


3.3 Quicksort



- ▶ quicksort
- ▶ selection
- ▶ duplicate keys
- ▶ system sorts

Two classic sorting algorithms

Critical components in the world's computational infrastructure.

- Full scientific understanding of their properties has enabled us to develop them into practical system sorts.
- Quicksort honored as one of top 10 algorithms of 20th century in science and engineering.

Mergesort.

← last lecture

- Java sort for objects.
- Perl, Python stable sort.

Quicksort.


← this lecture

- Java sort for primitive types.
- C qsort, Unix, g++, Visual C++, Python.

Quicksort

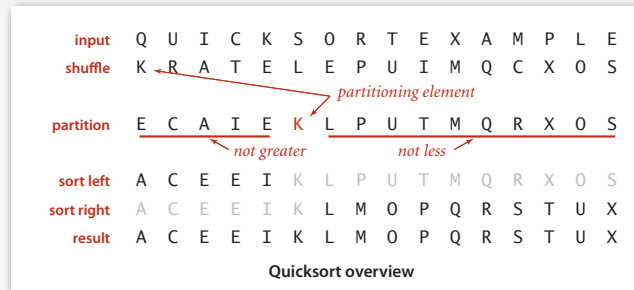
Basic plan.

- **Shuffle** the array.
- **Partition** so that, for some j
 - element $a[j]$ is in place
 - no larger element to the left of j
 - no smaller element to the right of j
- **Sort** each piece recursively.



Sir Charles Antony Richard Hoare
1980 Turing Award

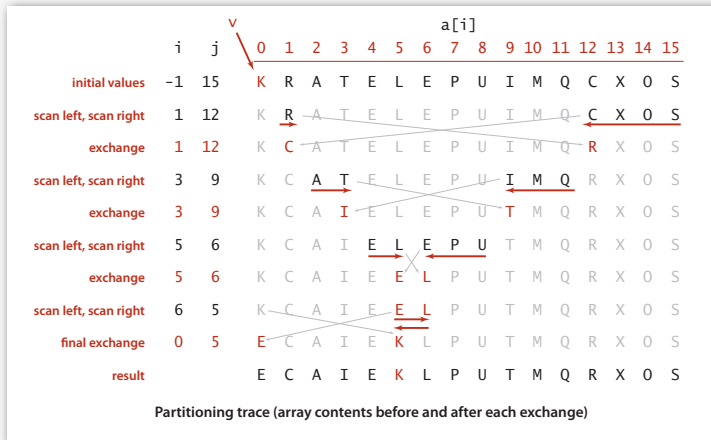
- ▶ quicksort
- ▶ selection
- ▶ duplicate keys
- ▶ system sorts



Quicksort partitioning

Basic plan.

- Scan i from left for an item that belongs on the right.
- Scan j from right for item item that belongs on the left.
- Exchange $a[i]$ and $a[j]$.
- Continue until pointers cross.






5

Quicksort: Java code for partitioning

```
private static int partition(Comparable[] a, int lo, int hi)
{
    int i = lo, j = hi+1;
    while (true)
    {
        while (less(a[++i], a[lo]))
            if (i == hi) break;
        while (less(a[lo], a[--j]))
            if (j == lo) break;
        if (i >= j) break;
        exch(a, i, j);
    }
    exch(a, lo, j);
    return j;
}
```

find item on left to swap
find item on right to swap
check if pointers cross
swap
swap with partitioning item
return index of item now known to be in place



6

Quicksort: Java implementation

```
public class Quick
{
    private static int partition(Comparable[] a, int lo, int hi)
    { /* see previous slide */ }

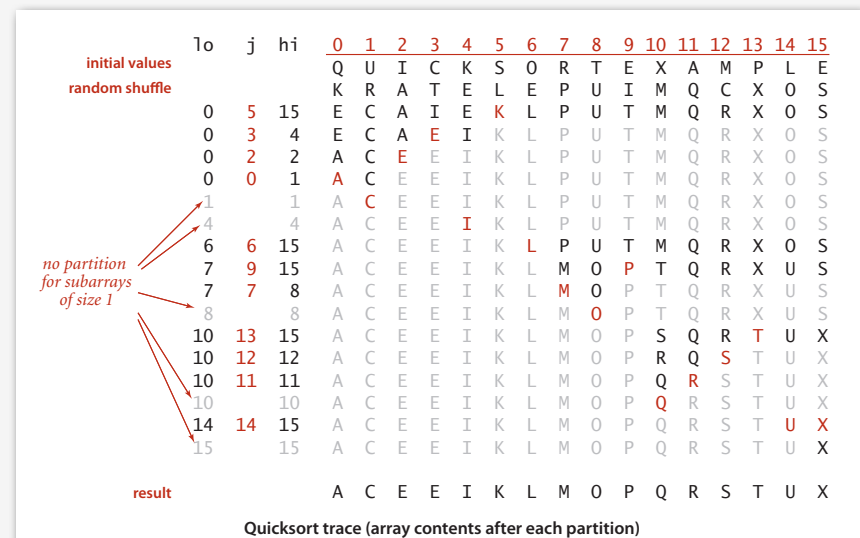
    public static void sort(Comparable[] a)
    {
        StdRandom.shuffle(a);
        sort(a, 0, a.length - 1);
    }

    private static void sort(Comparable[] a, int lo, int hi)
    {
        if (hi <= lo) return;
        int j = partition(a, lo, hi);
        sort(a, lo, j-1);
        sort(a, j+1, hi);
    }
}
```

shuffle needed for performance guarantee

7


Quicksort trace



8

Quicksort animation

50 random elements



<http://www.sorting-algorithms.com/quick-sort>

9

Quicksort: implementation details

Partitioning in-place. Using a spare array makes partitioning easier (and stable), but is not worth the cost.

Terminating the loop. Testing whether the pointers cross is a bit trickier than it might seem.

Staying in bounds. The $(j == 10)$ test is redundant (why?), but the $(i == hi)$ test is not.

Preserving randomness. Shuffling is needed for performance guarantee.

Equal keys. When duplicates are present, it is (counter-intuitively) best to stop on elements equal to the partitioning element.

10

Quicksort: empirical analysis

Running time estimates:

- Home pc executes 10^8 compares/second.
- Supercomputer executes 10^{12} compares/second.

| computer | insertion sort (N^2) | | | mergesort ($N \log N$) | | | quicksort ($N \log N$) | | |
|----------|--------------------------|-----------|-----------|--------------------------|----------|---------|--------------------------|---------|---------|
| | thousand | million | billion | thousand | million | billion | thousand | million | billion |
| home | instant | 2.8 hours | 317 years | instant | 1 second | 18 min | instant | 0.3 sec | 6 min |
| super | instant | 1 second | 1 week | instant | instant | instant | instant | instant | instant |

Lesson 1. Good algorithms are better than supercomputers.

Lesson 2. Great algorithms are better than good ones.

11

Quicksort: average-case analysis

Proposition I. The average number of compares C_N to quicksort an array of N elements is $\sim 2N \ln N$ (and the number of exchanges is $\sim \frac{1}{3} N \ln N$).

Pf. C_N satisfies the recurrence $C_0 = C_1 = 0$ and for $N \geq 2$:

$$C_N = (N+1) + \frac{C_0 + C_1 + \dots + C_{N-1}}{N} + \frac{C_{N-1} + C_{N-2} + \dots + C_0}{N}$$

↑ partitioning
↑ left
↑ right
↙ partitioning probability

- Multiply both sides by N and collect terms:

$$NC_N = N(N+1) + 2(C_0 + C_1 + \dots + C_{N-1})$$

- Subtract this from the same equation for $N-1$:

$$NC_N - (N-1)C_{N-1} = 2N + 2C_{N-1}$$

- Rearrange terms and divide by $N(N+1)$:

$$\frac{C_N}{N+1} = \frac{C_{N-1}}{N} + \frac{2}{N+1}$$

12


Quicksort: average-case analysis

- Repeatedly apply above equation:

$$\begin{aligned}
 \frac{C_N}{N+1} &= \frac{C_{N-1}}{N} + \frac{2}{N+1} \\
 &\stackrel{\text{previous equation}}{=} \frac{C_{N-2}}{N-1} + \frac{2}{N} + \frac{2}{N+1} \\
 &= \frac{C_{N-3}}{N-2} + \frac{2}{N-1} + \frac{2}{N} + \frac{2}{N+1} \\
 &= \frac{2}{1} + \frac{2}{2} + \frac{2}{3} + \dots + \frac{2}{N+1}
 \end{aligned}$$

- Approximate sum by an integral:

$$\begin{aligned}
 C_N &\sim 2(N+1) \left(1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{N} \right) \\
 &\sim 2(N+1) \int_1^N \frac{1}{x} dx
 \end{aligned}$$



- Finally, the desired result:

$$C_N \sim 2(N+1) \ln N \approx 1.39N \lg N$$

13

Quicksort: summary of performance characteristics

Worst case. Number of compares is quadratic.

- $N + (N-1) + (N-2) + \dots + 1 \sim N^2 / 2$.
- More likely that your computer is struck by lightning.

Average case. Number of compares is $\sim 1.39 N \lg N$.

- 39% more compares than mergesort.
- But** faster than mergesort in practice because of less data movement.

Random shuffle.

- Probabilistic guarantee against worst case.
- Basis for math model that can be validated with experiments.

Caveat emptor. Many textbook implementations go **quadratic** if input:

- Is sorted or reverse sorted.
- Has many duplicates (even if randomized!) [stay tuned]

14

Quicksort: practical improvements

Median of sample.

- Best choice of pivot element = median.
- Estimate true median by taking median of sample.

Insertion sort small subarrays.

- Even quicksort has too much overhead for tiny subarrays.
- Can delay insertion sort until end.

Optimize parameters.

- Median-of-3 random elements.
- Cutoff to insertion sort for ≈ 10 elements.

$\sim 12/7 N \ln N$ compares
 $\sim 12/35 N \ln N$ exchanges

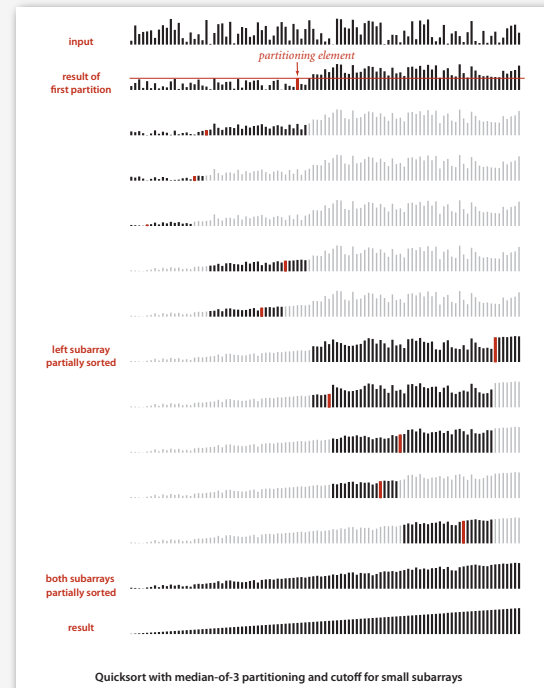
Non-recursive version.

- Use explicit stack.
- Always sort smaller half first.

guarantees $O(\log N)$ stack size

15

Quicksort with cutoff to insertion sort: visualization



Quicksort with median-of-3 partitioning and cutoff for small subarrays

16

- › quicksort
- › **selection**
- › duplicate keys
- › system sorts

17

Selection

Goal. Find the k^{th} largest element.

Ex. Min ($k = 0$), max ($k = N-1$), median ($k = N/2$).

Applications.

- Order statistics.
- Find the "top k ."

Use theory as a guide.

- Easy $O(N \log N)$ upper bound.
- Easy $O(N)$ upper bound for $k = 1, 2, 3$.
- Easy $\Omega(N)$ lower bound.

Which is true?

- $\Omega(N \log N)$ lower bound? ← is selection as hard as sorting?
- $O(N)$ upper bound? ← is there a linear-time algorithm for all k ?

18


Quick-select

Partition array so that:

- Element $a[j]$ is in place.
- No larger element to the left of j .
- No smaller element to the right of j .

Repeat in **one** subarray, depending on j ; finished when j equals k .

```
public static Comparable select(Comparable[] a, int k)
{
    StdRandom.shuffle(a);
    int lo = 0, hi = a.length - 1;
    while (hi > lo)
    {
        int j = partition(a, lo, hi);
        if (j < k) lo = j + 1;
        else if (j > k) hi = j - 1;
        else return a[k];
    }
    return a[k];
}
```



19

Quick-select: mathematical analysis

Proposition. Quick-select takes **linear** time on average.

Pf sketch.

- Intuitively, each partitioning step roughly splits array in half:
 $N + N/2 + N/4 + \dots + 1 \sim 2N$ compares.
- Formal analysis similar to quicksort analysis yields:

$$C_N = 2N + k \ln(N/k) + (N-k) \ln(N/(N-k))$$

Ex. $(2 + 2 \ln 2)N$ compares to find the median.

Remark. Quick-select uses $\sim N^2/2$ compares in worst case, but as with quicksort, the random shuffle provides a probabilistic guarantee.

20

Theoretical context for selection

Challenge. Design algorithm whose worst-case running time is linear.

Proposition. [Blum, Floyd, Pratt, Rivest, Tarjan, 1973] There exists a compare-based selection algorithm whose worst-case running time is linear.

Remark. But, algorithm is too complicated to be useful in practice.

Use theory as a guide.

- Still worthwhile to seek **practical** linear-time (worst-case) algorithm.
- Until one is discovered, use quick-select if you don't need a full sort.

21

Generic methods

In our `select()` implementation, client needs a cast.

```
Double[] a = new Double[N];
for (int i = 0; i < N; i++)
    a[i] = StdRandom.uniform();
Double median = (Double) Quick.select(a, N/2);
```

unsafe cast
required

The compiler also complains.

```
% javac Quick.java
Note: Quick.java uses unchecked or unsafe operations.
Note: Recompile with -Xlint:unchecked for details.
```

Q. How to fix?

22

Generic methods

Pedantic (safe) version. Compiles cleanly, no cast needed in client.

```
public class QuickPedantic
{
    public static <Key extends Comparable<Key>> Key select(Key[] a, int k)
    { /* as before */ }

    public static <Key extends Comparable<Key>> void sort(Key[] a)
    { /* as before */ }

    private static <Key extends Comparable<Key>> int partition(Key[] a, int lo, int hi)
    { /* as before */ }

    private static <Key extends Comparable<Key>> boolean less(Key v, Key w)
    { /* as before */ }

    private static <Key extends Comparable<Key>> void exch(Key[] a, int i, int j)
    { Key swap = a[i]; a[i] = a[j]; a[j] = swap; }
}
```

generic type variable
(value inferred from argument a[])

return type matches array type

can declare variables of generic type

<http://www.cs.princeton.edu/algs4/35applications/QuickPedantic.java.html>

Remark. Obnoxious code needed in system sort; not in this course (for brevity).

23

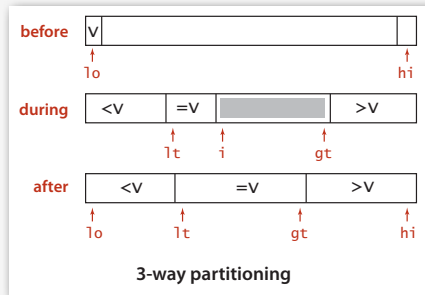
- › quicksort
- › selection
- › duplicate keys
- › system sorts

24

3-way partitioning: Dijkstra's solution

3-way partitioning.

- Let v be partitioning element $a[lo]$.
- Scan i from left to right.
 - $a[i]$ less than v : exchange $a[lt]$ with $a[i]$ and increment both lt and i
 - $a[i]$ greater than v : exchange $a[gt]$ with $a[i]$ and decrement gt
 - $a[i]$ equal to v : increment i



All the right properties.

- In-place.
- Not much code.
- Small overhead if no equal keys.

3-way partitioning: trace


| lt | i | gt | a[] | | | | | | | | | | | |
|----|---|----|-----|---|---|---|---|---|---|---|---|---|----|----|
| | | | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
| 0 | 0 | 11 | R | B | W | W | R | W | B | R | R | W | B | R |
| 0 | 1 | 11 | R | B | W | W | R | W | B | R | R | W | B | R |
| 1 | 2 | 11 | B | R | W | W | R | W | B | R | R | W | B | R |
| 1 | 2 | 10 | B | R | R | W | R | W | B | R | R | W | B | W |
| 1 | 3 | 10 | B | R | R | W | R | W | B | R | R | W | B | W |
| 1 | 3 | 9 | B | R | R | B | R | W | B | R | R | W | W | W |
| 2 | 4 | 9 | B | B | R | R | R | W | B | R | R | W | W | W |
| 2 | 5 | 9 | B | B | R | R | R | W | B | R | R | W | W | W |
| 2 | 5 | 8 | B | B | R | R | R | W | B | R | R | W | W | W |
| 2 | 5 | 7 | B | B | R | R | R | R | B | R | W | W | W | W |
| 2 | 6 | 7 | B | B | R | R | R | B | R | W | W | W | W | W |
| 3 | 7 | 7 | B | B | B | R | R | R | R | R | W | W | W | W |
| 3 | 8 | 7 | B | B | B | R | R | R | R | R | W | W | W | W |

3-way partitioning trace (array contents after each loop iteration)

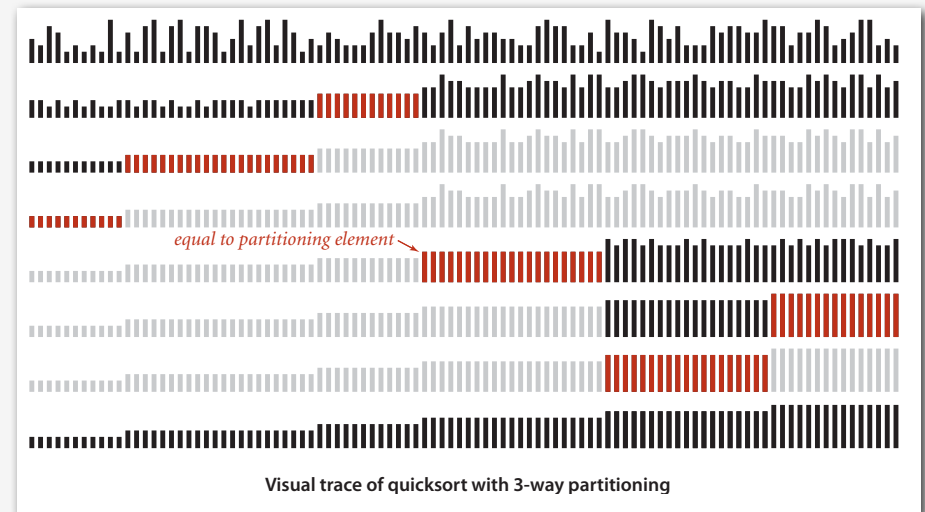
3-way quicksort: Java implementation

```
private static void sort(Comparable[] a, int lo, int hi)
{
    if (hi <= lo) return;
    int lt = lo, gt = hi;
    Comparable v = a[lo];
    int i = lo;
    while (i <= gt)
    {
        int cmp = a[i].compareTo(v);
        if (cmp < 0) exch(a, lt++, i++);
        else if (cmp > 0) exch(a, i, gt--);
        else i++;
    }

    sort(a, lo, lt - 1);
    sort(a, gt + 1, hi);
}
}
```



3-way quicksort: visual trace



Duplicate keys: lower bound

Sorting lower bound. If there are n distinct keys and the i^{th} one occurs x_i times, any compare-based sorting algorithm must use at least

$$\lg\left(\frac{N!}{x_1! x_2! \dots x_n!}\right) \sim -\sum_{i=1}^n x_i \lg \frac{x_i}{N}$$

$N \lg N$ when all distinct;
linear when only a constant number of distinct keys

compares in the worst case.

Proposition. [Sedgwick-Bentley, 1997]

Quicksort with 3-way partitioning is entropy-optimal.

Pf. [beyond scope of course]

proportional to lower bound

Bottom line. Randomized quicksort with 3-way partitioning reduces running time from linearithmic to linear in broad class of applications.

33

- ▶ selection
- ▶ duplicate keys
- ▶ comparators
- ▶ system sorts

34

Sorting applications

Sorting algorithms are essential in a broad variety of applications:

- Sort a list of names.
- Organize an MP3 library.
- Display Google PageRank results. obvious applications
- List RSS news items in reverse chronological order.

- Find the median.
- Find the closest pair.
- Binary search in a database. problems become easy once items are in sorted order
- Identify statistical outliers.
- Find duplicates in a mailing list.

- Data compression.
- Computer graphics.
- Computational biology.
- Supply chain management. non-obvious applications
- Load balancing on a parallel computer.
- ...

Every system needs (and has) a system sort!

35

Java system sorts

Java uses both mergesort and quicksort.

- `Arrays.sort()` sorts array of `Comparable` or any primitive type.
- Uses quicksort for primitive types; mergesort for objects.

```
import java.util.Arrays;

public class StringSort
{
    public static void main(String[] args)
    {
        String[] a = StdIn.readAll().split("\\s+");
        Arrays.sort(a);
        for (int i = 0; i < N; i++)
            StdOut.println(a[i]);
    }
}
```

Q. Why use different algorithms, depending on type?


36

Java system sort for primitive types

Engineering a sort function. [Bentley-McIlroy, 1993]

- Original motivation: improve `qsort()`.
- Basic algorithm = 3-way quicksort with cutoff to insertion sort.
- Partition on Tukey's ninther: median of the medians of 3 samples, each of 3 elements.

approximate median-of-9



Why use Tukey's ninther?

- Better partitioning than random shuffle.
- Less costly than random shuffle.

37

Achilles heel in Bentley-McIlroy implementation (Java system sort)

Based on all this research, Java's system sort is solid, **right?**

A killer input.

more disastrous consequences in C

- Blows function call stack in Java and crashes program.
- Would take quadratic time if it didn't crash first.

```
% more 250000.txt
0
218750
222662
11
166672
247070
83339
...

% java IntegerSort < 250000.txt
Exception in thread "main"
java.lang.StackOverflowError
  at java.util.Arrays.sort1 (Arrays.java:562)
  at java.util.Arrays.sort1 (Arrays.java:606)
  at java.util.Arrays.sort1 (Arrays.java:608)
  at java.util.Arrays.sort1 (Arrays.java:608)
  at java.util.Arrays.sort1 (Arrays.java:608)
  ...
```

250,000 integers between 0 and 250,000

Java's sorting library crashes, even if you give it as much stack space as Windows allows

38

Achilles heel in Bentley-McIlroy implementation (Java system sort)

McIlroy's devious idea. [A Killer Adversary for Quicksort]

- Construct malicious input **while** running system quicksort, in response to elements compared.
- If `v` is partitioning element, commit to $(v < a[i])$ and $(v < a[j])$, but don't commit to $(a[i] < a[j])$ or $(a[j] > a[i])$ until `a[i]` and `a[j]` are compared.

Consequences.

- Confirms theoretical possibility.
- Algorithmic complexity attack: you enter linear amount of data; server performs quadratic amount of work.

Remark. Attack is not effective if array is shuffled before sort.

Q. Why do you think system sort is deterministic?

39

System sort: Which algorithm to use?

Many sorting algorithms to choose from:

Internal sorts.

- Insertion sort, selection sort, bubblesort, shaker sort.
- Quicksort, mergesort, heapsort, samplesort, shellsort.
- Solitaire sort, red-black sort, splay sort, Dobosiewicz sort, psort, ...

External sorts. Poly-phase mergesort, cascade-merge, oscillating sort.

Radix sorts. Distribution, MSD, LSD, 3-way radix quicksort.

Parallel sorts.


- Bitonic sort, Batcher even-odd sort.
- Smooth sort, cube sort, column sort.
- GPU sort.

40

System sort: Which algorithm to use?

Applications have diverse attributes.

- Stable?
- Parallel?
- Deterministic?
- Keys all distinct?
- Multiple key types?
- Linked list or arrays?
- Large or small records?
- Is your array randomly ordered?
- Need guaranteed performance?



many more combinations of attributes than algorithms

Elementary sort may be method of choice for some combination.
Cannot cover **all** combinations of attributes.

- Q. Is the system sort good enough?
A. Usually.

41

Sorting summary

| | inplace? | stable? | worst | average | best | remarks |
|-------------|----------|---------|-----------|-------------|-----------|---|
| selection | x | | $N^2 / 2$ | $N^2 / 2$ | $N^2 / 2$ | N exchanges |
| insertion | x | x | $N^2 / 2$ | $N^2 / 4$ | N | use for small N or partially ordered |
| shell | x | | ? | ? | N | tight code, subquadratic |
| quick | x | | $N^2 / 2$ | $2 N \ln N$ | $N \lg N$ | $N \log N$ probabilistic guarantee fastest in practice |
| 3-way quick | x | | $N^2 / 2$ | $2 N \ln N$ | N | improves quicksort in presence of duplicate keys |
| merge | | x | $N \lg N$ | $N \lg N$ | $N \lg N$ | $N \log N$ guarantee, stable |
| ??? | x | x | $N \lg N$ | $N \lg N$ | $N \lg N$ | holy sorting grail |

42

Which sorting algorithm?

| | | | | | | | |
|------|------|------|------|------|------|------|------|
| lifo | find | data | data | data | data | hash | data |
| fifo | fifo | fifo | fifo | exch | fifo | fifo | exch |
| data | data | find | find | fifo | lifo | data | fifo |
| type | exch | hash | hash | find | type | link | find |
| hash | hash | heap | heap | hash | hash | leaf | hash |
| heap | heap | lifo | lifo | heap | heap | heap | heap |
| sort | less | link | link | leaf | link | exch | leaf |
| link | left | list | list | left | sort | node | left |
| list | leaf | push | push | less | find | lifo | less |
| push | lifo | root | root | lifo | list | left | lifo |
| find | push | sort | sort | link | push | find | link |
| root | root | type | type | list | root | path | list |
| leaf | list | leaf | leaf | sort | leaf | list | next |
| tree | tree | left | tree | tree | null | next | node |
| null | null | node | null | null | path | less | null |
| path | path | null | path | path | tree | root | path |
| node | node | path | node | node | exch | sink | push |
| left | link | tree | left | type | left | swim | root |
| less | sort | exch | less | root | less | null | sink |
| exch | type | less | exch | push | node | sort | sort |
| sink | sink | next | sink | sink | next | type | swap |
| swim | swim | sink | swim | swim | sink | tree | swim |
| next | next | swap | next | next | swap | push | tree |
| swap | swap | swim | swap | swap | swim | swap | type |

original ? ? ? ? ? ? sorted

43