#### Lock-free Contention Adapting Search Trees

#### Kjell Winblad, Konstantinos Sagonas and Bengt Jonsson



Department of Information Technology Uppsala University, Sweden



#### Introduction

- LFCA tree Related Work Evaluation
- Final Remarks

SPAA'2018

- A new lock-free data structure
  - The lock-free contention adapting search tree (LFCA tree)



#### Introduction

- LFCA tree
- Related Work
- Evaluation
- Final Remarks

SPAA'2018

- A new lock-free data structure
  - The lock-free contention adapting search tree (LFCA tree)
- Ordered sets and key-value stores



#### Introduction

- LFCA tree Related Work
- Evaluation
- Final Remarks

SPAA'2018

- A new lock-free data structure
  - The lock-free contention adapting search tree (LFCA tree)
- Ordered sets and key-value stores
- Insert and Remove (lock-free and linearizable)



#### Introduction

- LFCA tree Related Work
- Evaluation
- Final Remarks

SPAA'2018

- A new lock-free data structure
  - The lock-free contention adapting search tree (LFCA tree)
- Ordered sets and key-value stores
- Insert and Remove (lock-free and linearizable)
- Lookup (wait-free and linearizable)



#### Introduction

- LFCA tree Related Work
- Evaluation
- Final Remarks

SPAA'2018

- A new lock-free data structure
  - The lock-free contention adapting search tree (LFCA tree)
- Ordered sets and key-value stores
- Insert and Remove (lock-free and linearizable)
- Lookup (wait-free and linearizable)
- Range query (lock-free and linearizable)
  - atomic snapshot of all items inside a range



#### Introduction

- LFCA tree Related Work
- Evaluation
- Final Remarks

SPAA'2018

- A new lock-free data structure
  - The lock-free contention adapting search tree (LFCA tree)
- Ordered sets and key-value stores
- Insert and Remove (lock-free and linearizable)
- Lookup (wait-free and linearizable)
- Range query (lock-free and linearizable)
  - atomic snapshot of all items inside a range
- Adapts its synchronization granularity to the workload
  - Attempts to get the best of:
    - Coarse-grained synchronization
    - Fine-grained synchronization



#### Introduction

LFCA tree Related Work Evaluation

Final Remarks

#### Concurrent ordered sets with support for range queries

#### Important for:

- Big scale databases and data processing platforms
  - Fast updates to store incoming data
  - Concurrent range queries for analytics



#### Introduction

- LFCA tree Related Work Evaluation
- Final Remarks

#### Concurrent ordered sets with support for range queries

- Important for:
  - Big scale databases and data processing platforms
    - Fast updates to store incoming data
    - Concurrent range queries for analytics
- Challenging
  - Large range queries + updates = many collisions



#### Coarse-grained – Mutable reference to immutable data





Final Remarks



Im-Treap Herlihy [PPoPP'90] (general method)

#### Fine-grained – Immutable leaf nodes



k-ary [OPODIS'12]



# **Coarse-grained vs Fine-grained – Trade-offs**

Introduction LFCA tree Related Work Evaluation Final Remarks

SPAA'2018



Figure: 10% Inserts, 10% Removes, 55% Lookups, 25% Range queries,  $\approx$  500K items





#### LFCA tree

Related Work Evaluation Final Remarks

SPAA'2018





# LFCA tree (Insert 1)

n		<b>^</b>	а		~		<b>^</b>	n
	LI I	U		u	ີ	LI	U	

LFCA tree

Related Work Evaluation Final Remarks





SPAA'2018



SPAA'2018



Evaluation

SPAA'2018





SPAA'2018





Evaluation

SPAA'2018





Evaluation

SPAA'2018





Evaluation

SPAA'2018





Introduction LFCA tree

Evaluation

SPAA'2018





Introduction LFCA tree

Evaluation

SPAA'2018





Evaluation Final Remarks

SPAA'2018





Evaluation Final Remarks

# LFCA tree



SPAA'2018



Evaluation Final Remarks

SPAA'2018





Evaluation Final Remarks

SPAA'2018





Introduction LFCA tree

Evaluation

SPAA'2018





#### LFCA tree

Related Work Evaluation Final Remarks

#### Implementation Details

- Immutable treap (balanced binary search tree)
- Items stored in fat leaf nodes
  - Up to 64 items in arrays
  - Improves cache locality



# How do LFCA trees perform?

Introduction

LFCA tree

Related Work Evaluation Final Remarks

SPAA'2018



Figure: 10% Inserts, 10% Removes, 55% Lookups, 25% Range queries,  $\approx$  500K items



## **Evaluation**

Introduction

- LFCA tree
- **Related Work**
- Evaluation
- Final Remarks

Benchmark with a mix of

- Inserts
- Removes
- Lookups
- Range Queries of size up to max
- Platform
  - NUMA with four Intel Xeon E5-4650 CPUs (2.70GHz) 8 cores each with hyperthreading = **64 logical cores**
- Implementation in Java
- See paper for other benchmarks...



#### Data structures

- Introduction
- LFCA tree
- Related Work
- Evaluation
- Final Remarks

- KiWi Global version number counter for range queries Basin et al. [PPoPP'17]
- ChatterjeeSL General method for range queries Chatterjee [ICDCN'17]
- k-ary up to k items in immutable leaf nodes Brown and Avni [OPODIS'12]
- Lock-Based CA tree
  Sagonas and Winblad [JPDC'18], Winblad [ICCSW'17]
- SnapTree Fast snapshots by letting updates copy nodes Bronson et al. [PPoPP'10]

#### More...

See paper for a more detailed discussion...





 $\approx$  500K items, Inserts:10%, Removes:10%, Lookups:55%, Queries:25%-max:10 Lock-free Contention Adapting Search Trees

SPAA'2018 http://www.it.uu.se/research/group/languages/software/ca\_tree

- 14 -

Small Range Queries







 $\approx$  500K items, Inserts:10%, Removes:10%, Lookups:55%, Queries:25%-max:1000 Lock-free Contention Adapting Search Trees

SPAA'2018

http://www.it.uu.se/research/group/languages/software/ca\_tree

- 15 -





 $\approx$  500K items, Inserts:10%, Removes:10%, Lookups:55%, **Queries:25%-max:100000** 

Lock-free Contention Adapting Search Trees http://www.it.uu.se/research/group/languages/software/ca\_tree - 16 -

SPAA'2018



## **Final Remarks**



Introduction LFCA tree Related Work Evaluation

**Final Remarks** 

SPAA'2018

#### LFCA tree

 Lock-free search tree with range query support that adapts to contention



 Automatically balances the trade-offs of coarse-grained and fine-grained synchronization



 Scales better than related concurrent data structures over a wide range of scenarios





Lock-free Contention Adapting Search Trees http://www.it.uu.se/research/group/languages/software/ca\_tree - 18 -

SPAA'2018





SPAA'2018



(a) Range queries (parallel updates) (b) Updates (parallel range queries)

Figure: On the left, throughput for the 16 range query threads, and on the right, throughput for the threads doing only inserts and removes.



SPAA'2018



Figure: Single-item operations only. Throughput (operations/ $\mu$ s) on the y-axis and thread count on the x-axis. The sub-figures are ordered in increasing amount of lookups.