhull' <- rpar (force rhull)
18                                      _ <- rseq lhull'
19                                      _ <- rseq rhull'
20                                      return (lhull', rhull') in
21                          let (lhull', rhull') = runEval $ parRes in
22
23              -- merge the left hull, the pivot and the right hull
24              lhull' `mergeHulls` (Hull (Cap [pivot] [pivot]) (Cap [pivot] [pivot]))
25              `mergeHulls` rhull'

```

This implementation outperforms the others mentioned so far and is, therefore, the final parallel approach. For a full listing of the code with all the helper functions, see Annex B.

4 Parameter Fitting




Figure 7: Runtime for different numbers of threads at different granularities (10^7 -point dataset)




Figure 8: Optimal leaf size for different numbers of threads

The approach detailed above requires the user to set one parameter: the chunk size or leaf size - the number of points beyond which employing parallelism in the convex hull computation becomes less efficient than the sequential method. Fig. 7 shows how the runtime of the algorithm differs for different leaf sizes given different numbers of threads. Taking the optimal size for each number of threads, we obtain Fig. 8. The optimal chunk size follows a function resembling a reverse-sigmoid. So, if one were to automate setting this parameter, fitting this graph with a reverse-sigmoid should yield good results.

5 Results

Running the sequential implementation and the parallel implementation (with different numbers of threads, always with the optimal chunk/leaf size) yields the following runtimes and speedups:

No. threads	RT	Speedup
sequential	41.37	-
2	27.74	x1.49
3	25.02	x1.65
4	23.06	x1.79
5	22.30	x1.86
6	21.37	x1.94
7	20.97	x1.97
8	21.13	x1.96
9	21.71	x1.90

Table 1: Speedup

Here, we can see that the optimal speedup is close to x2 and is obtained for 7 threads.

Using Threadscope, we can visualize how the load is balanced throughout the entire run:




Figure 9: Overall Runtime of Algorithm (Threadscope)

and just throughout the parallel section:



Figure 10: Parallel Section of the Algorithm (Threadscope)

We can see that, after the data is read from disk (which still takes a considerable amount of time due to the size of the input), the parallel section is fairly-well balanced. In the case depicted above, which uses only 4 threads, all threads are performing work consistently throughout the parallel section of the execution. However, the overall profile is not smooth because of the large amount of garbage collection - garbage is generated by the partitioning of the dataset, the sorting, and the stack in Graham's algorithm.

6 Sample Input/Output

Sample Input/Output

```
1 2.0 0.0
2 0.0 2.0
3 1.0 3.0
4 0.0 4.0
5 3.0 3.0
6 2.0 6.0
7 4.0 2.0
8 4.0 4.0
```

```
1 (0.0, 2.0)
2 (2.0, 0.0)
3 (4.0, 2.0)
4 (4.0, 4.0)
5 (2.0, 6.0)
6 (0.0, 4.0)
```

Note: the implementation presented here produces only the vertices of the convex hull. Other tools must be used for visualization.




Figure 11: Dataset Visualization




Figure 12: Convex Hull Visualization

7 Conclusion

This report presents how Graham's scan can be efficiently parallelized, resulting in an almost-x2 speedup over the entire runtime of the program. This shows that, even though Graham's scan has a linear complexity, the stack based algorithm preceded by a sorting of the dataset can be made much more efficient by introducing parallelism at all stages (besides I/O). The best performance was obtained for 7 threads and a parallel/sequential threshold of 1,000 points. This allowed the processing of 10,000,000 points in 20.97 seconds.

The workload is distributed well amongst the working threads, which leads to fairly good balancing. However, the garbage collection is a noticeable setback to this approach, constantly fragmenting the execution of the threads and resulting in long, frequent garbage collection breaks. The question of how garbage production could be minimized could be considered as further work.

A Usage

The following command runs the parallel algorithm with n threads:

```
1 stack run -- +RTS -Nn -s -ls -RTS p <input-file> <output-file>
```

As a side effect, this command also prints runtime statistics to standard error and produces an associated .eventlog file which can be inspected with Threadscope. To run the sequential implementation mentioned in the beginning of this report, use the command:

```
1 stack run -- +RTS -Nn -s -ls -RTS s <input-file> <output-file>
```

Input files must contain a sequence of points, each represented as a space-separated pair of doubles per line. To test the implementation, run:

```
1 stack test
```

B Code Listing

app/Main.hs

```
1 module Main (main) where
2
3 import qualified Data.Text as B (words, lines)
4 import qualified Data.Text.IO as B (readFile)
5 import qualified Data.Text.Read as B (double)
6 import Types (Point(..))
7 import qualified Lib (convexHull)
8 import qualified ParLib (convexHull)
9 import System.Environment (getArgs, getProgName)
10 import System.Exit (exitFailure)
11 import System.IO (hPutStrLn, stderr, withFile, IOMode(WriteMode))
12
13 main :: IO ()
14 main = do
15     args <- getArgs
16     case args of
17     [mode, inputFilename, outputFilename] -> do
18         contents <- B.readFile inputFilename
19         let pts = map (\p -> case p of
20             [xstr, ystr] -> (case (B.double xstr, B.double ystr) of
21                 (Right (x, _), Right (y, _)) ->
22                     Point x y
23                 _ -> error "error while parsing input"
24             )
25             - -> error "malformed input"
26         )
27             (map B.words $ B.lines contents)
28         let hull = case mode of
29             "p" -> ParLib.convexHull pts
30             "s" -> Lib.convexHull pts
31             - -> error "invalid mode"
32         withFile outputFilename WriteMode
33             (\h -> do mapM_ (\p -> hPutStrLn h $ show p) $ hull)
34
35     - -> do
36         pn <- getProgName
37         hPutStrLn stderr $ "Usage: " ++
38             pn ++
39             "<mode(p/s)> <input-filename> <output-filename>"
40         exitFailure
```

```

1 module Types
2   ( Point(..)
3     , Cap(..)
4     , Hull(..)
5     ) where
6
7 import Control.Parallel.Strategies
8 import Control.DeepSeq
9 -- a cap is either the top or the bottom half of a hull
10 -- consists of its left-right and its right-left traversal
11 -- a hull is made out of its upper-hull/cap and its lower one
12 data Cap = Cap [Point] [Point]
13 data Hull = Hull Cap Cap
14 instance NFData Hull where
15   rnf (Hull (Cap upperLR upperRL)
16            (Cap lowerLR lowerRL)) = (rnf upperLR) 'seq'
17                                     (rnf upperRL) 'seq'
18                                     (rnf lowerLR) 'seq'
19                                     (rnf lowerRL)
20
21 data Point = Point Double Double
22 instance Eq Point where
23   (Point ax ay) == (Point bx by) = (ax == ay) && (bx == by)
24 instance Ord Point where
25   (Point ax ay) 'compare' (Point bx by) =
26     case (ax 'compare' bx) of
27       LT -> LT
28       GT -> GT
29       EQ -> ay 'compare' by
30 instance Show Point where
31   show (Point x y) = "(" ++ (show x) ++ ", " ++ (show y) ++ ")"
32 instance NFData Point where
33   rnf (Point x y) = (rnf x) 'seq' (rnf y)

```

```

1 module ParLib
2   ( convexHull
3     ) where
4
5 import qualified Data.List as List
6 import Control.Parallel.Strategies
7 import Control.DeepSeq
8 import Types (Point(..), Cap(..), Hull(..))
9
10 {- get the convexity of an ordered triplet of points
11    - by computing a quantity proportional to the signed
12    - area of the triangle formed by these points
13    - this is > 0 when convex, < 0 when concave and
14    - = 0 for collinear points
15    -}
16 convexity :: Point -> Point -> Point -> Double
17 convexity (Point ax ay)
18           (Point bx by)
19           (Point cx cy) = (ax * by + bx * cy + cx * ay) -
20                           (ax * cy + bx * ay + cx * by)
21
22 {- find the upper hull of an xy-sorted set of points
23    - maintain a stack and enforce convexity at insertion
24    -}
25 upperHull :: [Point] -> [Point]
26 upperHull xs = List.foldl insertConvex [] xs -- for every point
27   where insertConvex :: [Point] -> Point -> [Point]
28         insertConvex stack@(a : pop@(b : _)) p =
29           if convexity b a p < 0

```

```

30         then p : stack           -- add on top of stack
31         else insertConvex pop p -- pop and retry
32     insertConvex stack p =
33         p : stack
34
35 {- find the lower hull of an xy-sorted set of points
36  - maintain a stack and enforce cocavity at insertion
37  -}
38 lowerHull :: [Point] -> [Point]
39 lowerHull xs = List.foldl insertConcav [] xs -- for every point
40     where insertConcav :: [Point] -> Point -> [Point]
41           insertConcav stack@(a : pop@(b : _)) p =
42               if convexity b a p > 0
43               then p : stack           -- add on top of stack
44               else insertConcav pop p -- pop and retry
45           insertConcav stack p =
46               p : stack
47
48 {- compute the lower and upper hull using the naive
49  - Graham's scan algorithm and then wrap them into
50  - a "Hull" instance
51  -}
52 convexHullNaive :: [Point] -> Hull
53 convexHullNaive xs =
54     let lower = lowerHull xs in
55     let upper = upperHull xs in
56     Hull (Cap (reverse lower) lower) (Cap (reverse upper) upper)
57
58 {- find the common upper tangent of two hulls
59  - the input consists of the list starting at the
60  - rightmost point of the left hull, followed by
61  - the list starting at the leftmost point of the
62  - right hull
63  - algorithm: shift the tangent incrementally
64  - until the optimum is reached
65  -}
66 upperTangent :: [Point] -> [Point] -> ([Point], [Point])
67 upperTangent xl xr
68     -- can shift tangent to the right
69     | (l : _) <- xl
70     , (r : xrs@(rnext : _)) <- xr
71     , convexity l r rnext > 0 = upperTangent xl xrs
72     -- can shift tangent to the left
73     | (l : xls@(lnext : _)) <- xl
74     , (r : _) <- xr
75     , convexity lnext l r > 0 = upperTangent xls xr
76     -- optimum reached
77     | otherwise = (xl, xr)
78
79 {- similar to upperTangent, but checks for concavity
80  - instead of convexity
81  -}
82 lowerTangent :: [Point] -> [Point] -> ([Point], [Point])
83 lowerTangent xl xr
84     -- can shift tangent to the right
85     | (l : _) <- xl
86     , (r : xrs@(rnext : _)) <- xr
87     , convexity l r rnext < 0 = lowerTangent xl xrs
88     -- can shift tangent to the left
89     | (l : xls@(lnext : _)) <- xl
90     , (r : _) <- xr
91     , convexity lnext l r < 0 = lowerTangent xls xr
92     -- optimum reached
93     | otherwise = (xl, xr)
94
95 {- to merge two hulls, compute their common upper tangent
96  - and lower tangent, then reconstruct the resulting hull
97  -}

```

```

98 mergeHulls :: Hull -> Hull -> Hull
99 mergeHulls (Hull (Cap _ lowerRL) (Cap _ upperRL))
100           (Hull (Cap lowerLR _) (Cap upperLR _)) =
101   -- compute upper/lower tangents
102   let (lowerL, lowerR) = lowerTangent lowerRL lowerLR in
103   let (upperL, upperR) = upperTangent upperRL upperLR in
104   -- combine into a 'Hull' instance
105   let lower = (reverse lowerL) ++ lowerR in
106   let upper = (reverse upperL) ++ upperR in
107   Hull (Cap lower $ reverse lower) (Cap upper $ reverse upper)
108
109 {- recursively partition the input array; call naive
110  - convex hull algorithm on leaves; merge left and right
111  - hulls at internal nodes
112  -}
113 convexHullHelper :: [Point] -> Hull
114 convexHullHelper xs =
115   if length xs <= 200
116   then convexHullNaive $ quicksort xs
117   else -- partition list of points into left/right halves
118        let pivot = head xs in
119        let (ls, _, rs) = partition xs pivot [] [] [] in
120
121        -- compute the convex hulls of the two halves in parallel
122        let lhull = convexHullHelper ls in
123        let rhull = convexHullHelper rs in
124        let parRes = do lhull' <- rpar (force lhull)
125                        rhull' <- rpar (force rhull)
126                        _ <- rseq lhull'
127                        _ <- rseq rhull'
128                        return (lhull', rhull') in
129        let (lhull', rhull') = runEval $ parRes in
130
131        -- merge the left hull, the pivot and the right hull
132        lhull' `mergeHulls` (Hull (Cap [pivot] [pivot]) (Cap [pivot] [pivot]))
133        `mergeHulls` rhull'
134
135 {- very simple quicksort method used in the leaves of
136  - the helper above; for some reason, this is faster
137  - than the library List.sort
138  -}
139 quicksort :: (NFData a, Ord a) => [a] -> [a]
140 quicksort [] = []
141 quicksort (x : xs) =
142   let l = quicksort [y | y <- xs, y <= x] in
143   let r = quicksort [y | y <- xs, y > x] in
144   l ++ [x] ++ r
145
146 {- partition function used in the convexHullHelper;
147  - splits a list into elements less than, equal to,
148  - and greater than a given pivot
149  -}
150 partition :: Ord a => [a] -> a -> [a] -> [a] -> [a] -> ([a], [a], [a])
151 partition [] _ lt eq gt = (lt, eq, gt)
152 partition (x : xs) pivot lt eq gt =
153   case x `compare` pivot of
154   LT -> partition xs pivot (x : lt) eq gt
155   EQ -> partition xs pivot lt (x : eq) gt
156   GT -> partition xs pivot lt eq (x : gt)
157
158 {- convert a Hull to list-of-point representation
159  - since the upper/lower hulls have two common endpoints,
160  - we have to remove these duplicates before concatenating
161  - their associated lists; because of Haskell laziness,
162  - the use of 'init' and '++' should not decrease the
163  - performance by a lot
164  -}
165 toList :: Hull -> [Point]

```

```

166 toList (Hull (Cap lowerLR _) (Cap _ upperRL)) =
167   (init lowerLR) ++ (init upperRL)
168
169 {- entry point into the convex-hull library
170  - calls the underlying helper method and converts the output to
171  - a list-of-point format
172  -}
173 convexHull :: [Point] -> [Point]
174 convexHull xs =
175   toList $ convexHullHelper xs

```

app/Lib.hs

```

1  module Lib
2    (convexHull
3     ) where
4
5  import qualified Data.List as List
6  import Types (Point(..))
7
8  {- get the convexity of an ordered triplet of points
9   - by computing a quantity proportional to the signed
10  - area of the triangle formed by these points
11  - this is > 0 when convex, < 0 when concave and
12  - = 0 for collinear points
13  -}
14 convexity :: Point -> Point -> Point -> Double
15 convexity (Point ax ay)
16           (Point bx by)
17           (Point cx cy) = (ax * by + bx * cy + cx * ay) -
18                           (ax * cy + bx * ay + cx * by)
19
20 {- find the upper hull of an xy-sorted set of points
21  - maintain a stack and enforce convexity at insertion
22  -}
23 upperHull :: [Point] -> [Point]
24 upperHull xs = List.foldl insertConvex [] xs
25   where insertConvex :: [Point] -> Point -> [Point]
26         insertConvex stack@(a : pop@(b : _)) p =
27           if convexity b a p < 0
28             then p : stack
29             else insertConvex pop p
30         insertConvex stack p =
31           p : stack
32
33 {- find the lower hull of an xy-sorted set of points
34  - maintain a stack and enforce concavity at insertion
35  -}
36 lowerHull :: [Point] -> [Point]
37 lowerHull xs = List.foldl insertConcav [] xs
38   where insertConcav :: [Point] -> Point -> [Point]
39         insertConcav stack@(a : pop@(b : _)) p =
40           if convexity b a p > 0
41             then p : stack
42             else insertConcav pop p
43         insertConcav stack p =
44           p : stack
45
46 {- compute the lower and upper hull and concatenate
47  - them; the two hulls have two common endpoints, so
48  - remove these from the upper hull before appending
49  -}
50 convexHull :: [Point] -> [Point]
51 convexHull xs =
52   let xsSorted = quicksort xs in
53   let lower = lowerHull xsSorted in
54   let upper = upperHull xsSorted in

```

```
55     combine lower upper
56     where combine :: [Point] -> [Point] -> [Point]
57           combine as bs =
58             let ta = List.reverse .tail $ as in
59               let tb = init bs in
60                 ta ++ tb
61
62 quicksort :: Ord a => [a] -> [a]
63 quicksort [] = []
64 quicksort (x : xs) =
65     let l = quicksort [y | y <- xs, y <= x] in
66     let r = quicksort [y | y <- xs, y > x] in
67     l ++ [x] ++ r
```

References

- [1] Miller Russ Chen, Weiyang. Parallel Implementation of the Convex Hull Problem. <https://cse.buffalo.edu/faculty/miller/Courses/CSE633/Weiyang-Chen-Spring-2020.pdf>.
- [2] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.
- [3] Petr Felkel. Convex Hulls. https://cw.fel.cvut.cz/b181/_media/courses/cg/lectures/04-convexhull.pdf.