Brief Announcement: Crystalline: Fast and Memory Efficient Wait-Free Reclamation

Ruslan Nikolaev *, rnikola@psu.edu, Penn State University, USA

Binoy Ravindran, binoy@vt.edu, Virginia Tech, USA

** The work was done while this author worked at Virginia Tech*

Memory Reclamation

- **Non-blocking** data structures do not use simple mutual exclusion
  - A *concurrent* thread may hold an **obsolete** pointer to an object which is about to be freed by *another* thread
  - *Safe memory reclamation* (SMR) schemes are typically used for unmanaged code (C/C++)
- Reclamation workload **balancing**
  - Read operations dominate, but data is still modified
    - In typical SMR schemes, most threads are not actively reclaiming memory
  - The problem have not received adequate attention in the literature
- **Synchronous** vs. **asynchronous** reclamation
  - In typical SMR schemes, threads *periodically* examine which objects marked for deletion can be safely freed
  - Reference counting: an *arbitrary* thread with the last reference frees an object

Memory Reclamation

- **Reference counting**
  - Impractical due to very high overheads when accessing objects
  - Hyaline [PODC’19 BA, PLDI’21] is an approach where reference counters are only used when objects are retired
    - **Pros**: asynchronous and exhibits high performance, protects against **stalled** threads
    - **Cons**: can still use unbounded memory (i.e., blocking) when threads starve
- We present **Crystalline**
  - **Crystalline-L** is based on Hyaline-1S but is lock-free even when threads starve
  - **Crystalline-W** further extends Crystalline-L to make it wait-free

Crystalline-L

- **Background (Hyaline)**
  - Threads explicitly annotate each operation
  - When objects are detached from a data structure, they are first **retired** and then **freed** when it is safe to do so
  - Hyaline-1S is a variant that bounds memory usage for **stalled** threads by explicitly tracking *local pointers* via a special **protect** method using the global era clock
    - Each allocated object is assigned a "birth era"
    - Not lock-free unless operations are periodically restarted for starving threads
    - ***Example***: one "unlucky" thread is stuck traversing a list because it keeps growing
- **Crystalline-L** adopts a different API
  - Hyaline-1S's API enables retrieving an unbounded number of local pointers
  - Alternative APIs used in Hazard Pointers [TPDS'04] or Hazard Eras [SPAA'17] explicitly differentiate each local pointer reservation in **protect**

Crystalline-L: Challenges

- Hyaline-1S aggregates objects in a **batch**
  - Can only **retire** the entire batch
  - Each thread has its own **retirement list,** and each object from the batch is inserted to the corresponding list
  - One of the objects keeps a per-batch reference counter
  - Needs at least MAX_THREADS+1 objects per a batch
- Crystalline-L handles MAX_IDX local pointers
  - The above problem is further aggravated
  - Needs at least MAX_THREADS×MAX_IDX+1 objects per a batch

Crystalline-L: Solution

- The required number of objects is **much lower** in practice
  - Each object is appended to the respective list only if the list’s era overlaps with the batch’s minimum birth era
- Crystalline-L uses **dynamic batches**
  - retire first checks how many lists are to be changed for the batch to be fully retired and records the location of the corresponding (per-thread) lists
  - If the number of objects in the batch suffices, retire completes by appending the objects to their corresponding lists
  - Otherwise, retire is repeated later when more objects are available
    - But the number of iterations is still **bounded** by the worst-case number of objects

Crystalline-W: Challenges

- Crystalline-L is only lock-free because
  - **retire** has an unbounded loop: **protect** or another **retire** contends on the same list
    - Does not let a CAS loop in retire to converge
  - **protect** has an unbounded loop which must converge on the era value
    - The era clock unconditionally increments when a new object is allocated

Crystalline-W: Solution

- The first problem with **retire**
  - When “traversing” retirement lists, i.e., dereferencing a thread from each batch that appears in its retirement list, **next** pointers in the corresponding list are ***tainted*** with SWAP
  - **retire** attaches new objects with SWAP rather than a CAS loop
    - If the **next** field of the new object is intact, the old list is attached as a tail (using CAS)
    - If the **next** field of the new object is tainted, **retire** traverses the “docked tail” (i.e., the old list) on behalf of the thread that tainted **next**
    - Some corner cases exist but are handled in wait-free fashion

Crystalline-W: Solution

- The second problem with **protect**
  - Adopts a mechanism similar to that of **Wait-Free Eras** [PPoPP’20]
    - The fast-path-slow-path approach to coordinate global era clock increments
    - Helping other threads before incrementing the era clock
  - Despite similarities, **Crystalline-W** diverges from Wait-Free Eras significantly
    - Cannot rescan retirement lists multiple times due to asynchronous reclamation
    - Uses *special* tricks: odd and even tags, an array of parent objects, “terminal” nodes in the retirement lists, etc.

Evaluation

None: no reclamation (leak memory)
HP: Hazard Pointers [TPDS'04]
HE: Hazard Eras [SPAA'17]
IBR: 2GE Interval-Based Reclamation [PPoPP'18]
WFE: Wait-Free Eras [PPoPP'20]
Hyaline: Hyaline-1 and Hyaline-1S [PODC'19 BA, PLDI'21]
EBR: Epoch-Based Reclamation

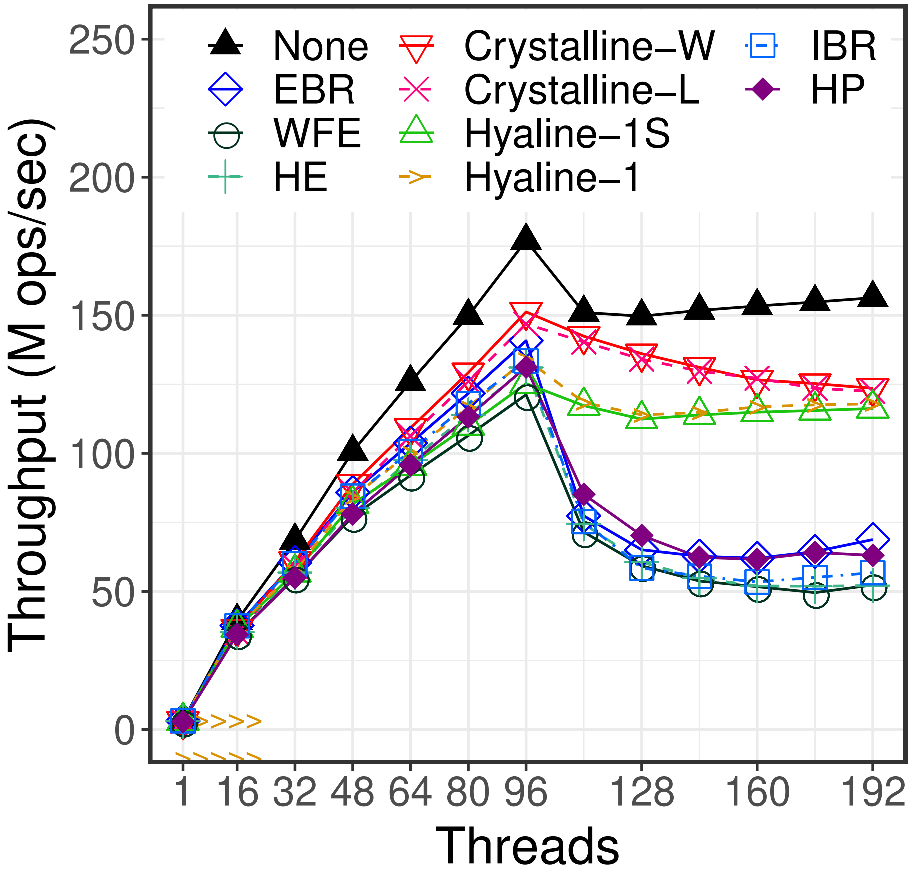

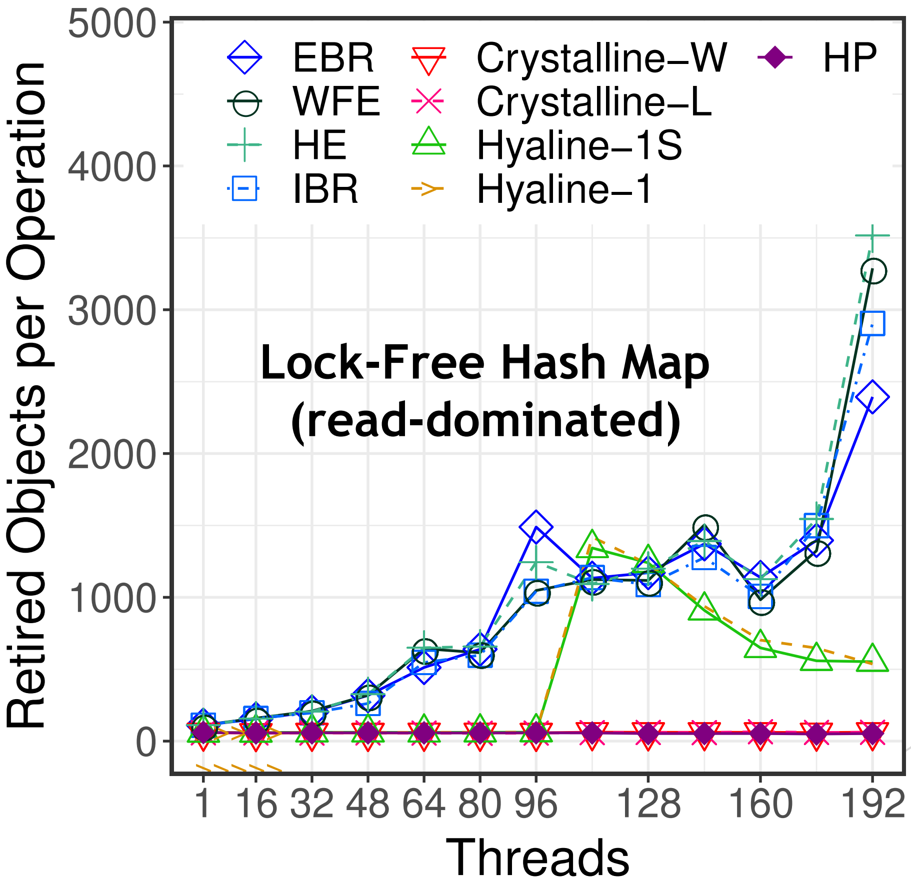

4 x Intel Xeon E7-8890 v4 2.20 GHz CPUs (96 cores), 256GB of RAM

More Details

- Code is open-source and available at:
  - https://github.com/rusnikola/wfsmr
- Full paper is available as an arXiv report:
  - https://arxiv.org/abs/2108.02763

THANK YOU!