

Quicksort

MergeSort, while it is clearly superior to the $O(n^2)$ sorting algorithms, suffers from two flaws

- the need to pass an additional temp array to the sorting algorithm. This is a pain at anytime, and could be a major problem if the array you are sorting is so big that you have trouble allocating another array of the same size.
- all of the copying back and forth between the data array and the temp array.

Quicksort is another divide-and-conquer algorithm that makes up for these deficiencies. You must be careful implementing QuickSort; small changes in the code can lead to incorrect sorts. When implemented correctly it is one of the best sorting algorithms known.

Quicksort was invented by the British computer scientist Tony Hoare in 1960. His original implementation was in Fortran and the code was so complex most people couldn't follow it. Once recursion became available (in ALGOL 60) Hoare's recursive version was so simple most people didn't take it seriously.

Quicksort was a featured part of two very important standard systems: it was a C-language library function in the original Unix implementation, and it was the sorting algorithm used in the original Java implementation.

The idea behind QuickSort is simple. To sort the elements of array A between indices *first* and *last*:

- a) Choose one of the elements of the array to be the "pivot" value. The algorithm works correctly regardless of which element you choose.
- b) Rearrange the array so that elements that are less than the pivot come before it, and the elements that are greater than or equal to the pivot come after it. The pivot must then be in its final location. Suppose that is at index i .
- c) Recursively sort the elements between *first* and $i-1$, and the elements between $i+1$ and *last*.

If there are n elements in the list the rearranging can be done in one pass with at most n steps.

Note that while the algorithm will correctly sort regardless of which pivot value is chosen, the choice of pivot does make a difference.

Suppose the pivot value we choose happens to be the lowest element. Then one of the recursive calls gets an array with no elements and the other gets all but one of the elements. If the array happens to be already sorted and we choose the leftmost entry to be the pivot, this will happen at each step and we end up doing n recursive calls, sorting n elements, $(n-1)$ elements, $(n-2)$ elements, and so forth element. The partitioning process takes $n + (n-1) + (n-2) + \dots$ steps, which altogether gets us $O(n^2)$ steps.

So in the worst case QuickSort is $O(n^2)$. If that was the end of the story we wouldn't be interested in it.

Think about the more typical case -- unless we are really unlucky, most of the time the pivot value should be around the middle of the data. This means that the recursive calls will each be to about half of the data. Just as with MergeSort, we would expect QuickSort to run in time $O(n \cdot \log(n))$ in the typical case. Even better than MergeSort, QuickSort doesn't have to copy all of the data on each step. In practice QuickSort tends to run significantly faster than MergeSort on typical data.

There are some things we can do to guard against especially unfavorable cases. Sorted, reverse-sorted, and nearly-sorted data all occur more often than you might think; using the first or last element of the list as the pivot doesn't do well in these cases. An easy remedy for this is to use the middle of the array as the pivot value. Some authors recommend taking the median of the left end, the right end, and the middle value. Steps like this don't avoid the worst-case $O(n^2)$ behavior, but they insure that the algorithm reaches its worst case only in very specific and unlikely situations.

The partition step of the algorithm needs to be coded carefully but the idea behind it is easy.

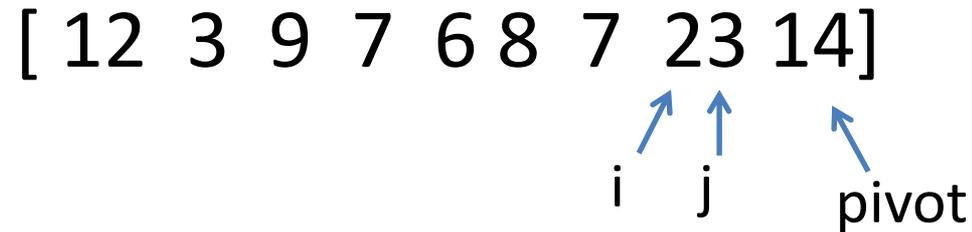
Suppose we start with the following data and choose as the pivot the middle value:

[23 3 9 7 14 8 7 12 6]


pivot

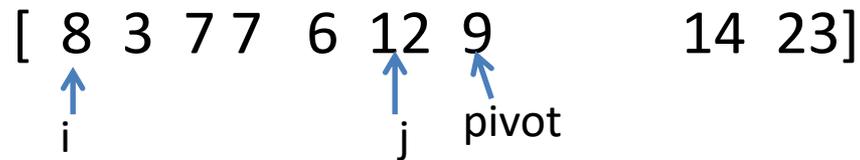
Start by interchanging the pivot with the last element, so it is out of the way. Let i and j be the first and last indices in the array

We repeat this step until i and j cross. For this example that happens immediately, since all of the elements with index less than j are less than the pivot:

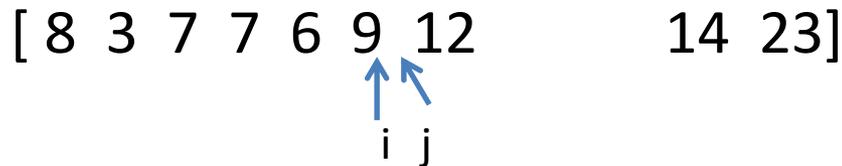


We swap the element at index i with the pivot (in the last entry) and recursively sort the portion of the array before index i and the portion of the array after index i . The latter is only one element, so there is nothing to do. For the portion before i we have:

We swap the values at i and j:



Again we move indices i and j inward; they cross at the element 12, which we swap with the pivot element:



We repeat this to sort the two portions [8 3 7 7 6] and [12]

On the next slide is the complete code for QuickSort:

```

private static <E extends Comparable<? super E>> void
    QuickSort(E[] A, int first, int last) {
        if (first < last) {
            int mid = (first+last)/2;
            E pivot = A[mid];
            swap(A, mid, last);
            int i = first;
            int j = last;
            while ( i < j) {
                while (A[i].compareTo(pivot) < 0)
                    i += 1;
                while (i < j && pivot.compareTo(A[j]) <= 0)
                    j -= 1;
                if (i < j)
                    swap(A, i, j);
            }
            swap(A, last, i);
            QuickSort(A, first, i-1);
            QuickSort(A, i+1, last);
        }
    }
}

```