



<http://algs4.cs.princeton.edu>

2.4 PRIORITY QUEUES

- ▶ *API and elementary implementations*
- ▶ *binary heaps*
- ▶ *heapsort*
- ▶ *event-driven simulation*



<http://algs4.cs.princeton.edu>

2.4 PRIORITY QUEUES


- ▶ *API and elementary implementations*
- ▶ *binary heaps*
- ▶ *heapsort*
- ▶ *event-driven simulation*

Collections

A **collection** is a data type that stores a group of items.

| data type | core operations | data structure |
|-----------------------|-----------------------|---------------------------------------|
| stack | PUSH, POP | <i>linked list, resizing array</i> |
| queue | ENQUEUE, DEQUEUE | <i>linked list, resizing array</i> |
| priority queue | INSERT, DELETE-MAX | <i>binary heap</i> |
| symbol table | PUT, GET, DELETE | <i>binary search tree, hash table</i> |
| set | ADD, CONTAINS, DELETE | <i>binary search tree, hash table</i> |

“ Show me your code and conceal your data structures, and I shall continue to be mystified. Show me your data structures, and I won't usually need your code; it'll be obvious. ” — Fred Brooks



Priority queue

Collections. Insert and delete items. Which item to delete?

Stack. Remove the item most recently added.

Queue. Remove the item least recently added.

Randomized queue. Remove a random item.

Priority queue. Remove the item with
the **largest (or smallest) key**.


Generalizes: stack, queue, randomized queue.

| <i>operation</i> | <i>argument</i> | <i>return value</i> |
|-------------------|-----------------|---------------------|
| <i>insert</i> | P | |
| <i>insert</i> | Q | |
| <i>insert</i> | E | |
| <i>remove max</i> | | Q |
| <i>insert</i> | X | |
| <i>insert</i> | A | |
| <i>insert</i> | M | |
| <i>remove max</i> | | X |
| <i>insert</i> | P | |
| <i>insert</i> | L | |
| <i>insert</i> | E | |
| <i>remove max</i> | | P |

Priority queue API

Requirement. Items are generic; they must also be Comparable.

Key must be Comparable
(bounded type parameter)




```
public class MaxPQ<Key extends Comparable<Key>>
```

| | |
|--------------------|--|
| MaxPQ() | <i>create an empty priority queue</i> |
| MaxPQ(Key[] a) | <i>create a priority queue with given keys</i> |
| void insert(Key v) | <i>insert a key into the priority queue</i> |
| Key delMax() | <i>return and remove a largest key</i> |
| boolean isEmpty() | <i>is the priority queue empty?</i> |
| Key max() | <i>return a largest key</i> |
| int size() | <i>number of entries in the priority queue</i> |

Note. Duplicate keys allowed; delMax() picks any maximum key.

Priority queue: applications

- Event-driven simulation. [customers in a line, colliding particles]
- Numerical computation. [reducing roundoff error]
- Discrete optimization. [bin packing, scheduling]
- Artificial intelligence. [A* search]
- Computer networks. [web cache]
- Operating systems. [load balancing, interrupt handling]
- Data compression. [Huffman codes]
- Graph searching. [Dijkstra's algorithm, Prim's algorithm]
- Number theory. [sum of powers]
- Spam filtering. [Bayesian spam filter]
- Statistics. [online median in data stream]



| | | |
|---|---|---|
| 8 | 4 | 7 |
| 1 | 5 | 6 |
| 3 | 2 | |

Priority queue: client example

Challenge. Find the largest M items in a stream of N items.

- Fraud detection: isolate \$\$ transactions.
- NSA monitoring: flag most suspicious documents.

N huge, M large

Constraint. Not enough memory to store N items.

```
% more transactions.txt
Turing      6/17/1990    644.08
vonNeumann  3/26/2002    4121.85
Dijkstra    8/22/2007    2678.40
vonNeumann  1/11/1999    4409.74
Dijkstra    11/18/1995   837.42
Hoare       5/10/1993    3229.27
vonNeumann  2/12/1994    4732.35
Hoare       8/18/1992    4381.21
Turing      1/11/2002    66.10
Thompson    2/27/2000    4747.08
Turing      2/11/1991    2156.86
Hoare       8/12/2003    1025.70
vonNeumann  10/13/1993   2520.97
Dijkstra    9/10/2000    708.95
Turing      10/12/1993   3532.36
Hoare       2/10/2005    4050.20
```

```
% java TopM 5 < transactions.txt
Thompson    2/27/2000    4747.08
vonNeumann  2/12/1994    4732.35
vonNeumann  1/11/1999    4409.74
Hoare       8/18/1992    4381.21
vonNeumann  3/26/2002    4121.85
```

sort key

Priority queue: client example

Challenge. Find the largest M items in a stream of N items.

- Fraud detection: isolate \$\$ transactions.
- NSA monitoring: flag most suspicious documents.

N huge, M large

Constraint. Not enough memory to store N items.

*Transaction data
type is Comparable
(ordered by \$\$)*

use a min-oriented pq

```
MinPQ<Transaction> pq = new MinPQ<Transaction>();  
while (StdIn.hasNextLine())  
{  
    String line = StdIn.readLine();  
    Transaction transaction = new Transaction(line);  
    pq.insert(transaction);  
    if (pq.size() > M)  
        pq.delMin();  
}
```

*pq now contains
largest M items*

Priority queue: unordered and ordered array implementation

| <i>operation</i> | <i>argument</i> | <i>return value</i> | <i>size</i> | <i>contents (unordered)</i> | | | | | <i>contents (ordered)</i> | | | | | | | | | |
|-------------------|-----------------|---------------------|-------------|-----------------------------|---|---|---|---|---------------------------|---|---|---|---|---|---|---|---|---|
| <i>insert</i> | P | | 1 | P | | | | | | P | | | | | | | | |
| <i>insert</i> | Q | | 2 | P | Q | | | | | P | Q | | | | | | | |
| <i>insert</i> | E | | 3 | P | Q | E | | | | E | P | Q | | | | | | |
| <i>remove max</i> | | Q | 2 | P | E | | | | | E | P | | | | | | | |
| <i>insert</i> | X | | 3 | P | E | X | | | | E | P | X | | | | | | |
| <i>insert</i> | A | | 4 | P | E | X | A | | | A | E | P | X | | | | | |
| <i>insert</i> | M | | 5 | P | E | X | A | M | | A | E | M | P | X | | | | |
| <i>remove max</i> | | X | 4 | P | E | M | A | | | A | E | M | P | | | | | |
| <i>insert</i> | P | | 5 | P | E | M | A | P | | A | E | M | P | P | | | | |
| <i>insert</i> | L | | 6 | P | E | M | A | P | L | | A | E | L | M | P | | | |
| <i>insert</i> | E | | 7 | P | E | M | A | P | L | E | | A | E | E | L | M | P | P |
| <i>remove max</i> | | P | 6 | E | M | A | P | L | E | | | A | E | E | L | M | P | |

A sequence of operations on a priority queue

Priority queue: implementations cost summary

Challenge. Implement **all** operations efficiently.

| implementation | insert | del max | max |
|------------------------|----------|----------|----------|
| unordered array | 1 | N | N |
| ordered array | N | 1 | 1 |
| goal for today | $\log N$ | $\log N$ | $\log N$ |

order of growth of running time for priority queue with N items



<http://algs4.cs.princeton.edu>


2.4 PRIORITY QUEUES

- ▶ *API and elementary implementations*
- ▶ *binary heaps*
- ▶ *heapsort*
- ▶ *event-driven simulation*

Complete binary tree

Binary tree. Empty **or** node with links to left and right binary trees.

Complete tree. Perfectly balanced, except for bottom level.



complete binary tree with $N = 16$ nodes (height = 4)

Property. Height of complete binary tree with N nodes is $\lfloor \lg N \rfloor$.

Pf. Height increases only when N is a power of 2.

A complete binary tree in nature




Hyphaene Compressa - Doum Palm

© Shlomit Pinter

Complete binary tree: array representation

Array representation.


- Indices start at 1.
- Take nodes in **level** order.
- Children of node k at locations $2*k$ and $2*k+1$
- No explicit links needed!



Priority queues: quiz 1

What is the index of the parent of the item at index k in a binary heap?

- A. $k/2 - 1$
- B. $k/2$
- C. $k/2 + 1$
- D. *None of the above.*
- E. *I don't know.*



| | | | | | | | | | | | | |
|--------|---|---|---|---|---|---|---|---|---|---|----|----|
| i | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
| $a[i]$ | - | T | S | R | P | N | O | A | E | I | H | G |


Binary heap

Array representation.

- Indices start at 1.
- Take nodes in **level** order.
- Children of node k at locations $2*k$ and $2*k+1$
- No explicit links needed!

Max-Heap ordering.

- Keys in nodes.
- Parent's key no smaller than children's keys.
- “Just enough” ordering to support efficient priority queue operations.




Binary heap. Array representation of a heap-ordered complete binary tree.

Binary heap: properties

“Just enough” ordering to support efficient priority queue operations.

- Largest key is $a[1]$, which is the root of the binary tree.
- Can use array indices to move through the tree.
 - Children of node at k at locations $2*k$ and $2*k+1$.
 - Parent of node at k is at $k/2$.




- $insert()$ and $delMax()$ violate heap order, but easy to fix up.

Binary heap demo

Insert. Add node at end, then swim it up.

Remove the maximum. Exchange root with node at end, then sink it down.

heap ordered




Binary heap demo

Insert. Add node at end, then swim it up.

Remove the maximum. Exchange root with node at end, then sink it down.

heap ordered



Binary heap: promotion


Scenario. A key becomes **larger** than its parent's key.

To eliminate the violation:

- Exchange key in child with key in parent.
- Repeat until heap order restored.

```
private void swim(int k)
{
    while (k > 1 && less(k/2, k))
    {
        exch(k, k/2);
        k = k/2;
    }
}
```

parent of node at k is at k/2




Binary heap: insertion

Insert. Add node at end, then swim it up.

Cost. At most $1 + \lg N$ compares.

```
public void insert(Key x)
{
    pq[++N] = x;
    swim(N);
}
```



Binary heap: demotion

Scenario. A key becomes **smaller** than one (or both) of its children's.


To eliminate the violation:

why not smaller child?

- Exchange key in parent with key in larger child.
- Repeat until heap order restored.

```
private void sink(int k)
{
    while (2*k <= N)
    {
        int j = 2*k;
        if (j < N && less(j, j+1)) j++;
        if (!less(k, j)) break;
        exch(k, j);
        k = j;
    }
}
```

children of node at k
are $2*k$ and $2*k+1$






Binary heap: delete the maximum

Delete max. Exchange root with node at end, then sink it down.

Cost. At most $2 \lg N$ compares.

```
public Key delMax()
{
    Key max = pq[1];
    exch(1, N);
    pq[N--] = null; ← prevent loitering
    sink(1);
    return max;
}
```

remove the maximum




Priority queue: implementations cost summary

| implementation | insert | del max | max |
|--------------------|----------|----------|-----|
| unordered array | 1 | N | N |
| ordered array | N | 1 | 1 |
| binary heap | $\log N$ | $\log N$ | 1 |

order-of-growth of running time for priority queue with N items


DELETE-RANDOM FROM A BINARY HEAP

Goal. Delete a random key from a binary heap in logarithmic time.



DELETE-RANDOM FROM A BINARY HEAP

Goal. Delete a random key from a binary heap in logarithmic time.




Solution.

- Pick a random index r between 1 and N .
- Perform $\text{exch}(r, N--)$.
- Perform either $\text{sink}(r)$ or $\text{swim}(r)$.

DELETE-RANDOM FROM A BINARY HEAP

Goal. Delete a random key from a binary heap in logarithmic time.




Solution.

- Pick a random index r between 1 and N .
- Perform $\text{exch}(r, N--)$.
- Perform either $\text{sink}(r)$ or $\text{swim}(r)$.

Binary heap: practical improvements

Do "half-exchanges" in sink and swim.

- Reduces number of array accesses.
- Worth doing.

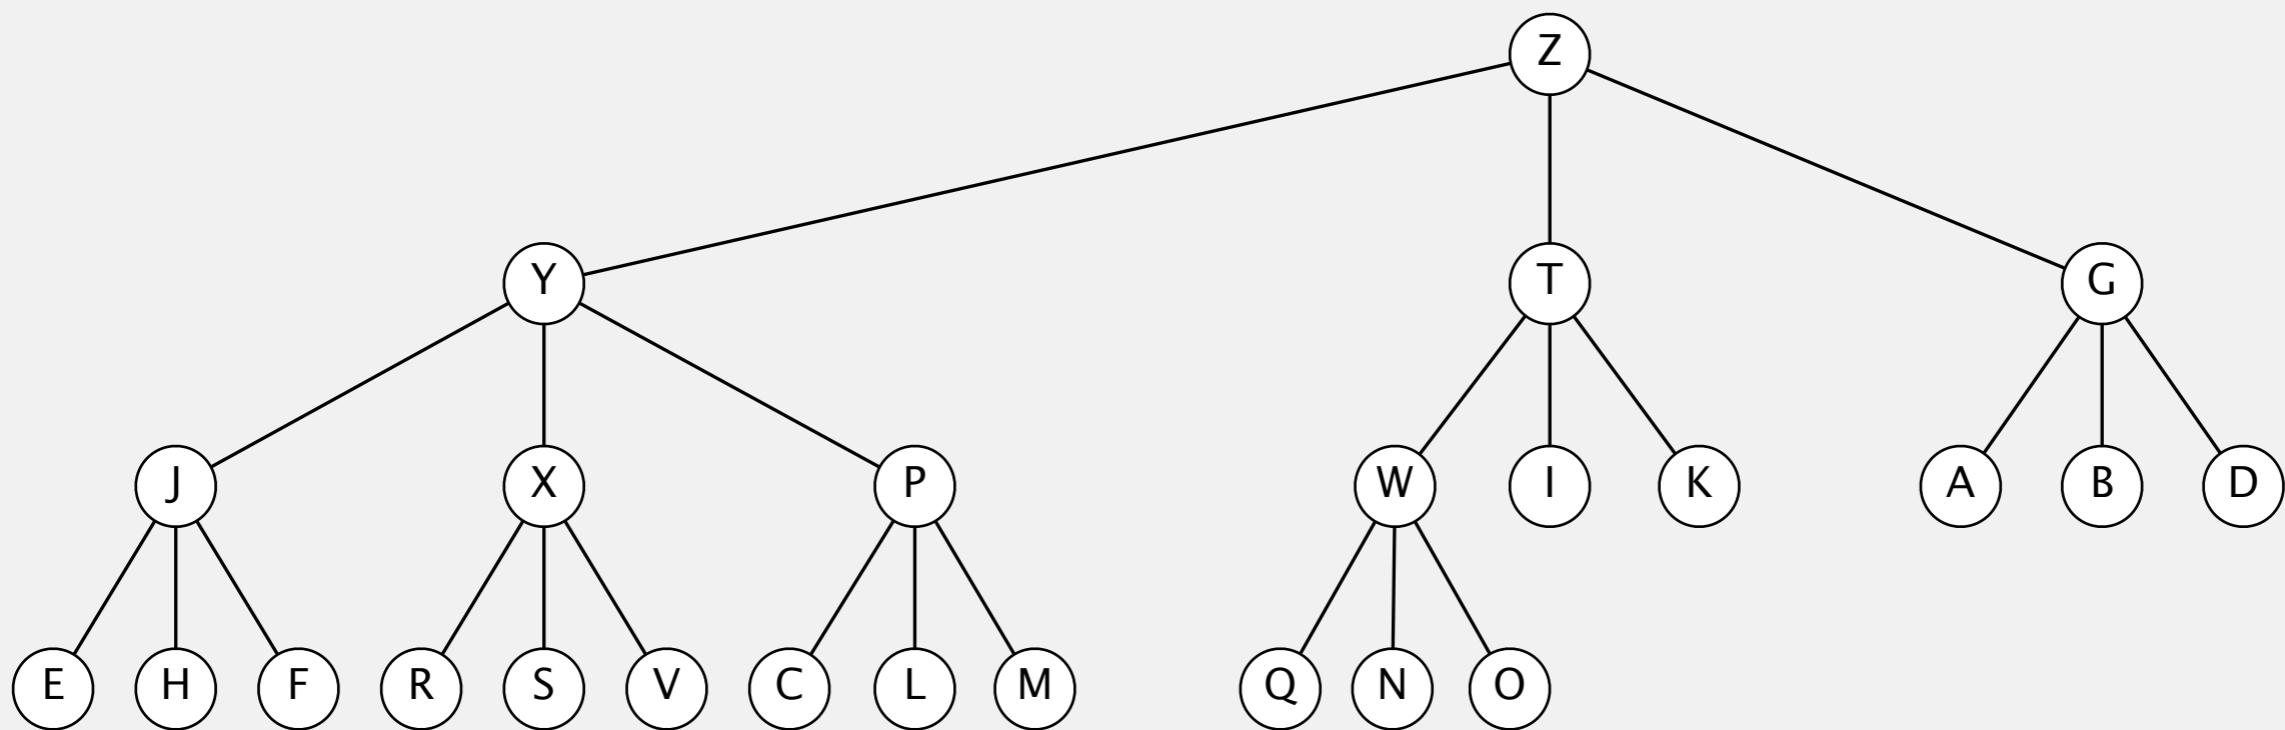


Binary heap: practical improvements

Multiway heaps.

- Complete d -way tree.
- Parent's key no smaller than **any** of its children's keys.

Fact. Height of complete d -way tree on N nodes is $\sim \log_d N$.



3-way heap

Priority queues: quiz 2

How many compares (in the worst case) to **insert** in a d -way heap?

- A. $\sim \log_2 N$
- B. $\sim \log_d N$
- C. $\sim d \log_2 N$
- D. $\sim d \log_d N$
- E. *I don't know.*

Priority queues: quiz 3

How many compares (in the worst case) to **delete-max** in a d -way heap?

- A. $\sim \log_2 N$
- B. $\sim \log_d N$
- C. $\sim d \log_2 N$
- D. $\sim d \log_d N$
- E. *I don't know.*

Priority queue: implementation cost summary

| implementation | insert | del max | max |
|------------------------|------------|------------------|-----|
| unordered array | 1 | N | N |
| ordered array | N | 1 | 1 |
| binary heap | $\log N$ | $\log N$ | 1 |
| d-ary heap | $\log_d N$ | $d \log_d N$ | 1 |
| Fibonacci | 1 | $\log N^\dagger$ | 1 |
| Brodal queue | 1 | $\log N$ | 1 |
| impossible | 1 | 1 | 1 |

← sweet spot: $d = 4$

← why impossible?

† amortized

order-of-growth of running time for priority queue with N items

Binary heap: considerations

Underflow and overflow.

- Underflow: throw exception if deleting from empty PQ.
- Overflow: add no-arg constructor and use resizing array.

leads to $\log N$
amortized time per op
(how to make worst case?)

Minimum-oriented priority queue.

- Replace `less()` with `greater()`.
- Implement `greater()`.

Other operations.

- Remove an arbitrary item.
- Change the priority of an item.

can implement efficiently with `sink()` and `swim()`
[stay tuned for Prim/Dijkstra]

Immutability of keys.

- Assumption: client does not change keys while they're on the PQ.
- Best practice: use immutable keys.



<http://algs4.cs.princeton.edu>

2.4 PRIORITY QUEUES

- ▶ *API and elementary implementations*
- ▶ *binary heaps*
- ▶ *heapsort*
- ▶ *event-driven simulation*

Priority queues: quiz 4

What are the properties of the following algorithm?

```
public void sort(String[] a)
{
    int N = a.length;
    MaxPQ<String> pq = new MaxPQ<String>();
    for (int i = 0; i < N; i++)
        pq.insert(a[i]);
    for (int i = N-1; i >= 0; i--)
        a[i] = pq.delMax();
}
```


- A.** $N \log N$ compares in the worst case.
- B.** In-place.
- C.** Stable.
- D.** *All of the above.*
- E.** *I don't know.*

Heapsort

Basic plan for in-place sort.


- View input array as a complete binary tree.
- Heap construction: build a max-heap with all N keys.
- Sortdown: repeatedly remove the maximum key.

keys in arbitrary order




| | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|----|----|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
| S | O | R | T | E | X | A | M | P | L | E |

build max heap
(in place)



| | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|----|----|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
| X | T | S | P | L | R | A | M | O | E | E |

sorted result
(in place)




| | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|----|----|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
| A | E | E | L | M | O | P | R | S | T | X |

Heapsort demo

Heap construction. Build max heap using bottom-up method.

we assume array entries are indexed 1 to N


array in arbitrary order



Heapsort demo

Sortdown. Repeatedly delete the largest remaining item.

array in sorted order









| | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|----|----|
| A | E | E | L | M | O | P | R | S | T | X |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |

Heapsort: heap construction

First pass. Build heap using bottom-up method.

```
for (int k = N/2; k >= 1; k--)  
    sink(a, k, N);
```




Heapsort: sortdown

Second pass.

- Remove the maximum, one at a time.
- Leave in array, instead of nulling out.

```

while (N > 1)
{
    exch(a, 1, N--);
    sink(a, 1, N);
}
    
```



Heapsort: Java implementation

```
public class Heap
{
    public static void sort(Comparable[] a)
    {
        int N = a.length;
        for (int k = N/2; k >= 1; k--)
            sink(a, k, N);
        while (N > 1)
        {
            exch(a, 1, N);
            sink(a, 1, --N);
        }
    }
}
```

but make static (and pass arguments)

```
private static void sink(Comparable[] a, int k, int N)
{ /* as before */ }
```

```
private static boolean less(Comparable[] a, int i, int j)
{ /* as before */ }
```

```
private static void exch(Object[] a, int i, int j)
{ /* as before */ }
```

```
}
```

but convert from 1-based
indexing to 0-base indexing

Heapsort: trace


| | | a[i] | | | | | | | | | | | |
|-----------------------|---|------|---|---|---|---|---|---|---|---|---|----|----|
| N | k | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
| <i>initial values</i> | | S | O | R | T | E | X | A | M | P | L | E | |
| 11 | 5 | S | O | R | T | L | X | A | M | P | E | E | |
| 11 | 4 | S | O | R | T | L | X | A | M | P | E | E | |
| 11 | 3 | S | O | X | T | L | R | A | M | P | E | E | |
| 11 | 2 | S | T | X | P | L | R | A | M | O | E | E | |
| 11 | 1 | X | T | S | P | L | R | A | M | O | E | E | |
| <i>heap-ordered</i> | | X | T | S | P | L | R | A | M | O | E | E | |
| 10 | 1 | T | P | S | O | L | R | A | M | E | E | X | |
| 9 | 1 | S | P | R | O | L | E | A | M | E | T | X | |
| 8 | 1 | R | P | E | O | L | E | A | M | S | T | X | |
| 7 | 1 | P | O | E | M | L | E | A | R | S | T | X | |
| 6 | 1 | O | M | E | A | L | E | P | R | S | T | X | |
| 5 | 1 | M | L | E | A | E | O | P | R | S | T | X | |
| 4 | 1 | L | E | E | A | M | O | P | R | S | T | X | |
| 3 | 1 | E | A | E | L | M | O | P | R | S | T | X | |
| 2 | 1 | E | A | E | L | M | O | P | R | S | T | X | |
| 1 | 1 | A | E | E | L | M | O | P | R | S | T | X | |
| <i>sorted result</i> | | A | E | E | L | M | O | P | R | S | T | X | |

Heapsort trace (array contents just after each sink)

Heapsort: mathematical analysis

Proposition. Heap construction makes $\leq N$ exchanges and $\leq 2N$ compares.

Pf sketch. [assume $N = 2^{h+1} - 1$]



binary heap of height $h = 3$

$$\begin{aligned} h + 2(h - 1) + 4(h - 2) + 8(h - 3) + \dots + 2^h(0) &= 2^{h+1} - h - 2 \\ &= N - (h + 1) \\ &\leq N \end{aligned}$$

a tricky sum
(see COS 340)

Heapsort: mathematical analysis

Proposition. Heap construction makes $\leq N$ exchanges and $\leq 2N$ compares.

Proposition. Heapsort uses $\leq 2N \lg N$ compares and exchanges.

algorithm can be improved to $\sim 1 N \lg N$
(but no such variant is known to be practical)

Significance. In-place sorting algorithm with $N \log N$ worst-case.

- Mergesort: no, linear extra space. ← in-place merge possible, not practical
- Quicksort: no, quadratic time in worst case. ← $N \log N$ worst-case quicksort possible, not practical
- Heapsort: yes!

Bottom line. Heapsort is optimal for both time and space, **but:**


- Inner loop longer than quicksort's.
- Makes poor use of cache.
- Not stable. ← can be improved using advanced caching tricks

Introsort

Goal. As fast as quicksort in practice; $N \log N$ worst case, in place.

Introsort.

- Run quicksort.
- Cutoff to heapsort if stack depth exceeds $2 \lg N$.
- Cutoff to insertion sort for $N = 16$.



Introspective Sorting and Selection Algorithms

David R. Musser*
Computer Science Department
Rensselaer Polytechnic Institute, Troy, NY 12180
musser@cs.rpi.edu

Abstract

Quicksort is the preferred in-place sorting algorithm in many contexts, since its average computing time on uniformly distributed inputs is $\Theta(N \log N)$ and it is in fact faster than most other sorting algorithms on most inputs. Its drawback is that its worst-case time bound is $\Theta(N^2)$. Previous attempts to protect against the worst case by improving the way quicksort chooses pivot elements for partitioning have increased the average computing time too much—one might as well use heapsort, which has a $\Theta(N \log N)$ worst-case time bound but is on the average 2 to 5 times slower than quicksort. A similar dilemma exists with selection algorithms (for finding the i -th largest element) based on partitioning. This paper describes a simple solution to this dilemma: limit the depth of partitioning, and for subproblems that exceed the limit switch to another algorithm with a better worst-case bound. Using heapsort as the “stopper” yields a sorting algorithm that is just as fast as quicksort in the average case but also has an $\Theta(N \log N)$ worst case time bound. For selection, a hybrid of Hoare’s FIND algorithm, which is linear on average but quadratic in the worst case, and the Blum-Floyd-Pratt-Rivest-Tarjan algorithm is as fast as Hoare’s algorithm in practice, yet has a linear worst-case time bound. Also discussed are issues of implementing the new algorithms as generic algorithms and accurately measuring their performance in the framework of the C++ Standard Template Library.

In the wild. C++ STL, Microsoft .NET Framework.

Sorting algorithms: summary

| | inplace? | stable? | best | average | worst | remarks |
|-------------|----------|---------|-----------------------|-------------------|-------------------|--|
| selection | ✓ | | $\frac{1}{2} N^2$ | $\frac{1}{2} N^2$ | $\frac{1}{2} N^2$ | N exchanges |
| insertion | ✓ | ✓ | N | $\frac{1}{4} N^2$ | $\frac{1}{2} N^2$ | use for small N or partially ordered |
| shell | ✓ | | $N \log_3 N$ | ? | $c N^{3/2}$ | tight code; subquadratic |
| merge | | ✓ | $\frac{1}{2} N \lg N$ | $N \lg N$ | $N \lg N$ | $N \log N$ guarantee; stable |
| timsort | | ✓ | N | $N \lg N$ | $N \lg N$ | improves mergesort when preexisting order |
| quick | ✓ | | $N \lg N$ | $2 N \ln N$ | $\frac{1}{2} N^2$ | $N \log N$ probabilistic guarantee; fastest in practice |
| 3-way quick | ✓ | | N | $2 N \ln N$ | $\frac{1}{2} N^2$ | improves quicksort when duplicate keys |
| heap | ✓ | | $3 N$ | $2 N \lg N$ | $2 N \lg N$ | $N \log N$ guarantee; in-place |
| ? | ✓ | ✓ | N | $N \lg N$ | $N \lg N$ | holy sorting grail |




<http://algs4.cs.princeton.edu>

2.4 PRIORITY QUEUES

- ▶ *API and elementary implementations*
- ▶ *binary heaps*
- ▶ *heapsort*
- ▶ *event-driven simulation*

Molecular dynamics simulation of hard discs

Goal. Simulate the motion of N moving particles that behave according to the laws of elastic collision.



Molecular dynamics simulation of hard discs

Goal. Simulate the motion of N moving particles that behave according to the laws of elastic collision.

Hard disc model.

- Moving particles interact via elastic collisions with each other and walls.
- Each particle is a disc with known position, velocity, mass, and radius.
- No other forces.

temperature, pressure,
diffusion constant

motion of individual
atoms and molecules

Significance. Relates macroscopic observables to microscopic dynamics.

- Maxwell-Boltzmann: distribution of speeds as a function of temperature.
- Einstein: explain Brownian motion of pollen grains.

Warmup: bouncing balls

Time-driven simulation. N bouncing balls in the unit square.

```
public class BouncingBalls
{
    public static void main(String[] args)
    {
        int N = Integer.parseInt(args[0]);
        Ball[] balls = new Ball[N];
        for (int i = 0; i < N; i++)
            balls[i] = new Ball();
        while(true)
        {
            StdDraw.clear();
            for (int i = 0; i < N; i++)
            {
                balls[i].move(0.5);
                balls[i].draw();
            }
            StdDraw.show(50);
        }
    }
}
```

↑
main simulation loop

```
% java BouncingBalls 100
```




Warmup: bouncing balls

```
public class Ball
{
    private double rx, ry;          // position
    private double vx, vy;          // velocity
    private final double radius;    // radius
    public Ball(...)
    { /* initialize position and velocity */ }

    public void move(double dt)
    {
        if ((rx + vx*dt < radius) || (rx + vx*dt > 1.0 - radius)) { vx = -vx; }
        if ((ry + vy*dt < radius) || (ry + vy*dt > 1.0 - radius)) { vy = -vy; }
        rx = rx + vx*dt;
        ry = ry + vy*dt;
    }
    public void draw()
    { StdDraw.filledCircle(rx, ry, radius); }
}
```

check for collision with walls




Missing. Check for balls colliding with **each other**.

- Physics problems: when? what effect?
- CS problems: which object does the check? too many checks?

Time-driven simulation

- Discretize time in quanta of size dt .
- Update the position of each particle after every dt units of time, and check for overlaps.
- If overlap, roll back the clock to the time of the collision, update the velocities of the colliding particles, and continue the simulation.




Time-driven simulation


Main drawbacks.

- $\sim N^2/2$ overlap checks per time quantum.
- Simulation is too slow if dt is very small.
- May miss collisions if dt is too large.
(if colliding particles fail to overlap when we are looking)

dt too small: excessive computation



dt too large: may miss collisions




Event-driven simulation

Change state only when something interesting happens.

- Between collisions, particles move in straight-line trajectories.
- Focus only on times when collisions occur.
- Maintain **PQ** of collision events, prioritized by time.
- Delete min = get next collision.

Collision prediction. Given position, velocity, and radius of a particle, when will it collide next with a wall or another particle?

Collision resolution. If collision occurs, update colliding particle(s) according to laws of elastic collisions.




Particle-wall collision

Collision prediction and resolution.

- Particle of radius s at position (r_x, r_y) .
- Particle moving in unit box with velocity (v_x, v_y) .
- Will it collide with a vertical wall? If so, when?

prediction (at time t)

$$\begin{aligned} dt &\equiv \text{time to hit wall} \\ &= \text{distance/velocity} \\ &= (1 - s - r_x) / v_x \end{aligned}$$




Predicting and resolving a particle-wall collision

Particle-particle collision prediction

Collision prediction.

- Particle i : radius s_i , position (rx_i, ry_i) , velocity (vx_i, vy_i) .
- Particle j : radius s_j , position (rx_j, ry_j) , velocity (vx_j, vy_j) .
- Will particles i and j collide? If so, when?



Particle-particle collision prediction

Collision prediction.

- Particle i : radius s_i , position (rx_i, ry_i) , velocity (vx_i, vy_i) .
- Particle j : radius s_j , position (rx_j, ry_j) , velocity (vx_j, vy_j) .
- Will particles i and j collide? If so, when?

$$\Delta t = \begin{cases} \infty & \text{if } \Delta v \cdot \Delta r \geq 0, \\ \infty & \text{if } d < 0, \\ -\frac{\Delta v \cdot \Delta r + \sqrt{d}}{\Delta v \cdot \Delta v} & \text{otherwise} \end{cases}$$

$$d = (\Delta v \cdot \Delta r)^2 - (\Delta v \cdot \Delta v) (\Delta r \cdot \Delta r - s^2), \quad s = s_i + s_j$$

$$\Delta v = (\Delta vx, \Delta vy) = (vx_i - vx_j, vy_i - vy_j)$$

$$\Delta r = (\Delta rx, \Delta ry) = (rx_i - rx_j, ry_i - ry_j)$$

$$\Delta v \cdot \Delta v = (\Delta vx)^2 + (\Delta vy)^2$$

$$\Delta r \cdot \Delta r = (\Delta rx)^2 + (\Delta ry)^2$$

$$\Delta v \cdot \Delta r = (\Delta vx)(\Delta rx) + (\Delta vy)(\Delta ry)$$

Important note: This is physics, so we won't be testing you on it!

Particle-particle collision resolution

Collision resolution. When two particles collide, how does velocity change?

$$\begin{aligned}vx_i' &= vx_i + Jx / m_i \\vy_i' &= vy_i + Jy / m_i \\vx_j' &= vx_j - Jx / m_j \\vy_j' &= vy_j - Jy / m_j\end{aligned}$$

← Newton's second law
(momentum form)

$$Jx = \frac{J \Delta r x}{s}, \quad Jy = \frac{J \Delta r y}{s}, \quad J = \frac{2 m_i m_j (\Delta v \cdot \Delta r)}{s (m_i + m_j)}$$

impulse due to normal force
(conservation of energy, conservation of momentum)

Important note: This is physics, so we won't be testing you on it!

Particle data type skeleton

```
public class Particle
{
    private double rx, ry;           // position
    private double vx, vy;           // velocity
    private final double radius;     // radius
    private final double mass;       // mass
    private int count;               // number of collisions

    public Particle( ... )           { ... }

    public void move(double dt) { ... }
    public void draw()           { ... }

    public double timeToHit(Particle that) { }
    public double timeToHitVerticalWall() { }
    public double timeToHitHorizontalWall() { }

    public void bounceOff(Particle that) { }
    public void bounceOffVerticalWall() { }
    public void bounceOffHorizontalWall() { }

}

    predict collision
    with particle or wall

    resolve collision
    with particle or wall
```

Particle-particle collision and resolution implementation

```
public double timeToHit(Particle that)
{
    if (this == that) return INFINITY;
    double dx = that.rx - this.rx, dy = that.ry - this.ry;
    double dvx = that.vx - this.vx; dvy = that.vy - this.vy;
    double dvdr = dx*dvx + dy*dvy;
    if( dvdr > 0) return INFINITY; ← no collision
    double dvdv = dvx*dvx + dvy*dvy;
    double drdr = dx*dx + dy*dy;
    double s = this.radius + that.radius;
    double d = (dvdr*dvdr) - dvdv * (drdr - s*s); ←
    if (d < 0) return INFINITY;
    return -(dvdr + Math.sqrt(d)) / dvdv;
}
```

```
public void bounceOff(Particle that)
{
    double dx = that.rx - this.rx, dy = that.ry - this.ry;
    double dvx = that.vx - this.vx, dvy = that.vy - this.vy;
    double dvdr = dx*dvx + dy*dvy;
    double s = this.radius + that.radius;
    double J = 2 * this.mass * that.mass * dvdr / (s * (this.mass + that.mass));
    double Jx = J * dx / s;
    double Jy = J * dy / s;
    this.vx += Jx / this.mass;
    this.vy += Jy / this.mass;
    that.vx -= Jx / that.mass;
    that.vy -= Jy / that.mass;
    this.count++;
    that.count++;
}
```


Important note: This is physics, so we won't be testing you on it!

Collision system: event-driven simulation main loop

Initialization.

- Fill PQ with all potential particle-wall collisions.
- Fill PQ with all potential particle-particle collisions.

“potential” since collision is invalidated
if some other collision intervenes



Main loop.

- Delete the impending event from PQ (min priority = t).
- If the event has been invalidated, ignore it.
- Advance all particles to time t , on a straight-line trajectory.
- Update the velocities of the colliding particle(s).
- Predict future particle-wall and particle-particle collisions involving the colliding particle(s) and insert events onto PQ.

Event data type

Conventions.

- Neither particle null \Rightarrow particle-particle collision.
- One particle null \Rightarrow particle-wall collision.
- Both particles null \Rightarrow redraw event.

```
private static class Event implements Comparable<Event>
{
    private final double time;           // time of event
    private final Particle a, b;        // particles involved in event
    private final int countA, countB;   // collision counts of a and b

    public Event(double t, Particle a, Particle b)           create event
    { ... }

    public int compareTo(Event that)                       ordered by time
    { return this.time - that.time; }

    public boolean isValid()                               valid if no intervening collisions
    { ... }                                                (compare collision counts)
}
```

Collision system implementation: main event-driven simulation loop

```
public void simulate()
{
    pq = new MinPQ<Event>();
    for(int i = 0; i < N; i++) predict(particles[i]);
    pq.insert(new Event(0, null, null));

    while(!pq.isEmpty())
    {
        Event event = pq.delMin();
        if(!event.isValid()) continue;
        Particle a = event.a;
        Particle b = event.b;

        for(int i = 0; i < N; i++)
            particles[i].move(event.time - t);
        t = event.time;

        if (a != null && b != null) a.bounceOff(b);
        else if (a != null && b == null) a.bounceOffVerticalWall();
        else if (a == null && b != null) b.bounceOffHorizontalWall();
        else if (a == null && b == null) redraw();

        predict(a);
        predict(b);
    }
}
```

initialize PQ with collision events and redraw event

get next event


update positions and time

process event

predict new events based on changes


Particle collision simulation: example 1

```
% java CollisionSystem 100
```




Particle collision simulation: example 2

```
% java CollisionSystem < billiards.txt
```



Particle collision simulation: example 3

```
% java CollisionSystem < brownian.txt
```



Particle collision simulation: example 4

```
% java CollisionSystem < diffusion.txt
```

