## Polylogarithmic Fully Retroactive Priority Queues via Hierarchical Checkpointing

Erik D. Demaine, Tim Kaler, Quanquan Liu, Aaron Sidford, Adam Yedidia





• Abstract data type where each element given a priority key.

- Abstract data type where each element given a priority key.
- Priority queue supports:
  - Updates:

- Abstract data type where each element given a priority key.
- Priority queue supports:
  - Updates:
    - Insert: Inserts elements.

- Abstract data type where each element given a priority key.
- Priority queue supports:
  - Updates:
    - Insert: Inserts elements.
    - Delete-min: Deletes the element with minimum key.

- Abstract data type where each element given a priority key.
- Priority queue supports:
  - Updates:
    - Insert: Inserts elements.
    - Delete-min: Deletes the element with minimum key.
  - Query:
    - Find-min: Finds the minimum key value element.

- Abstract data type where each element given a priority key.
- Priority queue supports:
  - Updates:
    - Insert: Inserts elements.
    - Delete-min: Deletes the element with minimum key.
  - Query:
    - Find-min: Finds the minimum key value element.



Normal priority queue events occur in the present.



 Normal priority queue events occur in the present.



- Normal priority queue events occur in the present.
- Timeline of update events maintains history of updates.



- Normal priority queue events occur in the present.
- Timeline of update events maintains history of updates.
- Updates occur in the present and are appended to the end of the timeline.



- Normal priority queue events occur in the present.
- Timeline of update events maintains history of updates.
- Updates occur in the present and are appended to the end of the timeline.



- Normal priority queue events occur in the present.
- Timeline of update events maintains history of updates.
- Updates occur in the present and are appended to the end of the timeline.
- Queries can be made in the present.



 Would like to insert(6) at the beginning of the timeline.



- Would like to insert(6) at the beginning of the timeline.
- Cannot insert(6) right now because do not know intervening updates.



- Would like to insert(6) at the beginning of the timeline.
- Cannot insert(6) right now because do not know intervening updates.
- In a normal priority queue, can replay all intervening updates.



- Would like to insert(6) at the beginning of the timeline.
- Cannot insert(6) right now because do not know intervening updates.
- In a normal priority queue, can replay all intervening updates.
- Costly if large number of updates.



- Update in the present are appended to the end of the timeline:
  - insert
  - delete-min
- Query the state of the structure in the present:
  - find-min





[DIL03]



### Partially Retroactive Priority Queue: Update



time

[DIL03]

Insert-Op(0, insert(6))

## Partially Retroactive Priority Queue: Query



time

min = 7

Demaine, Kaler, Liu, Sidford, Yedidia

[DIL03]





 $O(\log m) \text{ runtime} \qquad \qquad \text{time} \qquad \qquad \text{[DIL03]}$ 

m = number of updates in timeline.





Demaine, Kaler, Liu, Sidford, Yedidia

- Partial retroactivity: updates can be added or deleted from anywhere in timeline but queries can only be made in the present.
  [DIL03]
- Full retroactivity: queries can be made at any point in the timeline. [DIL03]
- Not always easy converting from partial to full retroactivity

	Normal	Partially Retroactive	Fully Retroactive
Updates	Present		
Queries	Present		

	Normal	Partially Retroactive	Fully Retroactive
Updates	Present	Past and Present	
Queries	Present	Present	

	Normal	Partially Retroactive	• Fully Retroactive
Updates	Present	Past and Present	Past and Present
Queries	Present	Present	Past and Present

	Normal	Partially Retroactive	• Fully Retroactive
Updates	Present	Past and Present	Past and Present
Queries	Present	Present	Past and Present

Can we minimize the cost of converting from partial to full retroactivity for retroactive data structures?

### **Previous Results**

#### [DIL03, GK09]:

Abstract Data Type	Partially Retroactive	Fully Retroactive
Dictionary	$O(\log m)$	$O(\log m)$
Queue	<i>O</i> (1)	$O(\log m)$
Stack	$O(\log m)$	$O(\log m)$
Deque	$O(\log m)$	$O(\log m)$
Union-Find	$O(\log m)$	$O(\log m)$
Priority queue	$O(\log m)$	$O(\sqrt{m}\log m)$

### **Previous Results**

#### [DIL03, GK09]:

Abstract Data Type	Partially Retroactive	Fully Retroactive
Dictionary	$O(\log m)$	$O(\log m)$
Queue	<i>O</i> (1)	$O(\log m)$
Stack	$O(\log m)$	$O(\log m)$
Deque	$O(\log m)$	$O(\log m)$
Union-Find	$O(\log m)$	$O(\log m)$
Priority queue	$O(\log m)$	$O(\sqrt{m}\log m)$

 $O(\sqrt{m})$  overhead in priority queues obtained through checkpointing Demaine, Kaler, Liu, Sidford, Yedidia

### **Our Results**

#### [DIL03, GK09, DKLSY15]:

Abstract Data Type	Partially Retroactive	Fully Retroactive
Dictionary	$O(\log m)$	$O(\log m)$
Queue	<i>O</i> (1)	$O(\log m)$
Stack	$O(\log m)$	$O(\log m)$
Deque	$O(\log m)$	$O(\log m)$
Union-Find	$O(\log m)$	$O(\log m)$
Priority queue	$O(\log m)$	$\tilde{O}(\log^2 m)$

 $\tilde{O}(\log m)$  overhead in priority queues obtained through hierarchical checkpointing Demaine, Kaler, Liu, Sidford, Yedidia



## Previous Results Checkpointing: Query



## Previous Results Checkpointing: Query


















Find-Min(8) emaine, Kaler, Liu, Sidford, Yedidia

• May need to update  $\sqrt{m}$  partially retroactive priority queues.



Insert-Op(3, U) emaine, Kaler, Liu, Sidford, Yedidia









Insert-Op(3, U)



Insert-Op(3, U)





Get-View(8)



# **Hierarchical Checkpointing**





# **Hierarchical Checkpointing**



#### time

Inner nodes contains partially retroactive priority queues spanning segment of time.

# **Hierarchical Checkpointing**



Each node contains updates in leaves of subtree.





### Delete-Op(u<sub>4</sub>)



#### Delete-Op(u<sub>4</sub>)



### Delete-Op(u<sub>4</sub>)



### Delete-Op(u<sub>4</sub>)



#### Delete-Op(u<sub>4</sub>)



update  $O(\log m)$  partially retroactive priority queue ancestors



 $O(\log^2 m)$  overhead with extra factor from tree rebuild



 $O(\log^3 m)$  update time in fully retroactive priority queue





Merge disjoint, contiguous intervals of time.



Merge disjoint, contiguous intervals of time.



 $O(\log m)$  partially retroactive priority queues to merge



# Checkpointing vs. Hierarchical Checkpointing

1. Checkpointing is a general transformation that can be applied to any partially retroactive data structure.

# Checkpointing vs. Hierarchical Checkpointing

- 1. Checkpointing is a general transformation that can be applied to any partially retroactive data structure.
- 2. Hierarchical checkpointing has a stronger bound on conversion with a stronger assumption on partially retroactive data structure.
# Checkpointing vs. Hierarchical Checkpointing

- 1. Checkpointing is a general transformation that can be applied to any partially retroactive data structure.
- 2. Hierarchical checkpointing has a stronger bound on conversion with a stronger assumption on partially retroactive data structure.
- 3. Hierarchical checkpointing requires ability to merge two versions of a structure to produce read-only structure that can be queried.

#### **Hierarchical Checkpointing**



Insert-Op(6, U<sub>6</sub>)



Insert-Op(6, U<sub>6</sub>)

#### **Hierarchical Checkpointing**



Get-View(8)

#### **Hierarchical Checkpointing**



Get-View(8)

- Q<sub>now</sub> contains all elements that are in a partially retroactive priority queue
- Q<sub>del</sub> contains set of elements that were deleted from a partially retroactive priority queue by delete-min operations
- We maintain Q<sub>now</sub> and Q<sub>del</sub> as weightbalanced B-trees with balance factor d = 8





time



time



time



time



time

#### All elements in $Q_{1, del}$ deleted.



time

#### All elements in $Q_{2, now}$ survive.





1. Find element in  $Q_{1, now} \cup Q_{2, del}$ where  $D = |Q_{2, del}|$  are less than it.



1. Find element in  $Q_{1, now} \cup Q_{2, del}$ where  $D = |Q_{2, del}|$  are less than it.



2. Use the element to split each tree into two parts containing surviving and deleted elements.

- 1. Find element in  $Q_{1, now} \cup Q_{2, del}$ where  $D = |Q_{2, del}|$  are less than it.
- 2. Use the element to split each tree into two parts containing surviving and deleted elements.



- 1. Find element in  $Q_{1, now} \cup Q_{2, del}$ where  $D = |Q_{2, del}|$  are less than it.
- 2. Use the element to split each tree into two parts containing surviving and deleted elements.
- 3. List of tree pieces comprise merged structure.



4. If two pieces originate from the same tree, concatenate them into one tree.

$$Q_{3, \text{ now}} = \left[ \bigtriangleup, \swarrow, \land, \land, \land, \land, \land, \land, \land \right]$$

4. If two pieces originate from the same tree, concatenate them into one tree.

$$Q_{3, \text{ now}} = \begin{bmatrix} & & & & & & & \\ & & & & & & & \\ Q_{3, \text{ now}} = \begin{bmatrix} & & & & & & & & \\ & & & & & & & & \\ \end{bmatrix}$$

- 4. If two pieces originate from the same tree, concatenate them into one tree.
- 5. Query the trees in merged Q<sub>now</sub> to find min element.

$$Q_{3, \text{ now}} = \begin{bmatrix} \checkmark, \land, \land, \land, \land, \land, \land, \land, \land &, \land \end{bmatrix}$$
$$Q_{3, \text{ now}} = \begin{bmatrix} \checkmark, \land, \land, \land, \land, \land, \land, \land &, \land &, \land \end{bmatrix}$$

1. Find element in  $Q_{1, now} \cup Q_{2, del}$  where  $D = |Q_{2, del}|$  are less than it.

- 1. Find element in  $Q_{1, now} \cup Q_{2, del}$  where  $D = |Q_{2, del}|$  are less than it.
  - Adaption of finding the (D+1)-st smallest element in k arrays [RA10] to k weight balanced binary search trees.

- 1. Find element in  $Q_{1, now} \cup Q_{2, del}$  where  $D = |Q_{2, del}|$  are less than it.
  - Adaption of finding the (D+1)-st smallest element in k arrays [RA10] to k weight balanced binary search trees.
  - $O(k \log m)$  time where  $k = 2 \log m$ .

- 1. Find element in  $Q_{1, now} \cup Q_{2, del}$  where  $D = |Q_{2, del}|$  are less than it.
  - Adaption of finding the (D+1)-st smallest element in k arrays [RA10] to k weight balanced binary search trees.
  - $O(k \log m)$  time where  $k = 2 \log m$ .
- 2. Use the element to split each tree into two parts containing surviving and deleted elements.

- 1. Find element in  $Q_{1, now} \cup Q_{2, del}$  where  $D = |Q_{2, del}|$  are less than it.
  - Adaption of finding the (D+1)-st smallest element in k arrays [RA10] to k weight balanced binary search trees.
  - $O(k \log m)$  time where  $k = 2 \log m$ .
- 2. Use the element to split each tree into two parts containing surviving and deleted elements.
  - $O(\log^2 m)$  assuming list of tree pieces size  $O(\log m)$ .

- 1. Find element in  $Q_{1, now} \cup Q_{2, del}$  where  $D = |Q_{2, del}|$  are less than it.
  - Adaption of finding the (D+1)-st smallest element in k arrays [RA10] to k weight balanced binary search trees.
  - $O(k \log m)$  ime where  $k = 2 \log m$
- 2. Use the element to split each tree into two parts containing surviving and deleted elements.
  - $O(\log^2 m)$  assuming list of tree pieces size  $O(\log m)$ .
- 3. List of tree pieces comprise merged structure.

4. If two pieces originate from the same tree, concatenate them into one tree.

- 4. If two pieces originate from the same tree, concatenate them into one tree.
  - If two pieces in the same list are from the same original tree, then they span disjoint, contiguous intervals.

- 4. If two pieces originate from the same tree, concatenate them into one tree.
  - If two pieces in the same list are from the same original tree, then they span disjoint, contiguous intervals.
  - $O(\log m \log \log m)$  time to sort and  $O(\log m)$  time to concatenate.

- 4. If two pieces originate from the same tree, concatenate them into one tree.
  - If two pieces in the same list are from the same original tree, then they span disjoint, contiguous intervals.
  - $O(\log m \log \log m)$  time to sort and  $O(\log m)$  time to concatenate.
- 5. Query the trees in merged  $Q_{now}$  to find min element.

- 4. If two pieces originate from the same tree, concatenate them into one tree.
  - If two pieces in the same list are from the same original tree, then they span disjoint, contiguous intervals.
  - $O(\log m \log \log m)$  time to sort and  $O(\log m)$  time to concatenate.
- 5. Query the trees in merged  $Q_{now}$  to find min element.
  - $O(\log^2 m)$  time

- 4. If two pieces originate from the same tree, concatenate them into one tree.
  - If two pieces in the same list are from the same original tree, then they span disjoint, contiguous intervals.
  - $O(\log m \log \log m)$  time to sort and  $O(\log m)$  time to concatenate.
- 5. Query the trees in merged  $Q_{now}$  to find min element.
  - $O(\log^2 m)$  time

 $O(\log^2 m)$  time per merge

- 4. If two pieces originate from the same tree, concatenate them into one tree.
  - If two pieces in the same list are from the same original tree, then they span disjoint, contiguous intervals.
  - $O(\log m \log \log m)$  time to sort and  $O(\log m)$  time to concatenate.
- 5. Query the trees in merged  $Q_{now}$  to find min element.
  - $O(\log^2 m)$  time

 $O(\log^2 m)$  time per merge  $O(\log m)$  merges

- 4. If two pieces originate from the same tree, concatenate them into one tree.
  - If two pieces in the same list are from the same original tree, then they span disjoint, contiguous intervals.
  - $O(\log m \log \log m)$  time to sort and  $O(\log m)$  time to concatenate.
- 5. Query the trees in merged  $Q_{now}$  to find min element.
  - $O(\log^2 m)$  time

 $O(\log^3 m)$  runtime

# **Splitting Trees**

- Merge two partially retroactive priority queues by splitting the Q<sub>now</sub> and Q<sub>del</sub> trees to obtain surviving and deleted elements.
- Obtain a set of tree pieces.
- List of tree pieces comprise merged  $Q_{\text{now}}$  and  $Q_{\text{del.}}$
- Query each of the trees in the merged  $Q_{now}$  to obtain the minimum element.
# Merging PRPQs Example

Find  $|Q_{1, del}|$  smallest elements to delete from  $Q_{1, now} U Q_{2, del}$ 



# Merging PRPQs Example

Find  $|Q_{1, del}|$  smallest elements to delete from  $Q_{1, now} U Q_{2, del}$ 

# Merging PRPQs Example

Find  $|Q_{1, del}|$  smallest elements to delete from  $Q_{1, now} U Q_{2, del}$ 

The merged partially retroactive priority queue contains list of tree pieces.

# **Optimizing the Merge**

- Query for the minimum element in list of trees.
- To prevent the list of trees from exceeding O(log m) trees, we can concatenate trees that originated from the same queue.
- Merging the  $O(\log m)$  trees with merge time  $T(k) = k \log m$  where  $k \le \log m$ results in a merge time of  $O(\log^3 m)$ .

## Hierarchical Checkpointing for FRPQs

 By applying our checkpointing framework directly to fully retroactive priority queues, we obtain

Operation	Runtime
Updates	$O(\log^3 m) \longrightarrow O(\log^2 m)$
Queries	$O(\log^2 m \log \log m)$
Finding time of deletion	$O(\log^3 m \log \log m)$
	$O(\log^2 m)$

## Hierarchical Checkpointing for FRPQs

By applying our checkpointing framework directly to fully retroactive priority queues, we obtain

Operation	Hierarchical Checkpointing Runtime	Improved Runtimes
Updates	$O(\log^3 m)$	
Queries	$O(\log^3 m)$	
Finding time of deletion	$O(\log^4 m)$	

# Log-Shaving: Updates

Simpler rebuild procedure when a subtree is unbalanced leads to  $O(\log^2 m)$  updates.

Operation	Hierarchical Checkpointing Runtime	Improved Runtimes
Updates	$O(\log^3 m)$	$O(\log^2 m)$
Queries	$O(\log^3 m)$	
Finding time of deletion	$O(\log^4 m)$	

## **Log-Shaving: Queries**

Merging the updates using a binary merge tree results in  $O(\log^2 m \log \log m)$  queries.

Operation	Hierarchical Checkpointing Runtime	Improved Runtimes
Updates	$O(\log^3 m)$	$O(\log^2 m)$
Queries	$O(\log^3 m)$	$O(\log^2 m \log \log m)$
Finding time of deletion	$O(\log^4 m)$	

## Log-Shaving: Find-delete-time

Modified binary search that maintains a counter of current surviving elements results in  $O(\log^2 m)$  runtime for find-deletion-time.

Operation	Hierarchical Checkpointing Runtime	Improved Runtimes
Updates	$O(\log^3 m)$	$O(\log^2 m)$
Queries	$O(\log^3 m)$	$O(\log^2 m \log \log m)$
Finding time of deletion	$O(\log^4 m)$	$O(\log^2 m)$

## **Time-Fusibility**

- Time-fusible partially retroactive data structures are structures that allow merging of timelines:
  - Produces a read-only data structure that contains updates from both timelines.
  - Exhibits substring closure.

## **Time-Fusibility**

- Time-fusible partially retroactive data structures are structures that allow merging of timelines:
  - Produces a read-only data structure that contains updates from both timelines.
  - Exhibits substring closure.
- Time-fusible data structures allow an O(log<sup>2</sup> m) overhead converting from partial to full retroactivity via hierarchical checkpointing.

• Hierarchical checkpointing procedure for time-fusible structures  $O(\log^2 m)$  overhead.

- Hierarchical checkpointing procedure for time-fusible structures  $O(\log^2 m)$  overhead.
  - Examples of other structures?

- Hierarchical checkpointing procedure for time-fusible structures  $O(\log^2 m)$  overhead.
  - Examples of other structures?
- Polylogarithmic fully retroactive priority queue with runtimes:

Operation	Improved Runtimes
Updates	$O(\log^2 m)$
Queries	$O(\log^2 m \log \log m)$
Finding time of deletion	$O(\log^2 m)$

- Hierarchical checkpointing procedure for time-fusible structures  $O(\log^2 m)$  overhead.
  - Examples of other structures?
- Polylogarithmic fully retroactive priority queue with runtimes:

Operation	Improved Runtimes
Updates	$O(\log^2 m)$
Queries	$O(\log^2 m \log \log m)$
Finding time of deletion	$O(\log^2 m)$

– Logarithmic or constant overhead?

# **Open Questions**

- Does there exist logarithmic separation between partial and full retroactivity for priority queues?
- Can we adapt hierarchical checkpointing to a broader class of structures?

### **PRPQ** Fusion

- (Diagram describing finding the split-key procedure)
- We first find the split-key in the PRPQ using a modified median finding algorithm
- Suppose for simplicity  $|Q_{now}| = |Q_{del}|$
- (Diagram showing find split key with two priority queues, i.e. two sets of Q\_now and Q\_del)

# Time-Fusibility of PRPQs

- Two partially retroactive priority queues may be fused
- Split-key "splits" elements into elements that are deleted and elements that survive
- All elements less than split-key are deleted in the fused structure
- All elements greater than split-key survive
- (Words above may be said and replaced with a diagram.)

### **Fusion Runtime**

- The runtime of the fusion operation for partially retroactive priority queues is  $O(\log m \log \log m)$
- (Above could be placed with the diagram of the fusion in the previous slide.)

# Substring Integrity

- Deleted element is moved from Q<sub>now</sub> to Q<sub>del</sub>
- Delete-min operations that do not move an element from Q<sub>now</sub> to Q<sub>del</sub> can be represented by the insertion of a key of ∞ weight