

Constant Time Priority Queues

Constant Time Priority Queues

Priority Queues on Parallel Machines

Gerth S. Brodal

BRICS

Aarhus University

The Problem

Sequential Priority Queues



Parallel Machines

problem: How can parallelism be exploited when performing single priority queue operations?

Operations

The common sequential priority queue operations

MAKE QUEUE

- creates an empty priority queue

INSERT (Q, e)

MELD (Q_1, Q_2)

FIND MIN (Q)

EXTRACT MIN (Q)

- deletes the minimum element from Q .

DELETE (Q, e)

DECREASE KEY (Q, e, e')

- replaces e by a smaller element

BUILD (e_1, \dots, e_n)

- creates a priority queue containing e_1, \dots, e_n

MULTI INSERT (Q, e_1, \dots, e_k)

Assume that a pointer to e is given

Model

- Comparison model
- Parallel machines: EREW or CREW PRAM
- n denotes the maximum size of a priority queue
- k number of elements involved in MULTIINSERT

Previous & New Results

Model	Pinotti, Pucci '92	Pinotti, Pucci '91	Chen, Hu '94	Ranade, et al. '94	This Talk
	EREW	CREW	EREW	Array	CREW
FINDMIN	1	1	1	1	1
INSERT	$\log \log n$	-	-	1	1
EXTRACTMIN	$\log \log n$	-	-	1	1
MELD	-	$\log \frac{n}{k} + \log \log k$	$\log \log \frac{n}{k} + \log k$	-	1
DELETE	-	-	-	-	1
DECREASEKEY	-	-	-	-	1
BUILD	$\log n$	$\frac{n}{k} \log k$	$\log \frac{n}{k} \log k$	-	$\log n$
MULTIINSERT	-	$\log \frac{n}{k} + \log k$	$\log \log \frac{n}{k} + \log k$	-	$\log k$
MULTIDELETE	-	$\log \frac{n}{k} + \log \log k$	$\log \log \frac{n}{k} + \log k$	-	-

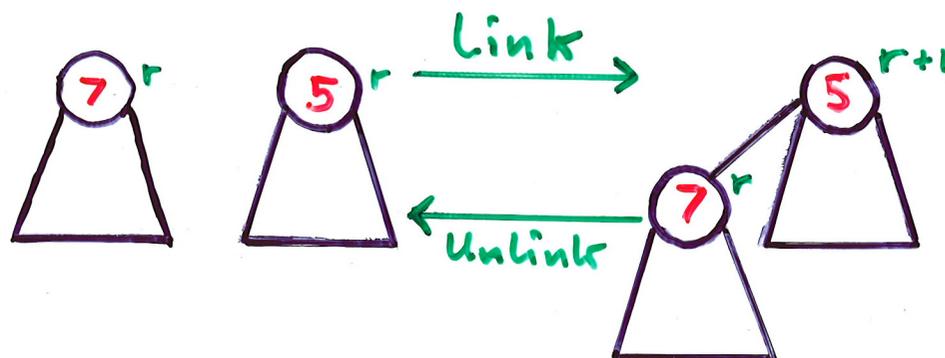
Delete the k smallest elements
 Assume that k is fixed
 Only one priority queue

A Simple Priority Queue

Basic Idea: Represent a priority queue by a forest of heap ordered binomial trees

- A rank 0 tree is a single node
- A rank r tree is obtained from two trees of rank $r-1$.

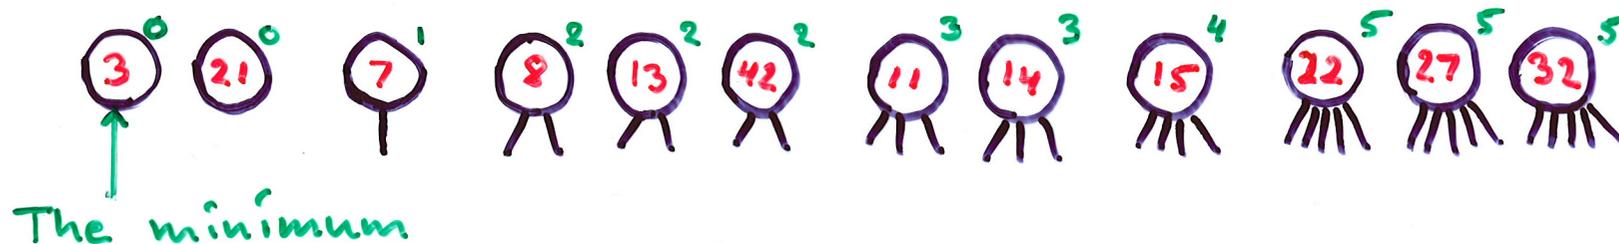
Primitive operations:



A Simple Priority Queue

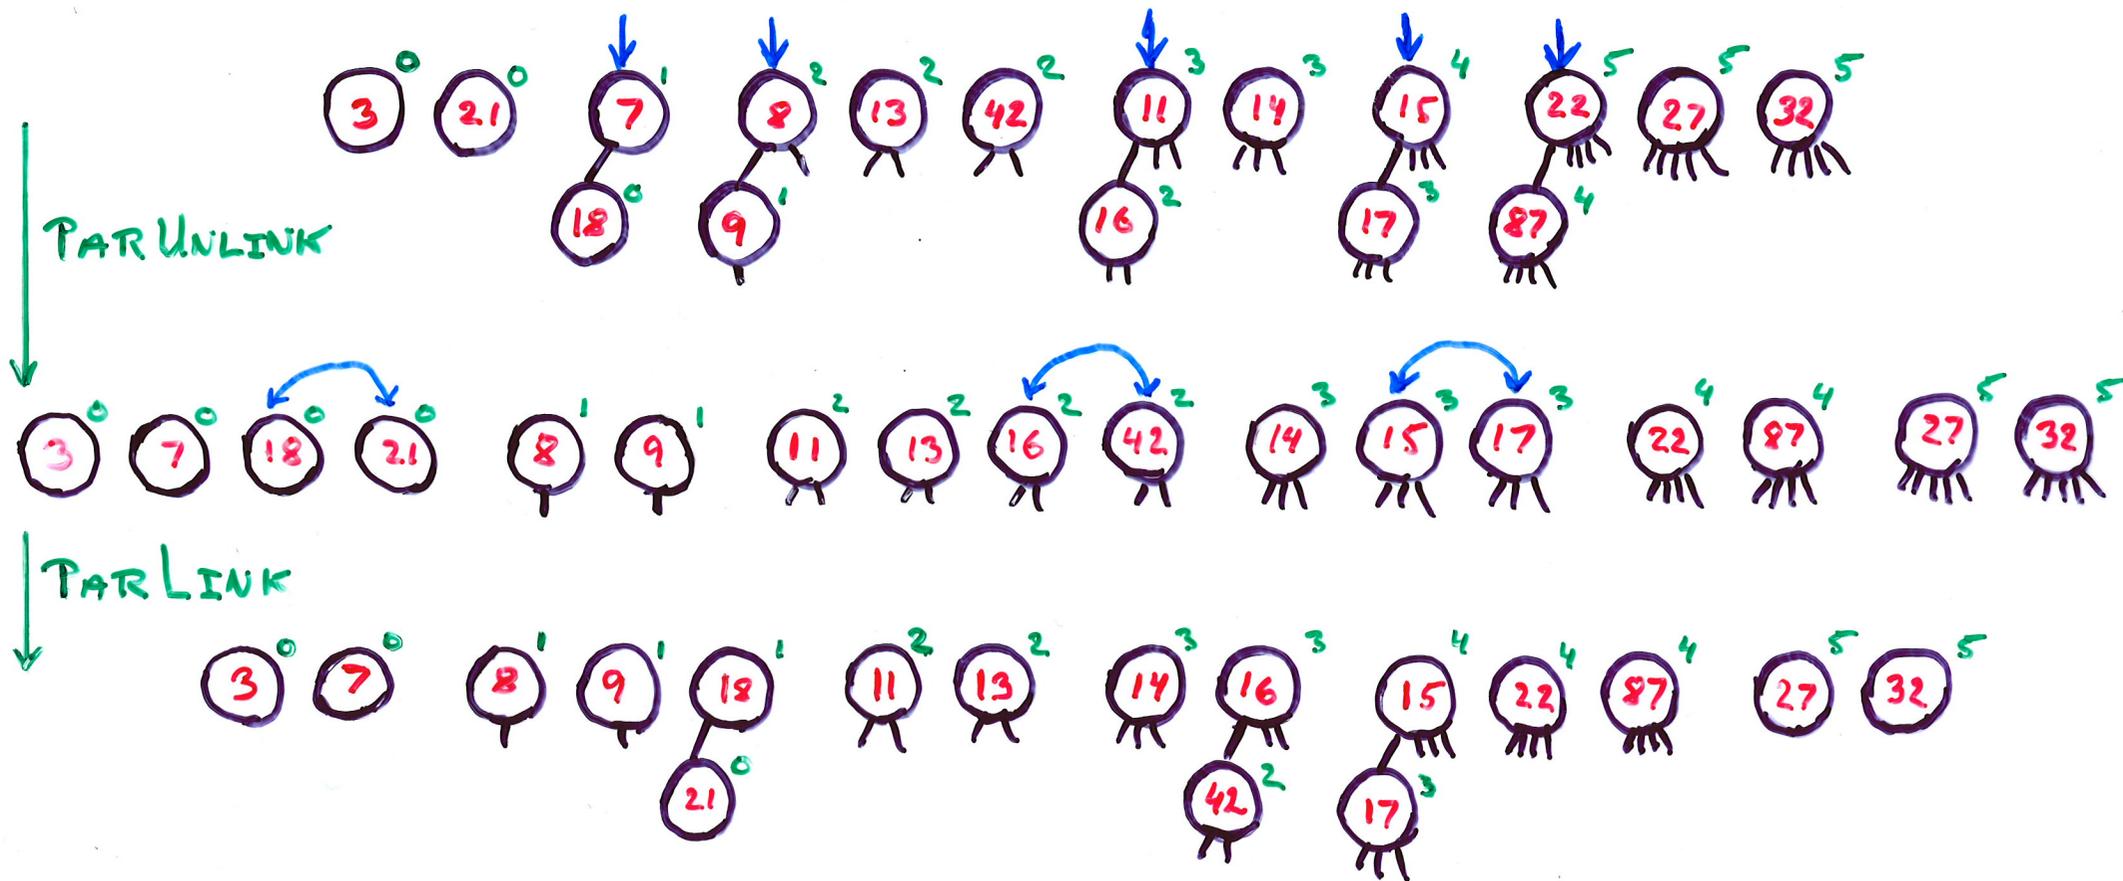
- The Invariant:
- There are 1, 2 or 3 trees of each rank
 - The minimum root of rank i is smaller than all roots of rank $\geq i$.

Example:



$$3 \leq 7 \leq 8 \leq 11 \leq 15 \leq 22$$

Parallel Linking and Unlinking



Notice: PAR UNLINK followed by PAR LINK is not the identity!

The Priority Queue Operations

FINDMIN (Q): return $\min(Q.L[0])$

INSERT (Q, e): $Q.L[0] := Q.L[0] \cup \{e\}$
PARLINK (Q)

MELD (Q₁, Q₂): for $p := 0$ to $\log n$ par do $Q_1.L[p] := Q_1.L[p] \cup Q_2.L[p]$
do 3 times PARLINK (Q₁)

EXTRACT MIN (Q): $e := \min(Q.L[0])$
 $Q.L[0] := Q.L[0] \setminus \{e\}$
| PARUNLINK (Q)
| PARLINK (Q)
return e

The above operations can all be implemented in
constant time with $\log n$ processors on a CREW PRAM

Building a Priority Queue

Step I: Build a normal Binomial Queue

Time $O(\log n)$ with $O\left(\frac{n}{\log n}\right)$ processors
on a EREW PRAM

Step II: for $i := 1$ to $\log n$ do
PARUNLINK,
PARLINK.

Time $O(\log n)$ with $O(\log n)$ processors
on a EREW PRAM

Building a Priority Queue

Step I: Build a normal Binomial Queue

Time $O(\log n)$ with $O\left(\frac{n}{\log n}\right)$ processors
on a EREW PRAM

Step II: for $i := 1$ to $\log n$ do
PARUNLINK,
PARLINK.

Time $O(\log n)$ with $O(\log n)$ processors
on a EREW PRAM

MULTI INSERT

MULTI INSERT $(Q, e_1, \dots, e_k) \equiv$ MELD $(Q, \text{BUILD}(e_1, \dots, e_k))$

Time $O(\log k)$ with $O\left(\frac{\log n + k}{\log k}\right)$ processors
on a CREW PRAM

Summary

(simple priority queues)

Model	CREW PRAM*	EREW PRAM
FINDMIN	1	1
INSERT, EXTRACTMIN, MELD	1	$\log \log n$
BUILD	$\log n$	$\log n$
MULTIINSERT	$\log k$	$\log k + \log \log n$

*The bounds also hold for the EREW PRAM, provided the processors know which priority queues are involved.

Summary

(simple priority queues)

Model	CREW PRAM*	EREW PRAM
FINDMIN	1	1
INSERT, EXTRACTMIN, MERGE	1	$\log \log n$
BUILD	$\log n$	$\log n$
MULTIINSERT	$\log k$	$\log k + \log \log n$

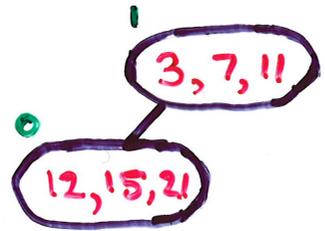
*The bounds also hold for the EREW PRAM, provided the processors know which priority queues are involved.

MULTI DELETE

MULTI DELETE was first considered by Pinotti, Pucci 91.

- Basic idea:
- Fix k
 - Store k elements (in sorted order) in each node (instead of only one)
 - All elements in a node are smaller than the elements at the sons.

Example: $k=3$.

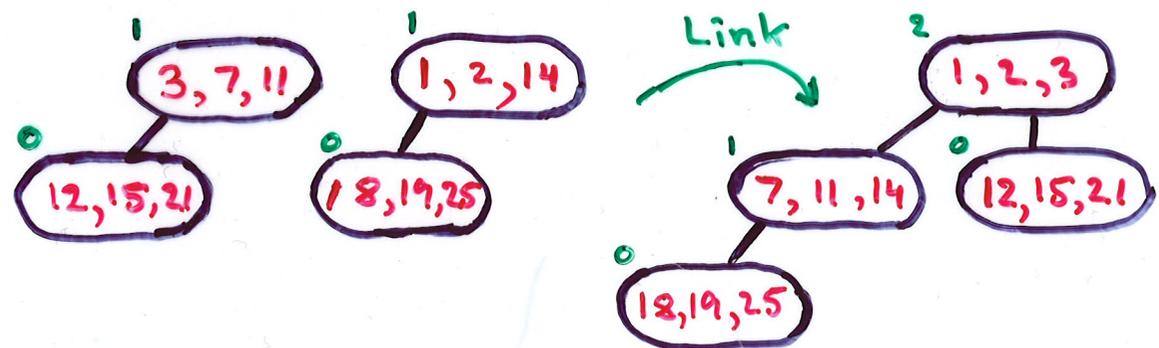


MULTI DELETE

MULTI DELETE was first considered by Pinotti, Pucci 91.

- Basic idea:
- Fix k
 - Store k elements (in sorted order) in each node (instead of only one)
 - All elements in a node are smaller than the elements at the sons.

Example: $k=3$.



The merging of the two root sets can be done in time $O(\log \log k)$ on a CREW PRAM
 — Kruskal 83

Summary (MULTI DELETE)

- Model: CREW PRAM
- MULTI INSERT : $O(T_{\text{SORT}}(k)) = O(\log k)$ - Cole 88
- MULTI DELETE, MERGE: $O(T_{\text{MERGE}}(k)) = O(\log \log k)$ - Kruskal 83
- BUILD : $O(T_{\text{SORT}}(k) + \log \frac{n}{k} T_{\text{MERGE}}(k)) = O(\log k + \log \frac{n}{k} \log \log k)$

DELETE

It is not obvious how the simple priority queues can support DELETE operations in constant time!

Idea: Combine ideas from

- Fibonacci Heaps,
- Relaxed Heaps

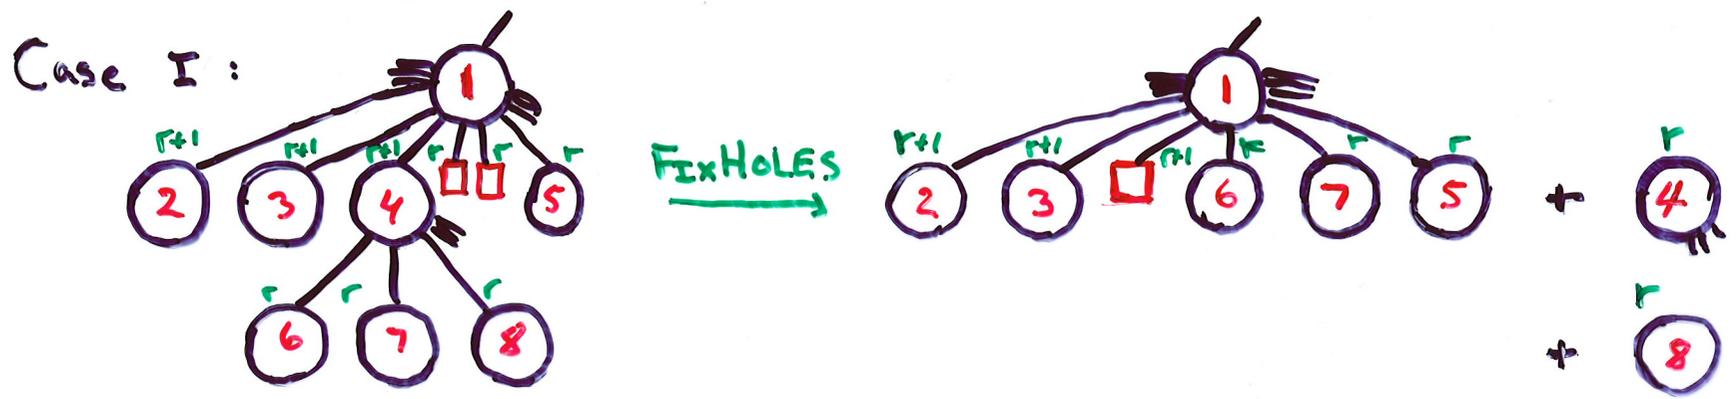
Fredman, Tarjan 84
Driscoll et al. 88

- Basic ideas:
- A node of rank r has 3 sons of each rank $0, \dots, r-1$
 - Subtrees can be missing - holes
 - At most two holes have equal rank in a priority queue

Reducing The Number of Holes

Two holes of rank r can be replaced by a rank $r+1$ hole

Assume w.l.o.g. the two holes are brothers:



- Performing the transformations in parallel for each possible $r \Rightarrow$ #holes of each rank is ≤ 2 .

The DELETE operation

Step I: Cut off the subtree at the node to be deleted
- Creates a hole.

Step II: Perform once parallel FixHoles.
- Bounds the number of holes by two of each rank.

Step III: Perform PARLINK $O(1)$ times
- Bounds the number of roots of each rank by $O(1)$.

- DELETE can be performed in constant time with $O(\log n)$ processors on a CREW PRAM

Summary (simple priority queues)

Model	CREW PRAM*	EREW PRAM
FINDMIN	1	1
INSERT, EXTRACTMIN, MELD	1	$\log \log n$
BUILD	$\log n$	$\log n$
MULTIINSERT	$\log k$	$\log k$
DELETE, DECREASE KEY	1*	$\log \log n$ + $\log \log n$ $\log \log n$

* The bounds also hold for the EREW PRAM, provided the processors know which priority queues are involved.

* On a EREW this bound becomes $\log \log n$

Conclusion

- Constant time priority queues
- Without DELETE the data structure is simple

Open Problems

- MULTI DELETE faster / less work?
- — for non fixed k ?
- DECREASE KEY (and MELD, INSERT) with $O(1)$ work?
- Applications...
- Can the ideas related to DELETE be used to simplify Relaxed Heaps?