# **Memory Reclamation**

Concurrent Algorithms Fall 2020 Igor Zablotchi



# Introduction

- So far in the course, we have assumed that memory is infinite
- This assumption needs not be true
  - In practice, memory is finite
  - Memory reclamation
- Topic of ongoing research

#### What is Memory Reclamation (MR)?

- Applications need memory
- Most realistic applications grow and shrink in memory
- Grow = allocate memory
- Shrink = free no-longer-useful memory

#### What is Memory Reclamation (MR)?

```
ds = new_data_structure(...);
node n = new_node(...);
insert(ds, n);
// use n in some way
remove(ds,n);
```



# **Freeing Memory is Necessary**

• Otherwise, applications might run out of memory or use too much memory

# **Automatic Garbage Collection**

- Some languages (e.g., Java) have automatic memory management
- Memory is allocated & freed without explicit programmer intervention
- Garbage collector decides automatically when a pointer should be freed

# **Explicit Memory Management**

- Other languages (e.g., C, C++) require the programmer to allocate & free memory explicitly
- Programmer needs to determine when to free some memory location
- This is our focus for this class

# **1-process MR is Easy**

- Allocate some memory
- Use it
- Free after last use

#### **1-process MR is Easy**



Process  $P_2$ 

 No easy way for a process to determine if a memory location will be used later by a different process











# Take-away So Far

- Memory reclamation = deciding when to free memory
- Necessary:
  - Most applications need to allocate + free
  - C, C++ are here to stay
  - No MR  $\rightarrow$  excessive memory use
- Challenging (concurrent case):
  - Need a way to determine when all processes are done with some memory location

# Outline

- Introduction
- Traditional MR Algorithms
  - Lock-free Reference Counting
  - Hazard Pointers
  - Epoch-based Reclamation
- QSense: A Hybrid MR Algorithm
- Conclusion

# **Lock-free Reference Counting**

- Main idea:
  - For each memory location, keep track of how many references are held to it.
  - When there are 0 references, safe to reclaim.



A linked list. No process has references. Each node has reference count = 1 (the reference from the previous node in the list).



Process P<sub>2</sub>

A thread is reading. The node that the thread is currently looking at has reference count = 2.



A thread is reading. The node that the thread is currently looking at has reference count = 2.



A thread is reading. The node that the thread is currently looking at has reference count = 2.



A thread has removed node  $O_3$  from the list.  $O_3$  now has reference count = 1 (the reference from the thread).



The thread has released its reference to  $O_{3.} O_3$  now has 0 references. Its memory can be freed.

## **Pros and cons of LFRC**

- ✓ Lock-free (wait-free version exists)
- ✓ Easy to understand & implement
- X Need to update reference counter on every access, even if read-only  $\rightarrow$  bad performance
- X Update of reference counter requires expensive atomic instructions → extremely bad performance!

# Outline

- Introduction
- Traditional MR Algorithms
  - Lock-free Reference Counting
  - Hazard Pointers
  - Epoch-based Reclamation
- QSense: A Hybrid MR Algorithm
- Conclusion

- Main idea:
  - Each process announces memory locations it plans to access: hazard pointers
  - Processes only free memory that is not protected by hazard pointers







- 0. Reachability
- Reachable node = can be found by following pointers from data structure root(s)



1. Announcing hazard pointers

Without hazard pointers

With hazard pointers

- 1. Read a reference p
- 2. Do something with p
- 3. (Release reference to p)

Read a reference p
 HP = p // protect p
 Check if p is still reachable. If yes, continue, otherwise

- 4. Do something with p
- 5. (Release reference to p)

- 2. Deleting elements
- Each process has a "limbo list" containing nodes that have been deleted but not yet freed
- After process p<sub>i</sub> deletes a node n from the data structure, it adds n to p<sub>i</sub>'s limbo list

- 3. Reclaiming memory
- When the limbo list grows to a certain size *R*, *p<sub>i</sub>* initiates a **scan**:
  - For each node *n* in the limbo list:
    - Look at HPs of all processes. Is any of them pointing to *n*?
    - If not, free *n*'s memory
    - (If yes, do nothing)

#### **HP Guarantees**

#### Constant time per node reclaimed + Bounded memory overhead

→ Great performance and reliability (in theory)

# **The Re-ordering Problem**

Modern architectures reorder instructions



# **The Re-ordering Problem**

Modern architectures reorder instructions

// read reference to n
Announce\_HP(n);

Check(n);
// continue using n





- Memory barriers prevent re-ordering
- But they are expensive (slow)

#### **HPs Need Barriers**

Modern architectures reorder instructions

// read reference to n
Announce\_HP(n);
Memory\_barrier();
Check(n);
// continue using n

#### **Barriers – Bad for Performance**



 $\rightarrow$  HP good in theory, slow in practice

## **Pros and Cons of HP**

- $\checkmark$  Limits memory use
- ✓ Lock-free
- X Need to update HP on every access, even if read-only  $\rightarrow$  bad performance
- **X** Need memory barriers  $\rightarrow$  bad performance
- **X** Complex to implement & use  $\rightarrow$  prone to errors

# Outline

- Introduction
- Traditional MR Algorithms
  - Lock-free Reference Counting
  - Hazard Pointers
  - Epoch-based Reclamation
- QSense: A Hybrid MR Algorithm
- Conclusion

# **Epoch-based Reclamation (EBR)**

- Main idea:
  - Processes keep track of each other's progress
  - After deleting an object, when all processes have made enough progress, memory can be freed

 Step 1: processes declare when they enter & exit critical sections



• Step 2: each process has an *epoch* (an integer, initially 0). The epoch is incremented by 1 when entering and exiting a critical section.



 $\rightarrow$  epoch is **odd** if inside critical section and **even** otherwise

 Step 3: After deleting an element, add it to a perprocess limbo list, together with current epochs of all processes



• Step 4: Periodically scan limbo list

Scan:

- cur\_vec = current epoch vector
- For each node *n* in the limbo list:
  - node\_vec = n's epoch vector
  - For each process i:
    - if node\_vec[i] is odd
      - if node\_vec[i] >= cur\_vec[i]
        - Continue to next node

• Free node

• Step 4: Periodically scan limbo list

Scan:

- cur\_vec = current epoch vector
- For each node *n* in the limbo list:
  - node\_vec = n's epoch vector
  - For each process i:
    - if node\_vec[i] is odd
      - if node\_vec[i] >= cur\_vec[i]
        - Continue to next node

• Free node



Only care about odd entries (processes inside crit. sec.)! Processes outside crit. sec. cannot access this node.

• Step 4: Periodically scan limbo list

Scan:

- cur\_vec = current epoch vector
- For each node *n* in the limbo list:
  - node\_vec = n's epoch vector
  - For each process i:
    - if node\_vec[i] is odd
      - if node\_vec[i] >= cur\_vec[i]
        - Continue to next node

• Free node



• Step 4: Periodically scan limbo list

Scan:

- cur\_vec = current epoch vector
- For each node *n* in the limbo list:
  - node\_vec = n's epoch vector
  - For each process i:
    - if node\_vec[i] is odd
      - if node\_vec[i] >= cur\_vec[i]
        - Continue to next node

• Free node



# **Pros and Cons of EBR**

- $\checkmark$  Small overhead  $\rightarrow$  very good performance
- $\checkmark$  Easy to use
- **X** Blocking (not lock-free)
  - $\rightarrow$  can invalidate lock- or wait-freedom of data structure
  - → if some process is delayed inside a critical section, memory cannot be reclaimed any more

# Outline

- Introduction
- Traditional MR Algorithms
  - Lock-free Reference Counting
  - Hazard Pointers
  - Epoch-based Reclamation
- QSense: A Hybrid MR Algorithm
- Conclusion

## HP and QSBR – Complementary

	Non-blocking	Small Overhead
EBR	×	
HP	$\checkmark$	X

# **A Hybrid Approach**



# A Hybrid Approach

- Keep track of both HPs and epochs
- When scanning:
  - If on fast path, use EBR-style scan
  - If on slow path, use HP-style scan

Ideally, we should only use memory barriers in the fallback path.

## **The Barrier Strikes Back**

n



R is reading n

- Read a pointer to a node *n* (Load)
- Assign HP to *n* (Store)
- If fallback mode is active (Load), then
  - Execute a memory barrier
- Recheck n (Load)
- Use *n* (Loads and Stores)



- D is deleting n
- Remove *n*
- If on fallback path
  - Scan hazard pointers
  - If no HPs for *n*, then
    - Free n
- Else [...]

### **The Barrier Strikes Back**



## **The Barrier Strikes Back**

It seems that we cannot turn memory barriers on and off.

But what if we could eliminate them altogether?

→ Cadence: HPs without Memory Barriers

# Cadence – Main Insight

context switch = memory barrier for process being switched out



Can we use this to replace memory barriers in the HP algorithm?

#### Cadence

Two main concepts: rooster processes and deferred reclamation

#### **Rooster Processes**



• • •

#### **Rooster Processes**



#### **Deferred Reclamation**



# **QSense: Hybrid MR**



### QSense Performance – Common Case



[50% reads, 50% updates]

#### **QSense Behavior with Delays**



## Recap

- What is memory reclamation?
- Traditional MR Techniques: LFRC, HP, EBR
- Cadence: HPs without memory barriers
- QSense: a hybrid of Cadence and EBR
  - Fast in the common case
  - Robust when necessary

# **Further Reading**

- T. E. Hart, P. E. McKenney, A. D. Brown, and J. Walpole. Performance of memory reclamation for lockless synchronization. Journal of Parallel and Distributed Computing, 67(12), 2007.
- J. D. Valois. Lock-free linked lists using compare-and-swap. PODC 1995.
- M.M. Michael, M.L. Scott. Correction of a memory management method for lock-free data structures. Technical Report TR599, Computer Science Department, University of Rochester. 1995.
- D. L. Detlefs, P. A. Martin, M. Moir, and G. L. Steele, Jr. Lock-free reference counting. PODC 2001.
- M. M. Michael. Hazard pointers: Safe memory reclamation for lock-free objects. IEEE Trans. Parallel Distrib. Syst., 15(6), 2004.
- O. Balmau, R. Guerraoui, M. Herlihy, and I. Zablotchi. Fast and Robust Memory Reclamation for Concurrent Data Structures. SPAA 2016.