

The challenges of implementing Dijkstra's algorithm

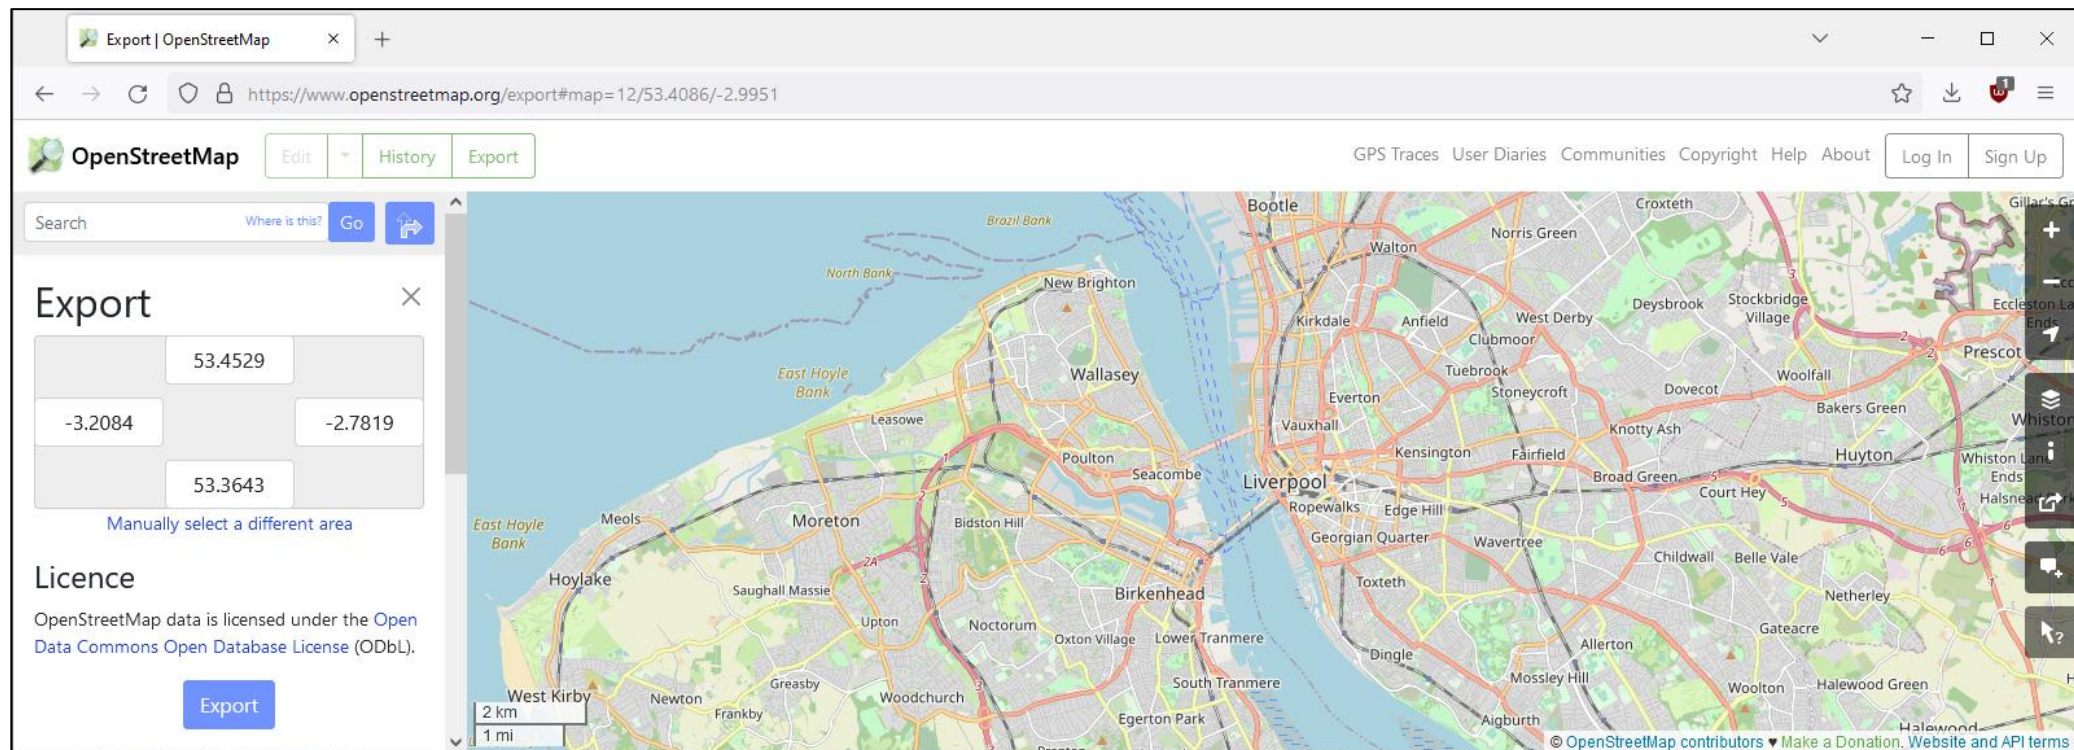
Gerth Stølting Brodal
Aarhus University

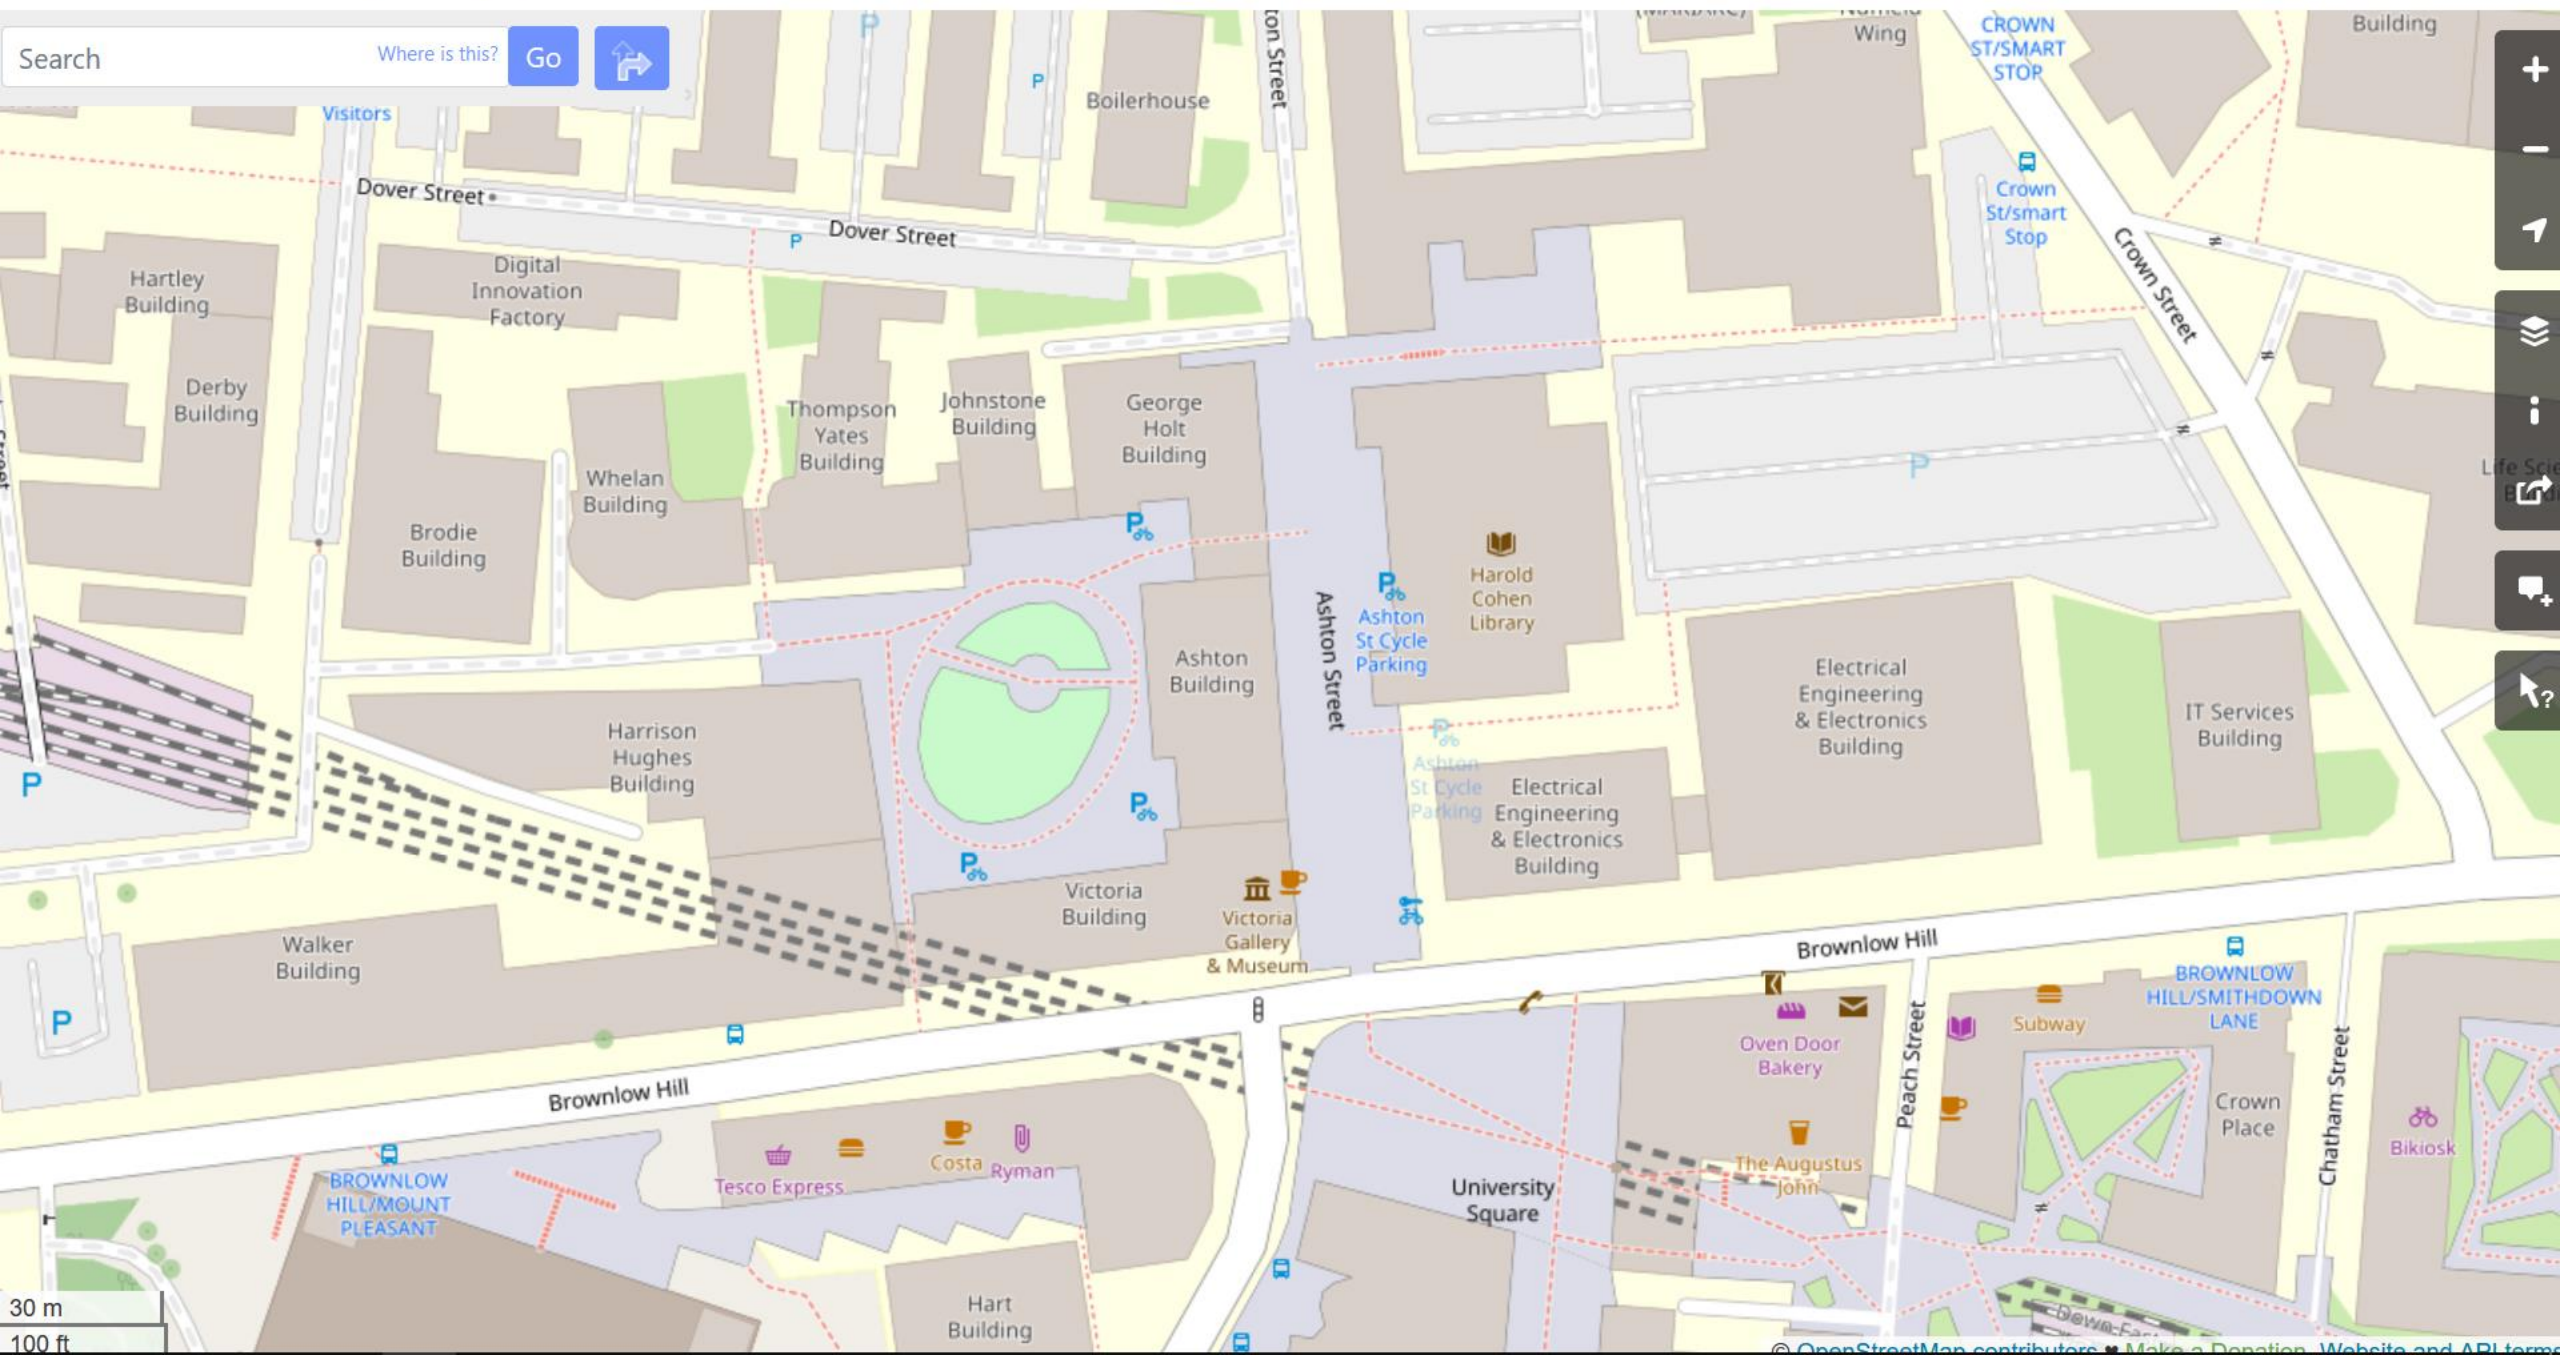
(Work presented at the 11th International Conference on Fun with Algorithms, FUN 2022)

Department of Computer Science, University of Liverpool, UK, January 9, 2023

Background

- Bachelorproject = shortest paths on Open Street Map graphs
- Students have trouble implementing Dijkstra's algorithm in Java™






30 m

100 ft

```
<node id="4606377719" visible="true" version="1" changeset="45143381" timestamp="2017-01-13T19:11:48Z"
user="Dyserth" uid="1892838" lat="53.4064488" lon="-2.9660774"/>
<way id="4961017" visible="true" version="11" changeset="66454995" timestamp="2019-01-19T14:25:32Z"
user="Redhochs" uid="1744940">
  <nd ref="348087292"/>
  <nd ref="4606377719"/>
  <nd ref="268780437"/>
  <nd ref="2700274352"/>
  <nd ref="268779145"/>
  <tag k="highway" v="pedestrian"/>
  <tag k="lit" v="yes"/>
  <tag k="name" v="Ashton Street"/>
  <tag k="surface" v="paving_stones"/>
</way>
<way id="826473129" visible="true" version="3" changeset="129789472" timestamp="2022-12-06T16:23:14Z"
user="pmyteh" uid="81653">
  <nd ref="107366213"/>
  <nd ref="3468209232"/>
  <nd ref="3476406048"/>
  <nd ref="4112993338"/>
  <nd ref="2188828954"/>
  <nd ref="534390991"/>
  <nd ref="7839585527"/>
  <nd ref="348087295"/>
  <tag k="bicycle" v="yes"/>
  <tag k="foot" v="yes"/>
  <tag k="highway" v="tertiary"/>
  <tag k="horse" v="yes"/>
  <tag k="lanes" v="3"/>
  <tag k="lit" v="yes"/>
  <tag k="maxspeed" v="30 mph"/>
  <tag k="name" v="Brownlow Hill"/>
  <tag k="oneway" v="no"/>
  <tag k="postal_code" v="L3"/>
  <tag k="sidewalk" v="both"/>
  <tag k="surface" v="asphalt"/>
</way>
```

Dijkstra's algorithm (1956)

- Non-negative edge weights
- Visits nodes in increasing distance from *source*



```

proc Dijkstra1(V, E, δ, s)
  dist[v] = +∞ for all v ∈ V \ {s}
  dist[s] = 0
  Insert(Q, ⟨dist[s], s⟩)
  while Q ≠ ∅ do
    ⟨d, u⟩ = ExtractMin(Q)
    for (u, v) ∈ E ∩ ({u} × V) do
      if dist[u] + δ(u, v) < dist[v] then
        dist[v] = dist[u] + δ(u, v)
        if v ∈ Q then
          DecreaseKey(Q, v, dist[v])
        else
          Insert(Q, ⟨v, dist[v]⟩)
  return dist
  
```

relax

Fibonacci heaps
(Fredman, Tarjan 1984)
⇒ $O(m + n \cdot \log n)$

```

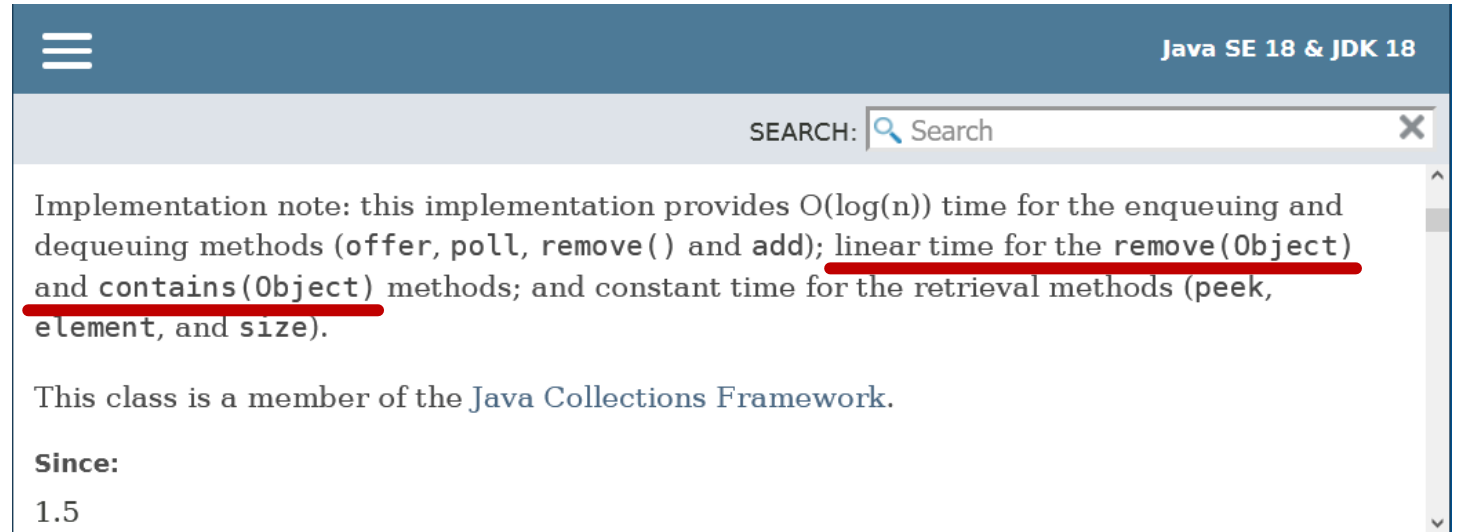
proc Dijkstra2(V, E, δ, s)
  dist[v] = +∞ for all v ∈ V \ {s}
  dist[s] = 0
  Insert(Q, ⟨dist[s], s⟩)
  while Q ≠ ∅ do
    ⟨d, u⟩ = ExtractMin(Q)
    for (u, v) ∈ E ∩ ({u} × V) do
      if dist[u] + δ(u, v) < dist[v] then
        dist[v] = dist[u] + δ(u, v)
        if v ∈ Q then
          Remove(Q, v)
          Insert(Q, ⟨dist[v], v⟩)
  return dist
  
```

Q = ⟨0,A⟩
~~⟨2,B⟩~~ ~~⟨4,C⟩~~
 relax
~~⟨3,C⟩~~ ~~⟨6,D⟩~~
~~⟨4,D⟩~~
~~⟨6,E⟩~~

$O(\log n)$ Remove
⇒ $O(m \cdot \log n)$

The Challenge - Java's builtin binary heap

- no decreasekey
- remove $O(n)$ time
⇒ Dijkstra $O(m \cdot n)$
- comparator function



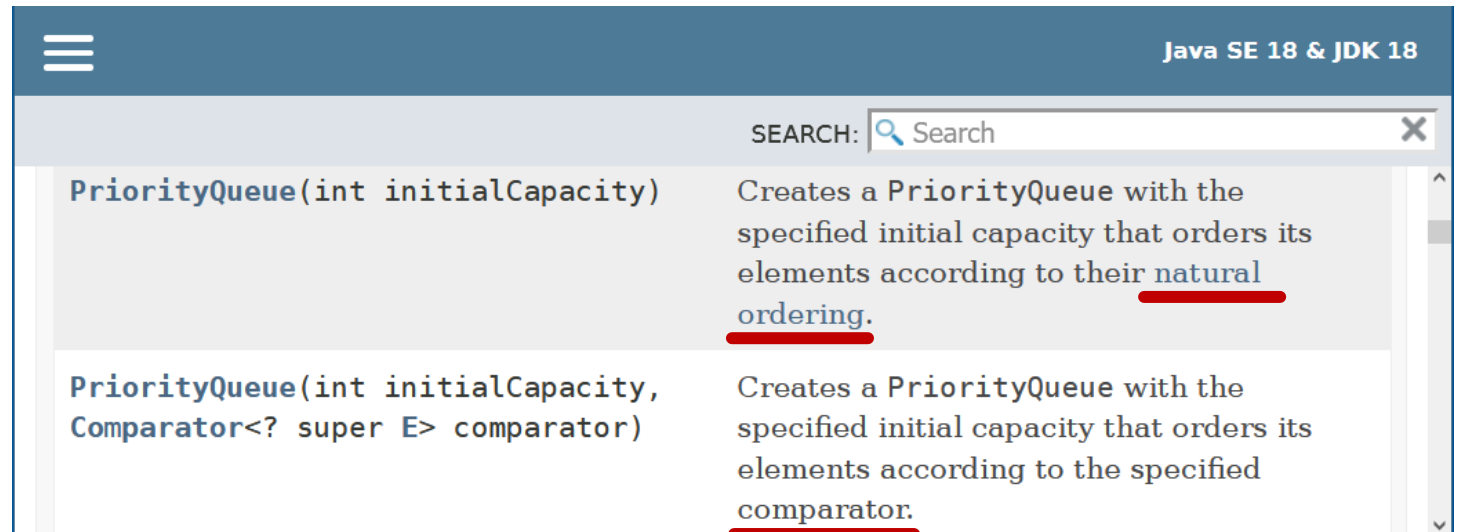
Java SE 18 & JDK 18

SEARCH:

Implementation note: this implementation provides $O(\log(n))$ time for the enqueueing and dequeuing methods (`offer`, `poll`, `remove()` and `add`); linear time for the `remove(Object)` and `contains(Object)` methods; and constant time for the retrieval methods (`peek`, `element`, and `size`).

This class is a member of the Java Collections Framework.

Since:
1.5




Java SE 18 & JDK 18

SEARCH:

<code>PriorityQueue(int initialCapacity)</code>	Creates a <code>PriorityQueue</code> with the specified initial capacity that orders its elements according to their <u>natural ordering</u> .
<code>PriorityQueue(int initialCapacity, Comparator<? super E> comparator)</code>	Creates a <code>PriorityQueue</code> with the specified initial capacity that orders its elements according to the specified <u>comparator</u> .

Repeated insertions

- **Relax** inserts new copies of item
- Skip **outdated** items



```
proc Dijkstra3(V, E, δ, s)
```

```
  dist[v] = +∞ for all v ∈ V \ {s}
```

```
  dist[s] = 0
```

```
  Insert(Q, ⟨dist[s], s⟩)
```

```
  while Q ≠ ∅ do
```

```
    ⟨d, u⟩ = ExtractMin(Q)
```

outdated? → if $d = \text{dist}[u]$ then

```
    for (u, v) ∈ E ∩ ({u} × V) do
```

```
      if  $\text{dist}[u] + \delta(u, v) < \text{dist}[v]$  then
```

```
         $\text{dist}[v] = \text{dist}[u] + \delta(u, v)$ 
```

relax
= reinsert → Insert(Q, ⟨dist[v], v⟩)

```
  return dist
```

Q = ⟨0, A⟩

~~⟨2, B⟩~~ ~~⟨4, C⟩~~

~~⟨3, C⟩~~ ~~⟨4, C⟩~~ ~~⟨6, D⟩~~

~~⟨4, C⟩~~ ~~⟨4, D⟩~~ ~~⟨6, D⟩~~

~~⟨4, D⟩~~ ~~⟨6, D⟩~~


~~⟨6, D⟩~~ ~~⟨6, E⟩~~

~~⟨6, E⟩~~

Using a visited set

```


proc Dijkstra4(V, E, δ, s)
  dist[v] = +∞ for all v ∈ V \ {s}
  dist[s] = 0
  visited = ∅
  Insert(Q, ⟨dist[s], s⟩)
  while Q ≠ ∅ do
    ⟨d, u⟩ = ExtractMin(Q)
    use bitvector → if u ∉ visited then
      visited = visited ∪ {u}
      for (u, v) ∈ E ∩ ({u} × V) do
        if dist[u] + δ(u, v) < dist[v] then
          dist[v] = dist[u] + δ(u, v)
          Insert(Q, ⟨dist[v], v⟩)
  return dist
  
```



A shaky idea...

```
proc Dijkstra4(V, E, δ, s)
  dist[v] = +∞ for all v ∈ V \ {s}
  dist[s] = 0
  visited = ∅
  Insert(Q, ⟨dist[s], s⟩)
  while Q ≠ ∅ do
    d never used → ⟨d, u⟩ = ExtractMin(Q)
    if u ∉ visited then
      visited = visited ∪ {u}
      for (u, v) ∈ E ∩ ({u} × V) do
        if dist[u] + δ(u, v) < dist[v] then
          dist[v] = dist[u] + δ(u, v)
          Insert(Q, ⟨dist[v], v⟩)
  return dist
```

- Q only store nodes (save space)
- Comparator
- Key = **current** distance $dist$



Heap invariants break



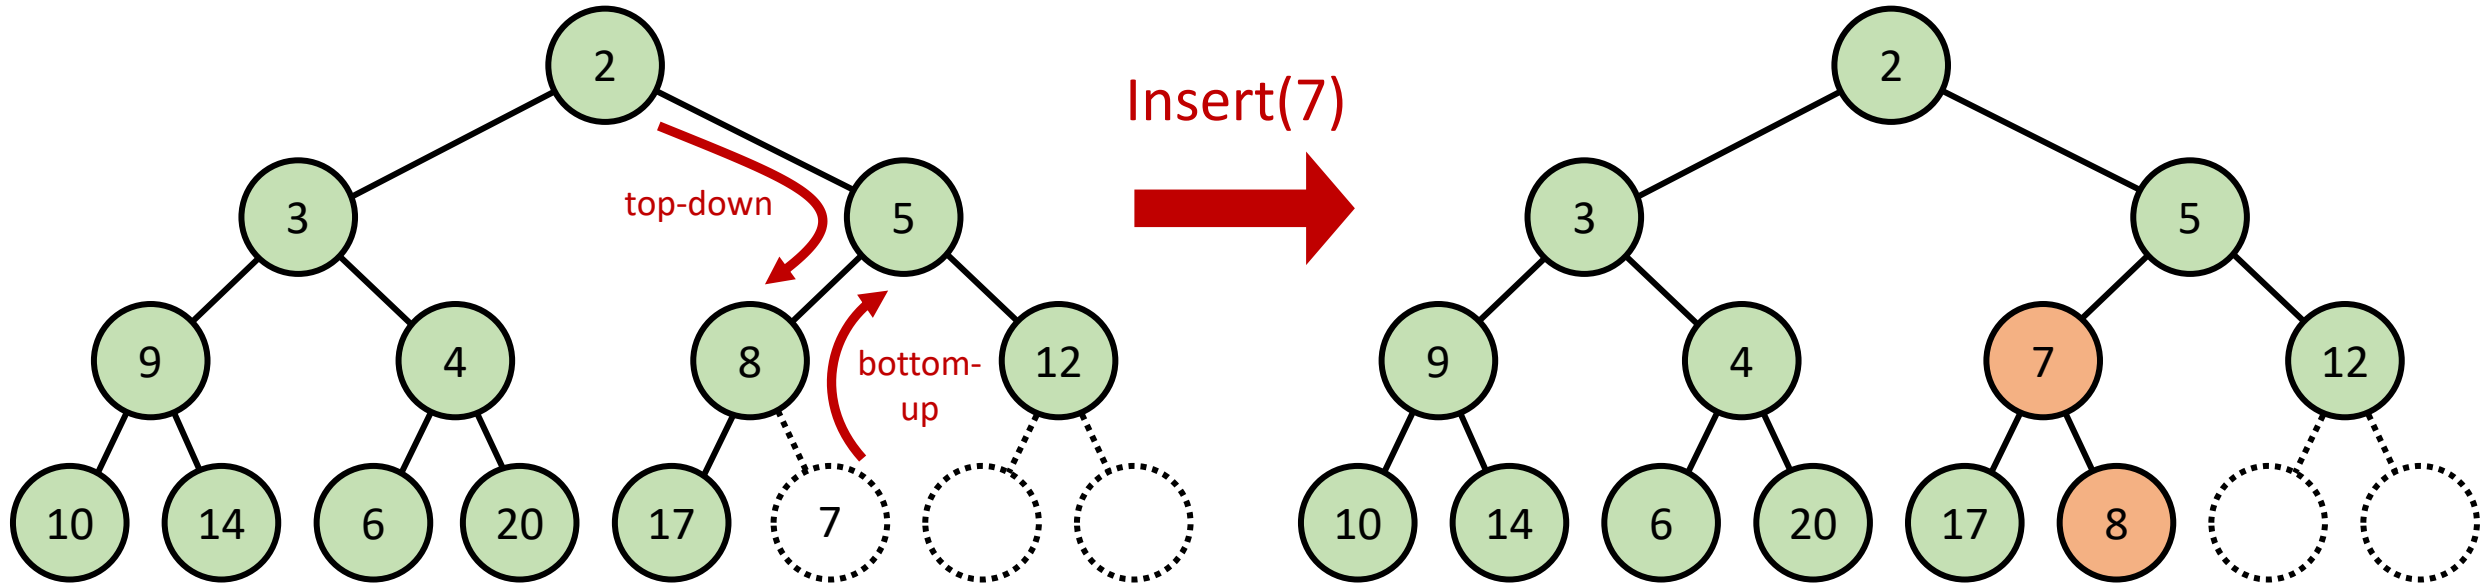
Experimental study

- Implemented Dijkstra₄ in Python
 - Stress test on random cliques
 - Binary heaps **failed** (default priority queue in Java and Python)
 - Skew heaps **worked**
 - Leftist heaps **worked**
 - Pairing heaps **worked**
 - Binomial queues **worked**
 - Post-order heaps **worked**
 - Binary heaps with top-down insertions **worked**
- } **Pointer based**
- } **Implicit (space efficient)**



```
visited = set()
Q = Queue()
Q.insert(Item(0, source))
while not Q.empty():
    u = Q.extract_min().value
    if u not in visited:
        visited.add(u)
        for v in G.out[u]:
            dist_v = dist[u] + G.weights[(u, v)]
            if dist_v < dist[v]:
                dist[v] = dist_v
                parent[v] = u
                Q.insert(Item(dist[v], v))
```

Binary heap insertions

– bottom-up vs top-down



Binary heaps using *dist* in a comparator fails



Definition

Priority Queues supporting Decreasing Keys

- Items = $\langle \text{key}, \text{value} \rangle$
- Over time keys can decrease – *priority queue is not informed*
- Items are compared w.r.t. their **current keys**
- The **original key** of an item is the key when it was inserted

Insert (item)

ExtractMin () returns an item with **current key less than or equal to all original keys** in the priority queue

Theorem 1

Dijkstra_4 correctly computes shortest paths when using *dist* as current key and a priority queue supporting decreasing keys




Theorem 2

The following priority queues support decreasing keys (out of the box)

- binary heaps with top-down insertions ((Williams 1964))
- leftist heaps (Crane 1972)
- binomial queues (Vuillemin, 1978)
- skew heaps (Sleator, Tarjan 1986)
- pairing heaps (Fredman, Sedgwick, Sleator, Tarjan 1986)
- post-order heaps (Harvey, Zatloukal 2004)

Proof of Theorem 2 - Basic idea


- **Decreased heap order**
 u ancestor of $v \Rightarrow$
current key $u \leq$ original key v
 - Root valid item to extract
 - Top-down merging two paths
preserves decreased heap order
- \Rightarrow **skew heaps** and **leftist heaps**
support decreasing keys



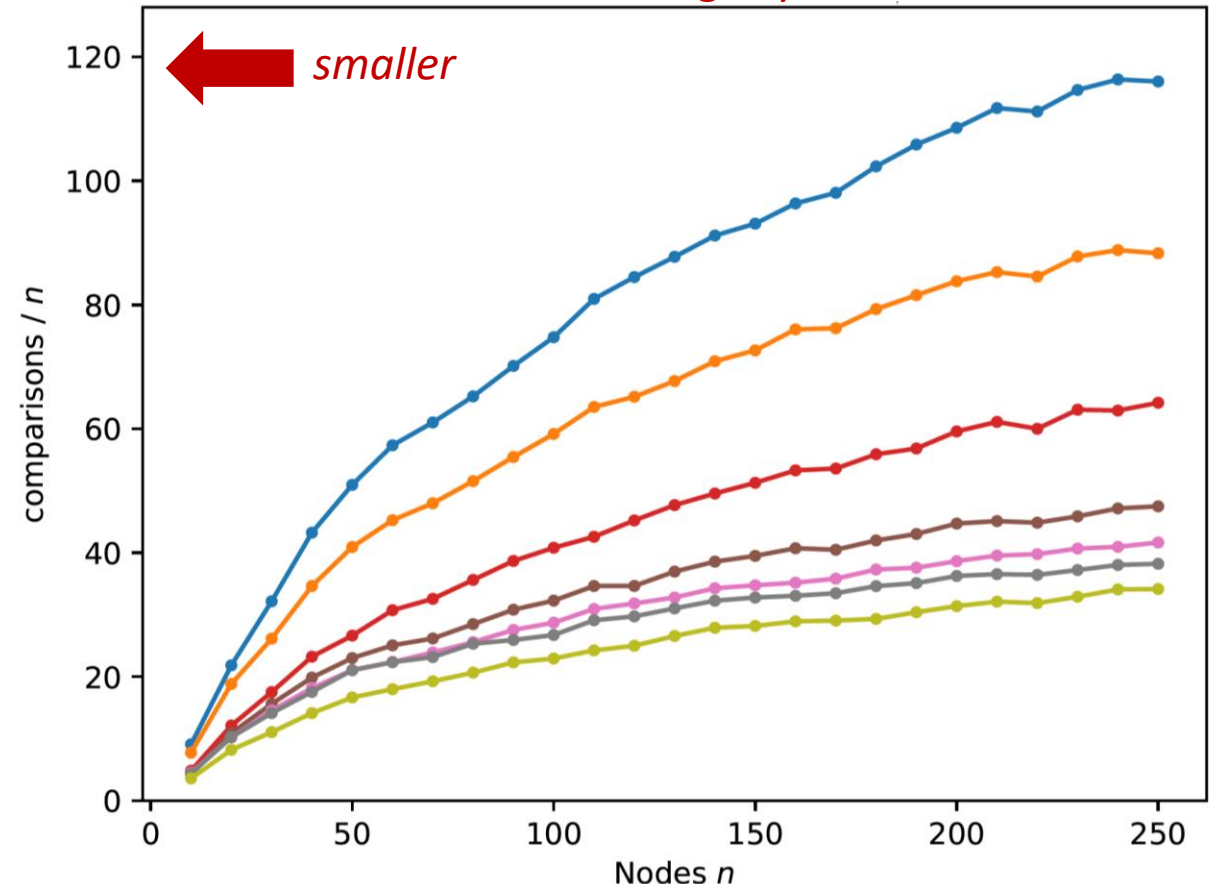
Experimental evaluation of various heaps

- Cliques with uniform random weights
- With decreasing keys less comparisons (outdated items removed earlier)

⟨key, value⟩ pairs




decreasing keys




Reduction in comparisons

comparisons decreasing keys / comparisons \langle key, value \rangle pairs



Postorder heap [Harvey and Zatloukal, FUN 2004]



- Insert amortized $O(1)$, ExtractMin amortized $O(\log n)$
- Implicit (space efficient)
- Best implicit comparison performance (and good time performance)

Conclusion

- Introduced notion of **priority queues with decreasing keys**
... as an approach to deal with outdated items in Dijkstra's algorithm
- Experiments identified priority queues supporting decreasing keys
... just had to prove it
- Builtin priority queues in Java and Python are binary heaps
... do not support decreasing keys
- **Binary heaps with top-down insertions** do support decreasing keys
... and also

**skew heaps, leftist heaps, pairing heaps,
binomial queues, post-order heaps**

