



Efficient and Reliable Lock-Free Memory Reclamation Based on Reference Counting

Anders Gidenstam, Marina Papatriantafilou,
Håkan Sundell and Philippos Tsigas

Distributed Computing and Systems group,
Department of Computer Science and Engineering,
Chalmers University of Technology



Outline

- Introduction
 - The Problem
 - Lock-free synchronization
- Our solution
 - Idea
 - Properties
- Experiments
- Conclusions

2



The Lock-Free Memory Reclamation Problem

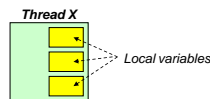
- Concurrent shared data structure
 - Dynamic use of shared memory
 - Concurrent and overlapping operations by threads or processes



3



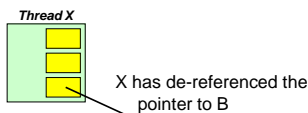
The Lock-Free Memory Reclamation Problem



4



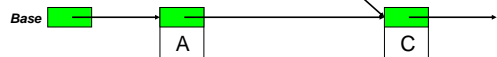
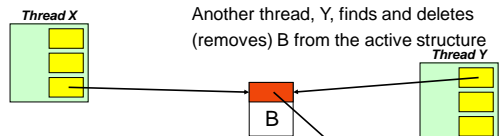
The Lock-Free Memory Reclamation Problem



5



The Lock-Free Memory Reclamation Problem

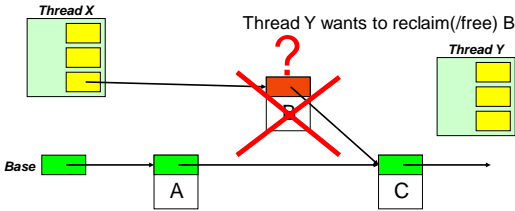


6



The Lock-Free Memory Reclamation Problem

Property I: A (de)-referenced node is not reclaimed

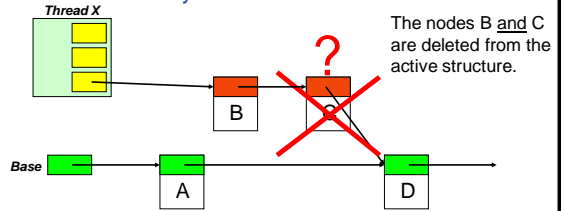


7



The Lock-Free Memory Reclamation Problem

Property II: Links in a (de)-referenced node should always be de-referencable.



8

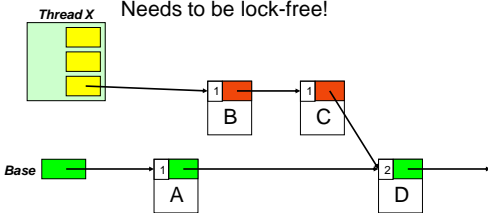


The Lock-Free Memory Reclamation Problem

Solutions?

- Garbage collection?
- Reference counting?

Needs to be lock-free!



9



Lock-free synchronization

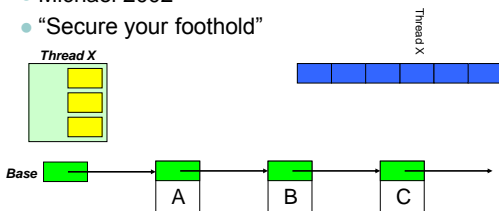
- A lock-free shared data structure
 - Allows concurrent operations without enforcing mutual exclusion (i.e. no locks)
 - Guarantees that at least one operation always makes progress
 - Avoids:
 - Blocking, deadlock and priority inversion
- Hardware synchronization primitives
 - Built into CPU and memory system
 - Typically: atomic read-modify-write instructions
 - Examples
 - Test-and-set, Compare-and-Swap, Load-Linked / Store-Conditional

10



Previous solutions

- Hazard Pointers
 - Michael 2002
 - "Secure your foothold"

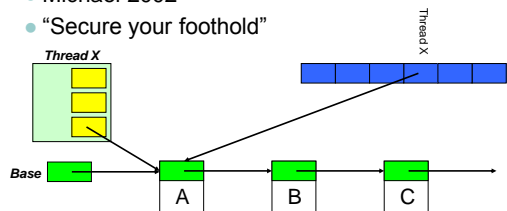


11



Previous solutions

- Hazard Pointers
 - Michael 2002
 - "Secure your foothold"

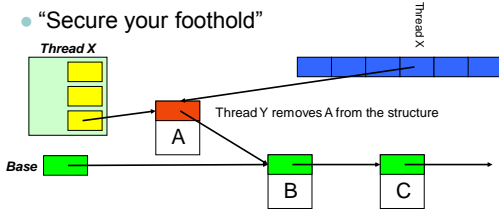


12



Previous solutions

- Hazard Pointers
 - Michael 2002
 - "Secure your foothold"

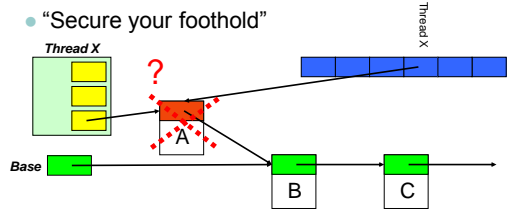


13



Previous solutions

- Hazard Pointers
 - Michael 2002
 - "Secure your foothold"

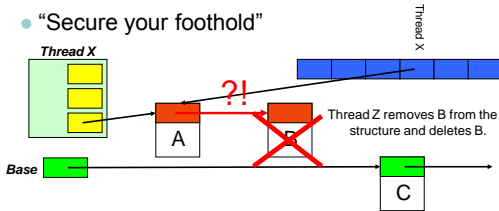


14



Previous solutions

- Hazard Pointers
 - Michael 2002
 - "Secure your foothold"

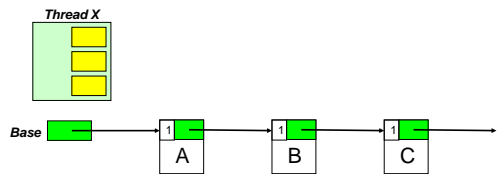


15



Previous solutions

- Lock-free Reference Counters
 - Valois + Michael & Scott 1995
 - Detlefs et al. 2001
 - Herlihy et al. 2002

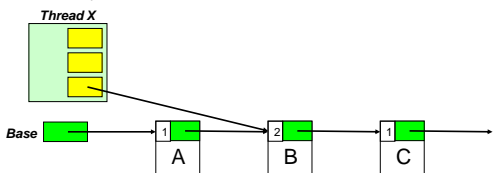


16



Previous solutions

- Lock-free Reference Counters
 - Valois + Michael & Scott 1995
 - Detlefs et al. 2001
 - Herlihy et al. 2002

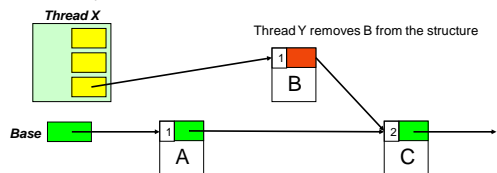


17



Previous solutions

- Lock-free Reference Counters
 - Valois + Michael & Scott 1995
 - Detlefs et al. 2001
 - Herlihy et al. 2002



18

Previous solutions

- o Lock-free Reference Counters
 - Valois + Michael & Scott 1995
 - Detlefs et al. 2001
 - Herlihy et al. 2002

19

Previous solutions

- o Lock-free Reference Counters
 - Valois + Michael & Scott 1995
 - Detlefs et al. 2001
 - Herlihy et al. 2002

Thread Z removes C from the structure.

20

Previous solutions

- o Lock-free Reference Counting
 - Valois + Michael & Scott 1995
 - Detlefs et al. 2001
 - Herlihy et al. 2002
- o Remaining issues
 - A slow thread might prevent reclamation
 - Cyclic garbage
 - Implementation practicality issues
 - Reference-count field MUST remain forever (Valois + Michael & Scott)
 - Needs double word CAS (Detlefs et al.)
 - Needs double width CAS (Herlihy, 2002)
 - Large overhead

21

Our approach – The basic idea

- o Combine the best of
 - Hazard pointers (Michael 2002)
 - Tracks references from threads
 - Fast de-reference
 - Upper bound on the amount of unreclaimed deleted nodes
 - Compatible with standard memory allocators
 - Reference counting
 - Tracks references from links in shared memory
 - Manages links within dynamic nodes
 - Safe to traverse links (also) in deleted nodes
- o Practical
 - Uses only single-word Compare-And-Swap

22

The basic idea

- o API
 - DeRefLink
 - ReleaseRef
 - CompareAndSwapRef
 - StoreRef
 - NewNode
 - DeleteNode

23

The basic idea

- o API
 - DeRefLink
 - ReleaseRef
 - CompareAndSwapRef
 - StoreRef
 - NewNode
 - DeleteNode

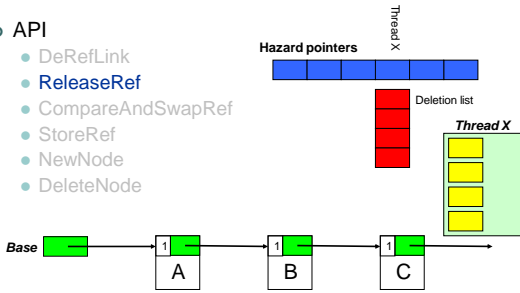
24



The basic idea

API

- DeRefLink
- **ReleaseRef**
- CompareAndSwapRef
- StoreRef
- NewNode
- DeleteNode



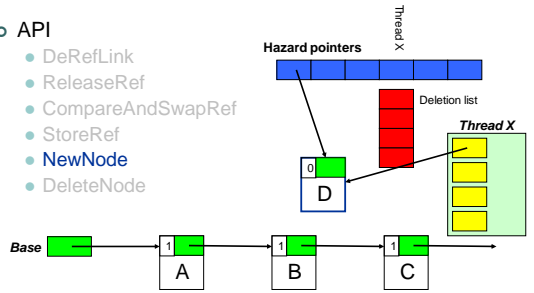
25



The basic idea

API

- DeRefLink
- ReleaseRef
- CompareAndSwapRef
- StoreRef
- **NewNode**
- DeleteNode



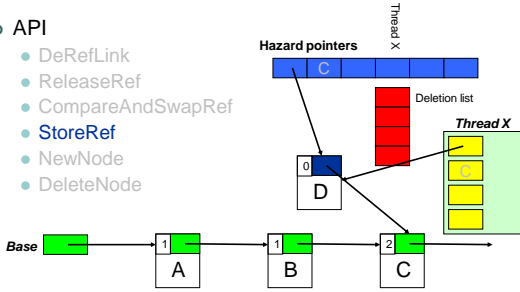
26



The basic idea

API

- DeRefLink
- ReleaseRef
- CompareAndSwapRef
- **StoreRef**
- NewNode
- DeleteNode



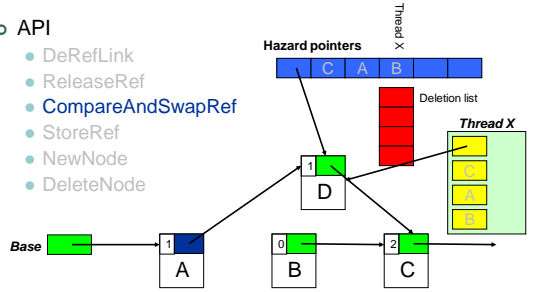
27



The basic idea

API

- DeRefLink
- ReleaseRef
- **CompareAndSwapRef**
- StoreRef
- NewNode
- DeleteNode



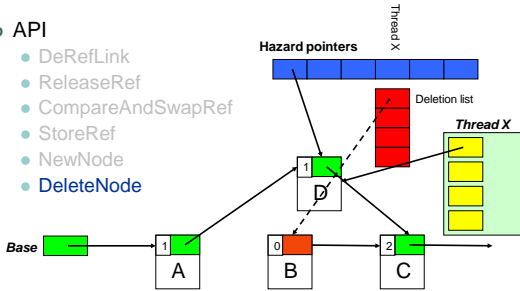
28



The basic idea

API

- DeRefLink
- ReleaseRef
- CompareAndSwapRef
- StoreRef
- NewNode
- **DeleteNode**



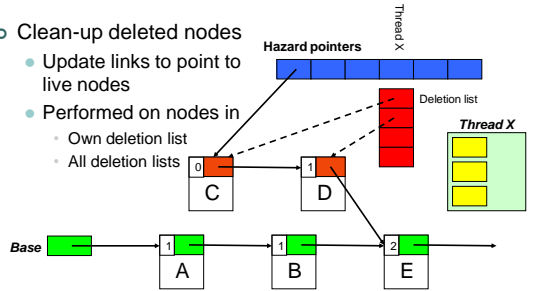
29



Breaking chains of garbage

API

- Clean-up deleted nodes
 - Update links to point to live nodes
 - Performed on nodes in
 - Own deletion list
 - All deletion lists



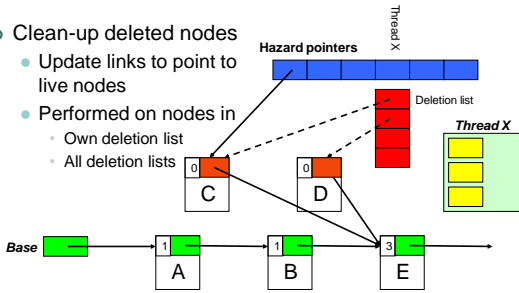
30



Breaking chains of garbage

○ Clean-up deleted nodes

- Update links to point to live nodes
- Performed on nodes in
 - Own deletion list
 - All deletion lists



31



Bound on unreclaimed nodes

- A deleted node can be reclaimed when
 - The reference count is zero and
 - No hazard pointer is pointing to it and
 - There is no ongoing clean-up of this node
- With a rate relative to the number of threads of
 - Scanning hazard pointers
 - Cleaning up nodes as needed
- Then the maximum size of each deletion list depends on
 - The number of hazard pointers
 - The number of links per node
 - The number of threads

32



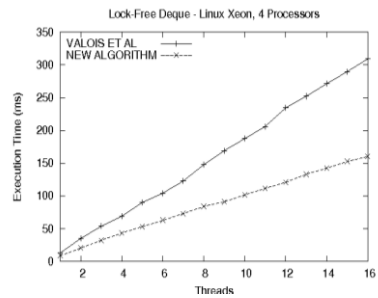
Experimental evaluation

- Lock-free deque (Sundell and Tsigas 2004) (deque – double-ended queue)
 - The algorithm needs traversal of deleted nodes
 - 10000 random operations/thread
- Tested memory reclamation schemes
 - Reference counting, Valois et al.
 - The new algorithm
- Systems
 - 4 processor Xeon PC / Linux (UMA)
 - 8 processor SGI Origin 2000 / IRIX (NUMA)

33



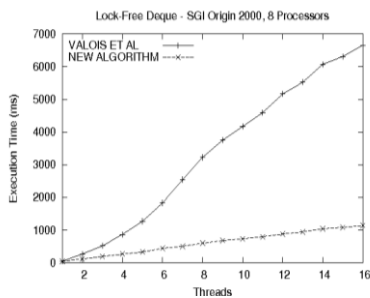
Experimental evaluation



34



Experimental evaluation



35



Conclusions

- First lock-free memory reclamation scheme that
 - Only uses atomic primitives available in contemporary architectures
 - Guarantees safety of
 - Local and
 - Global references
 - Has an upper bound on the amount of deleted but unreclaimed nodes
 - Allows arbitrary reuse of memory

36