

Extra Credit

The attached paper claims to describe a new sorting algorithm that sorts the items in an array in $O(n \log n)$ worst-case time using $O(\log n)$ extra storage and running faster than Quicksort experimentally. I am skeptical. I will offer extra credit points for convincing answers to one or both of the following questions:

1. What (if anything) is wrong with this algorithm?
2. If your answer to 1 is “yes”, is there any way to fix the algorithm so that it is correct and runs in $O(n \log n)$ worst-case time?

Note: This assignment is completely optional and is intended for fun over Spring break. If you don't find it fun, don't do it.