2020-11-26
KAIST Software Graduate Program

When is it safe to free memory in concurrent programming?

An opinionated survey of memory reclamation algorithms

Jeehoon Kang

(A joint work with Jaehwang Jung)

KAIST Korea Advanced Institute of Science and Technology

Concurrent Data Structures (CDS)

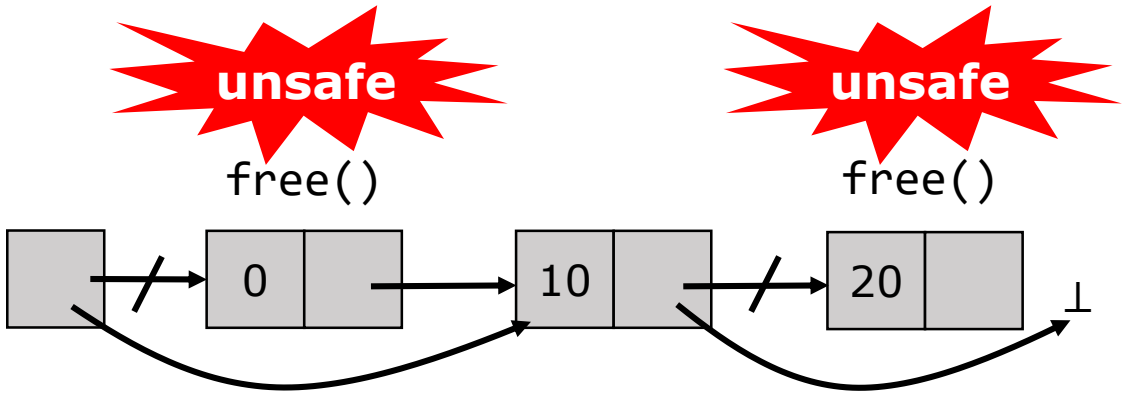

+ fast & non-blocking
- complex design
  - …
  - **memory management**

Memory Reclamation in CDS:
You Are Not Free to `free()` the Removed Block

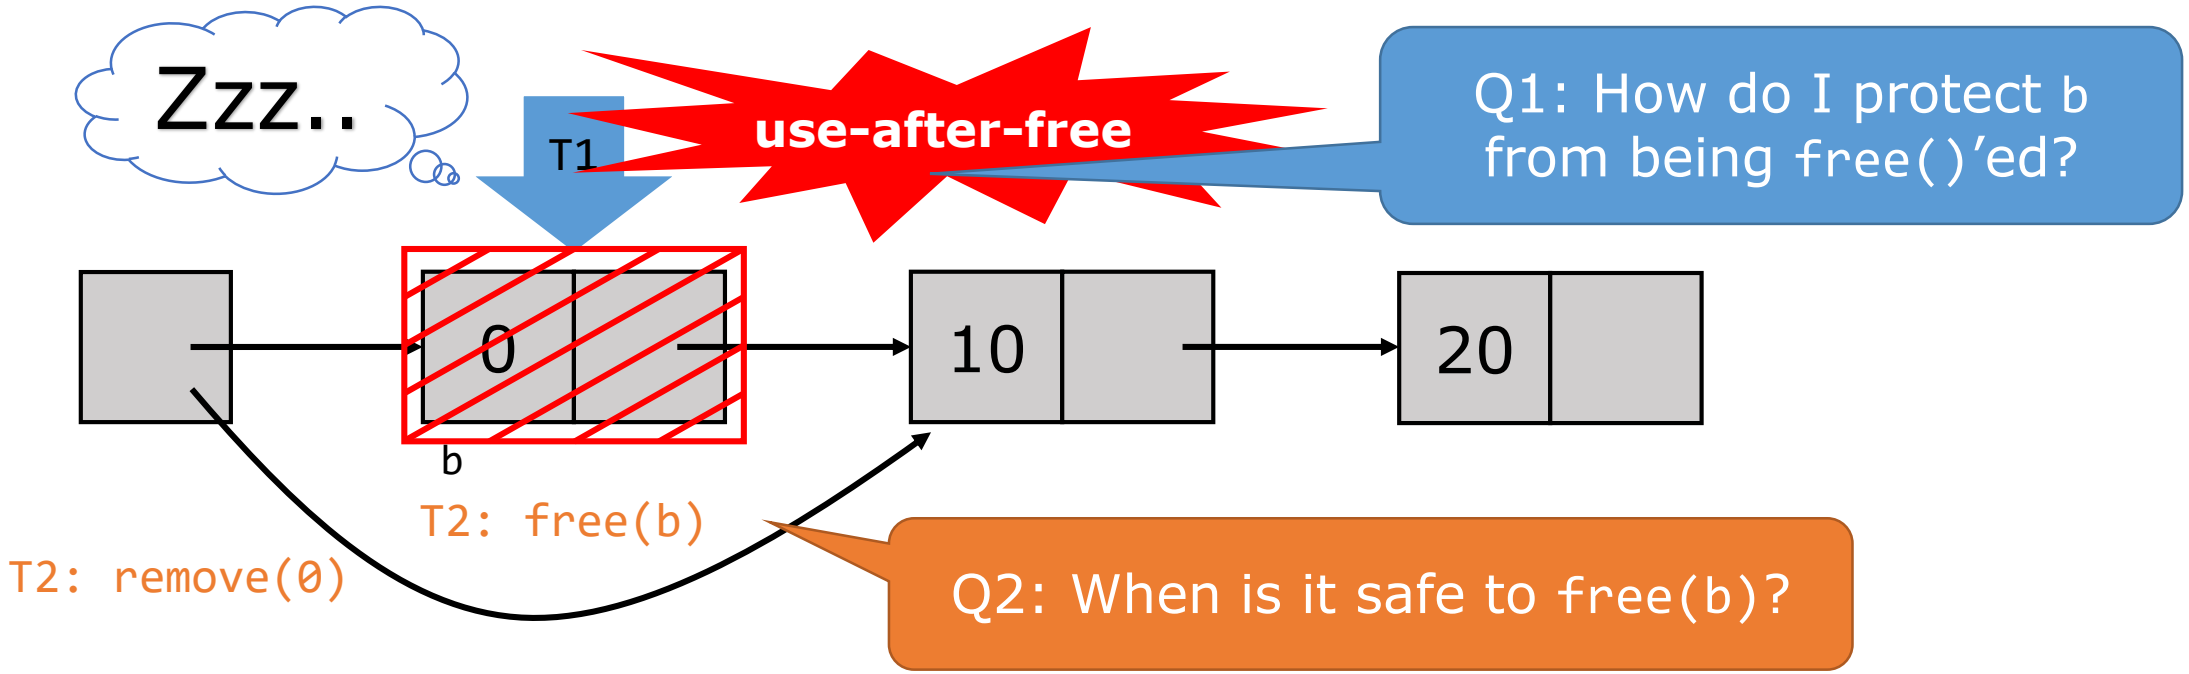

Solution: Safe Memory Reclamation **(SMR)** Algorithms

The SMR Algorithms Literature

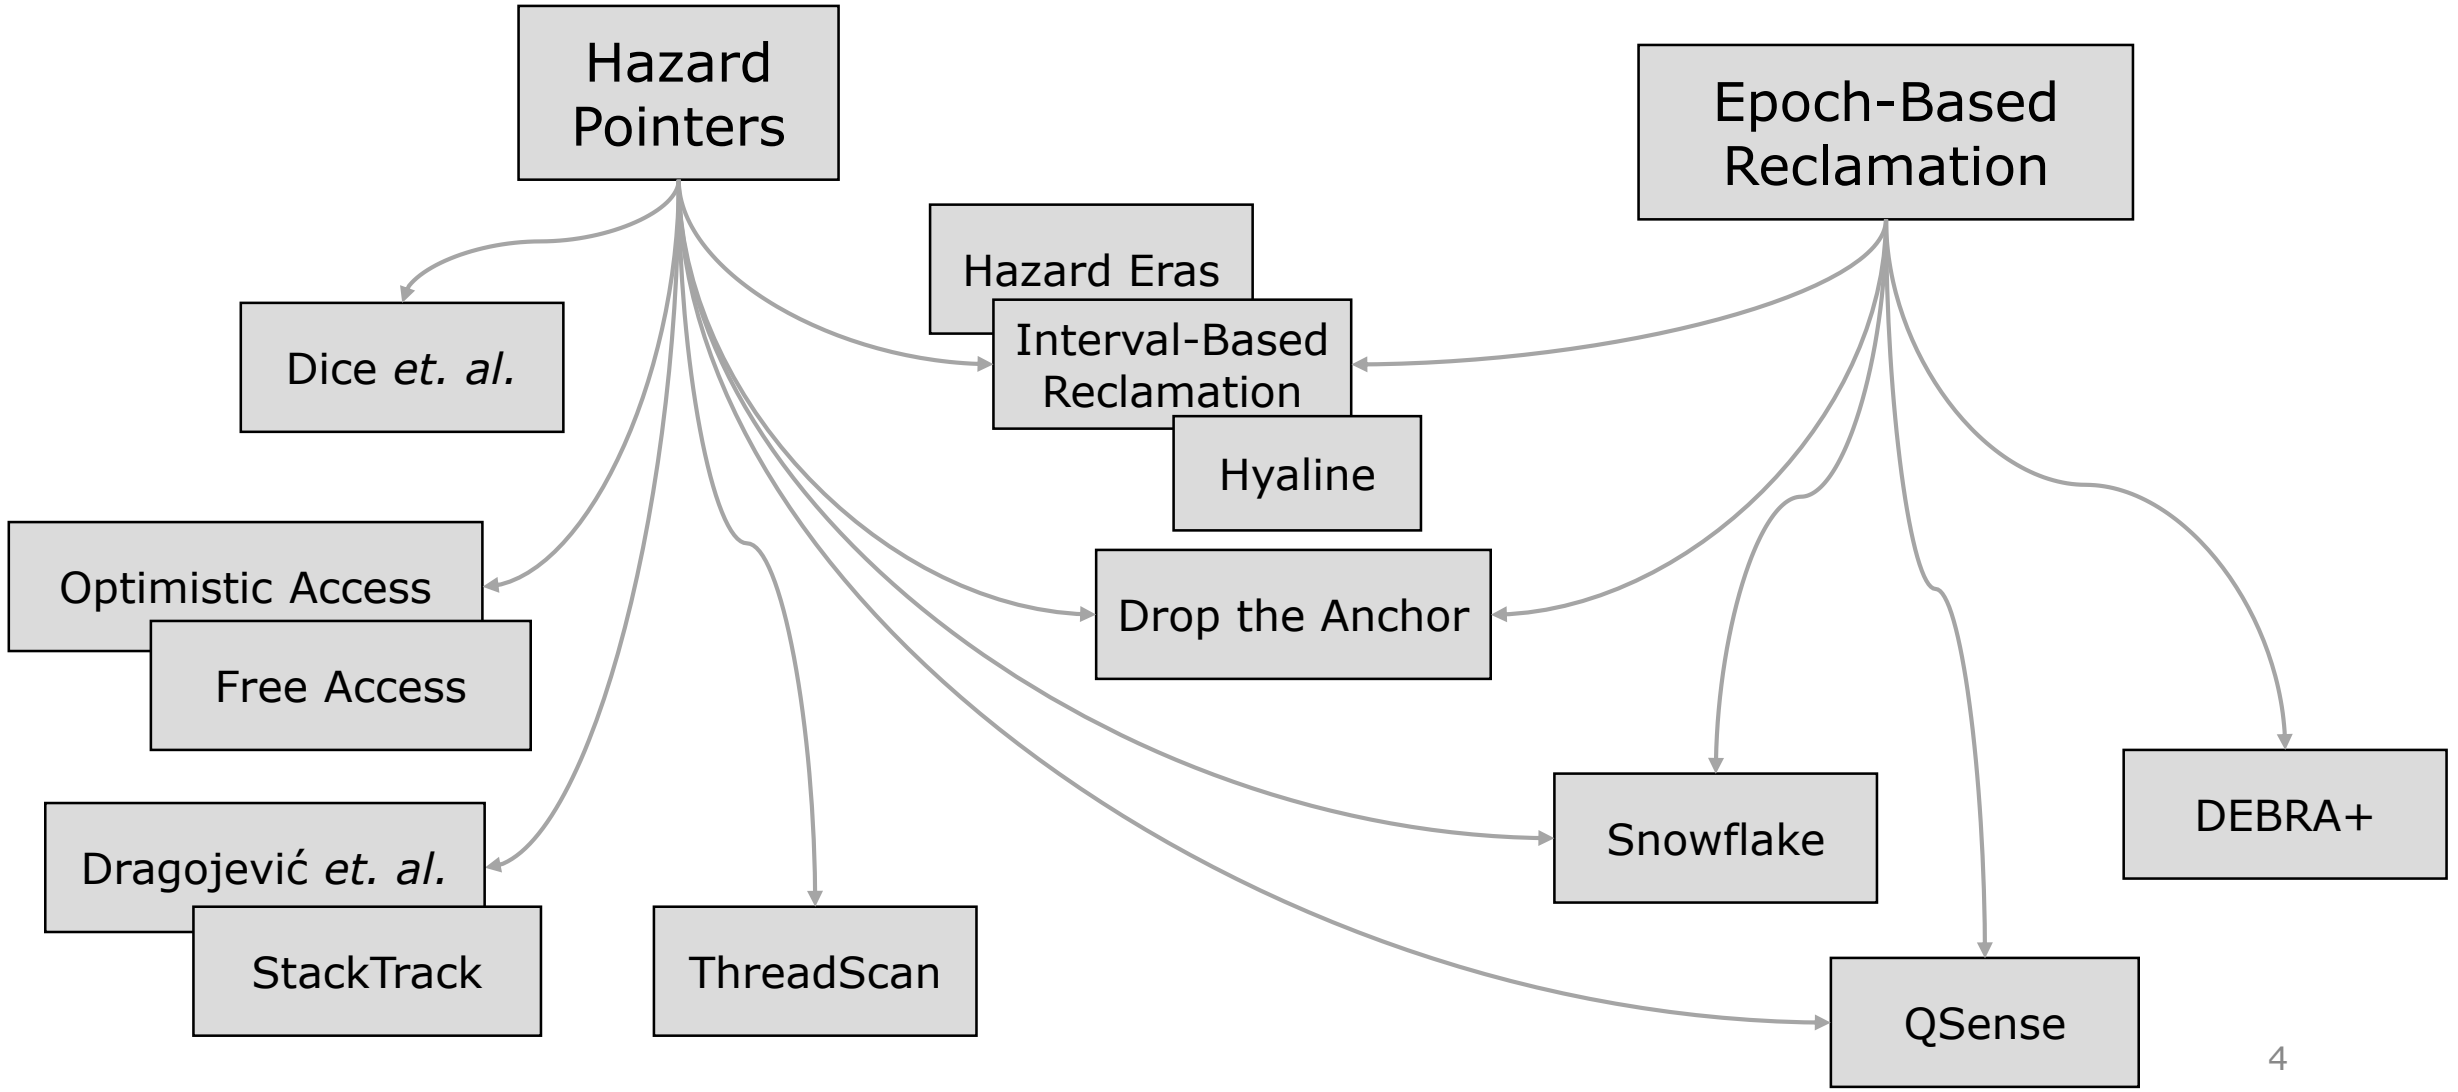

SMR Algorithm #1
Hazard Pointers (HP)

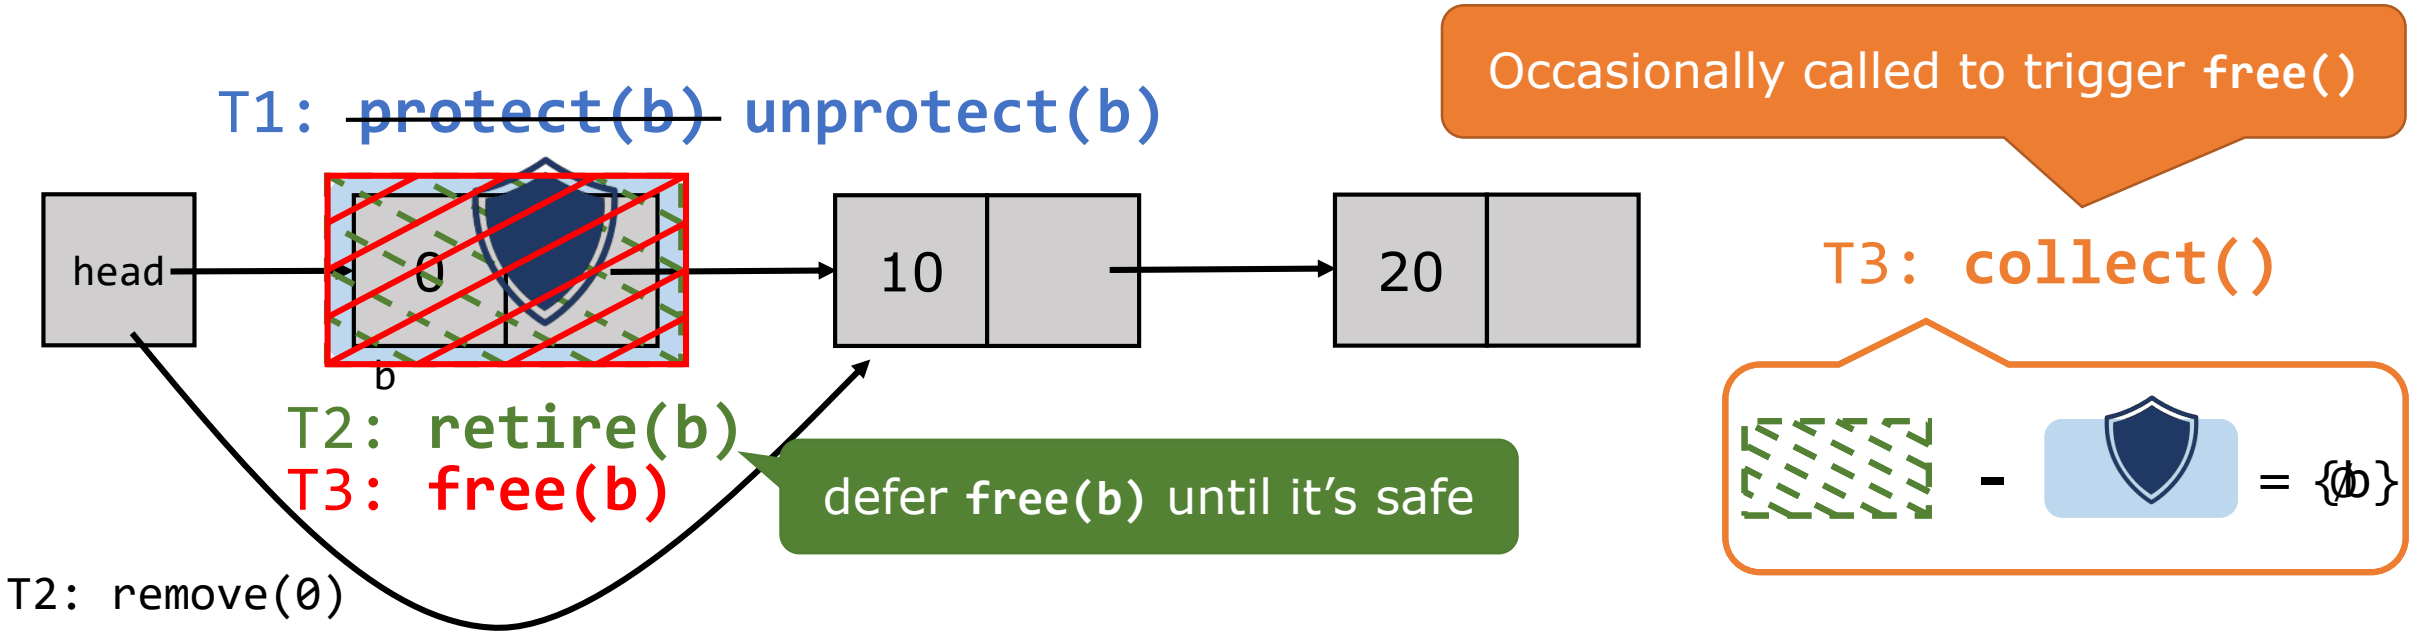

Q1: How do I protect b from being free()'ed?

A1: protect(b) (called *hazards*)

Q2: When is it safe to free(b)?

A2: When it is no longer protect()'ed.

HP's Example
HP-Protected Treiber's Stack

```
1  struct Node<T> { Node<T>* next; T data; };
2  struct Stack<T> {
3    Atomic<Node<T>*> head;        // nullptr, initially
4    void push(T data) {
5      auto node = new Node<T>{nullptr, data};
6      do {
7        auto cur = this->head.load();
8        node->next = cur;
9      } while (!this->head.cas(cur, node));
10   }
11   optional<T> Stack<T>::pop() {
12     auto cur = nullptr;
13     loop {
14       cur = this->head.load();
15       if (cur == nullptr) return {};
16       auto next = cur->next;
17       if (this->head.cas(cur, next)) {
18         free(cur); break;        // unsafe reclamation
19       }
20     }
21     return std::move(cur->data);
22   }
23 }
```

unsafe

protect(cur);
if (this->head.load() != cur)
    continue;

unprotect(cur);

retire(cur);

HP's Drawbacks
Not Fast & Not Widely Applicable

Not fast: requires per-`protect()` synchronization with expensive store-load fence

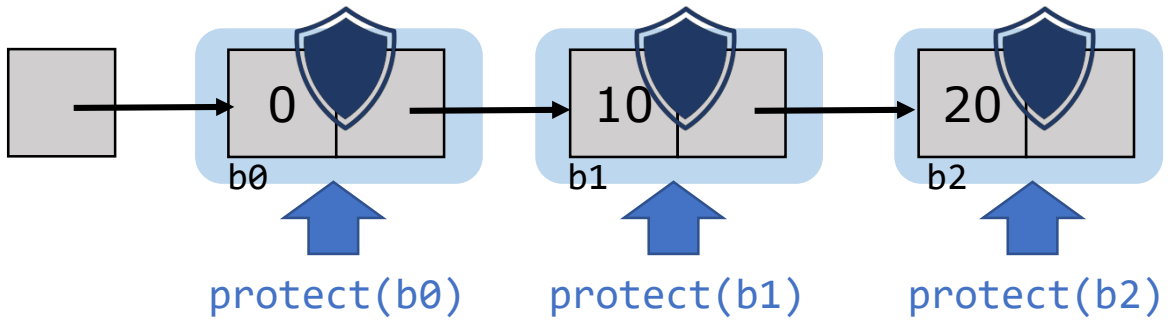

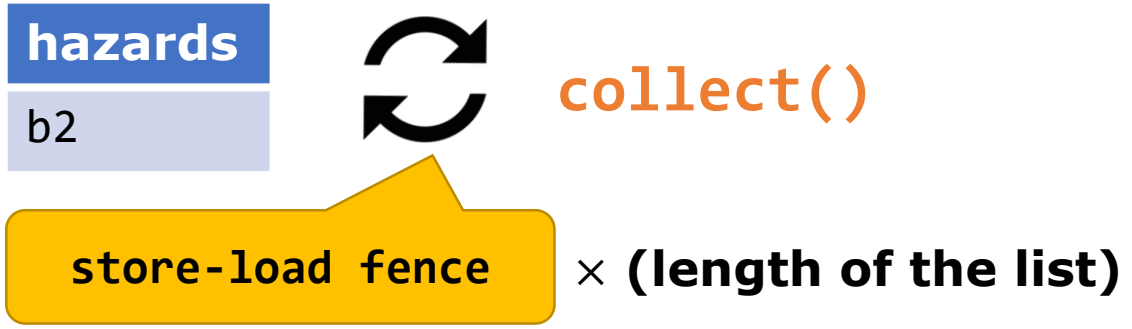

Not widely Applicable: doesn't support "chained retirement"

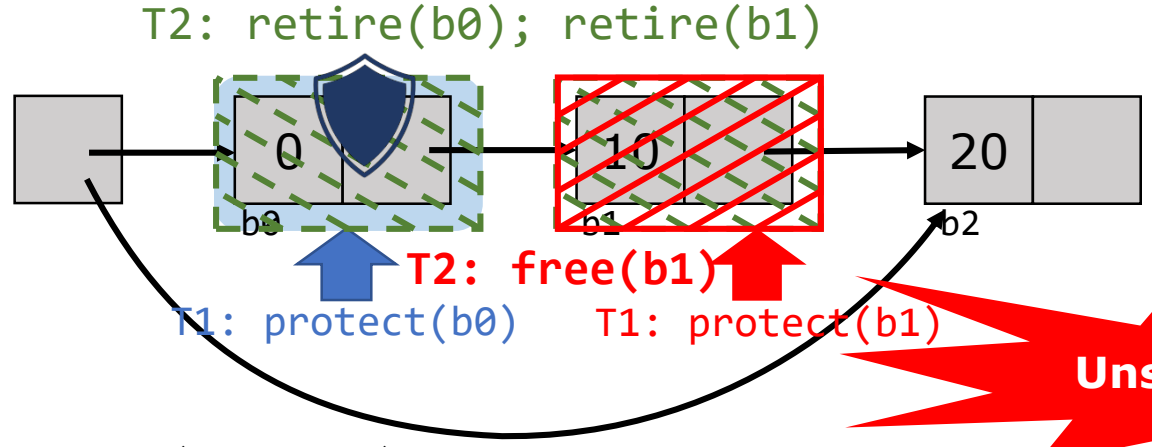

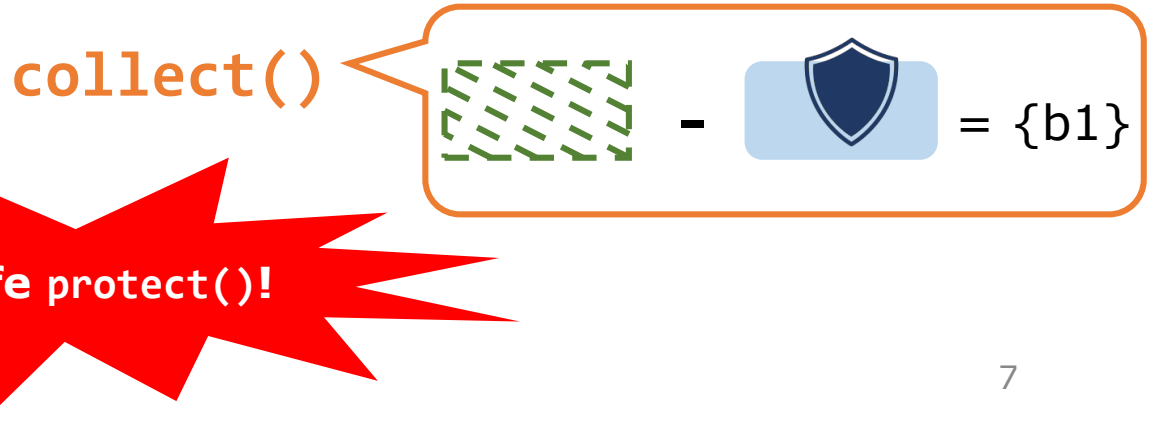

T2: remove([0, 10])

SMR Algorithm #2
Epoch-Based Reclamation (EBR)

Fast: synchronization for each active state

Widely applicable: Protection for all potential accesses

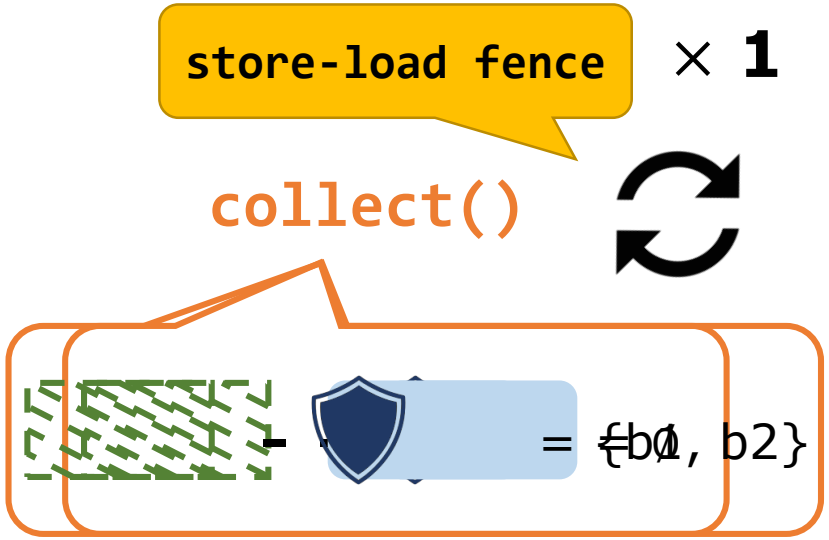

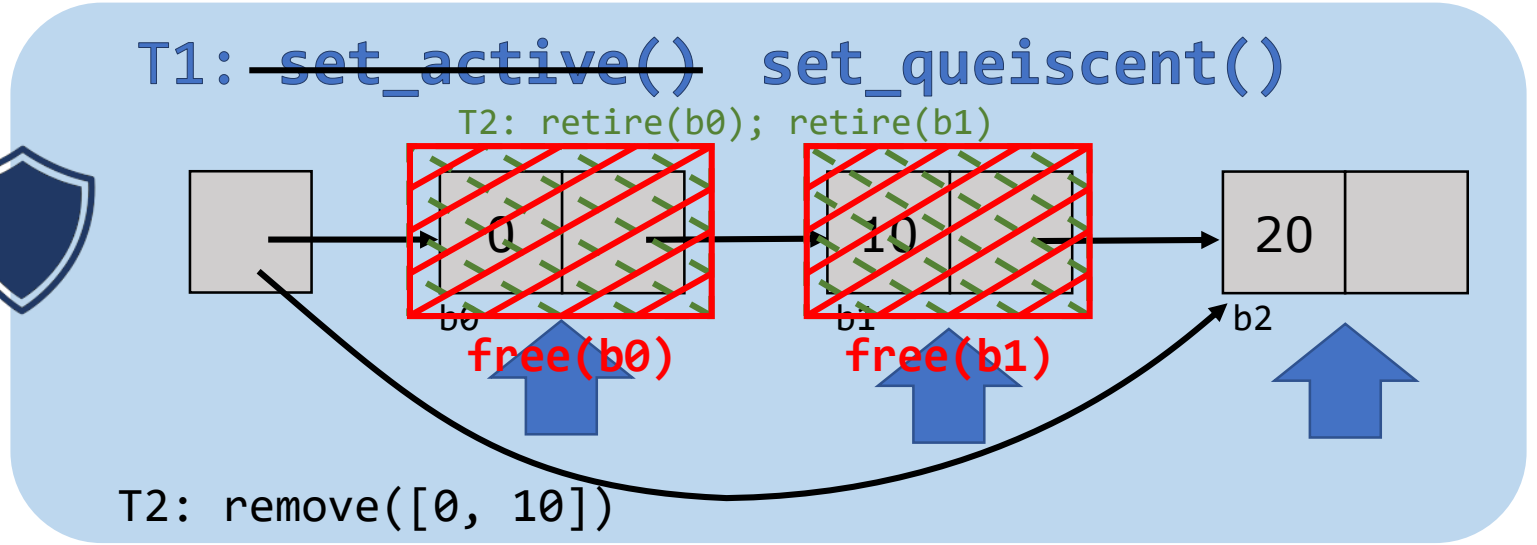

Q1: How do I protect `b` from being `free()`'ed?

A1: `set_active()` for

How to decide?

Q2: When is it safe to `free(b)`?

A2: When `b` can't be accessed by any active threads.

Epoch Consensus

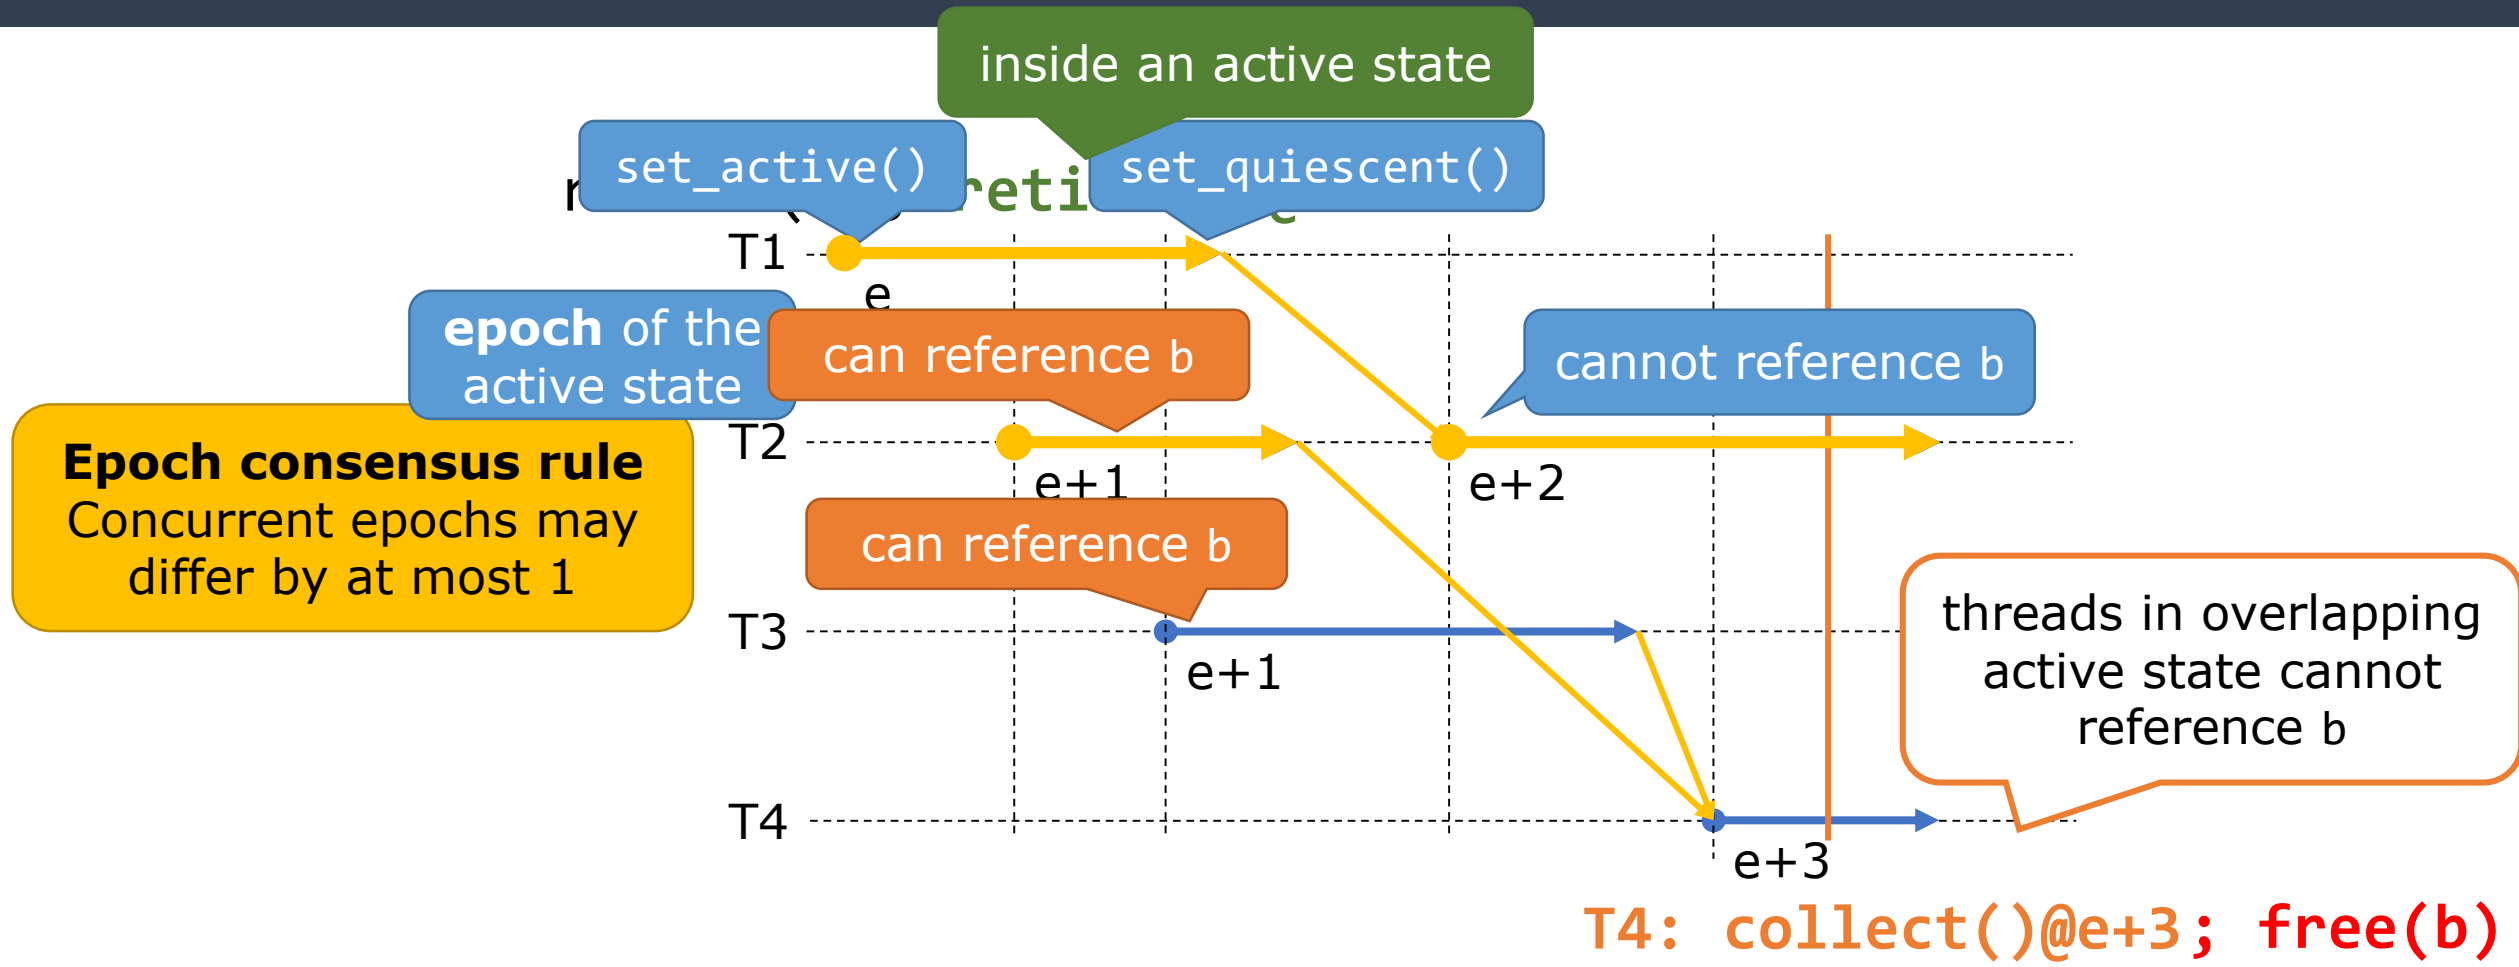

Q2: When is it safe to `free(b)`?

A2: At epoch `e+3`

EBR's Example
EBR-Protected Treiber's Stack

```
1  struct Node<T> { Node<T>* next; T data; };
2  struct Stack<T> {
3    Atomic<Node<T>*> head;        // nullptr, initially
4    void push(T data) {
5      auto node = new Node<T>{nullptr, data};
6      do {
7        auto cur = this->head.load();
8        node->next = cur;
9      } while (!this->head.cas(cur, node));
10   }
11   optional<T> Stack<T>::pop() {
12     auto cur = nullptr;
13     loop {
14       cur = this->head.load();
15       if (cur == nullptr) return {};
16       auto next = cur->next;
17       if (this->head.cas(cur, next)) {
18         free(cur); break;        // unsafe reclamation
19       }
20     }
21     return std::move(cur->data);
22   }
23 }
```

set_active();

retire(cur);

set_quiescent();

EBR's Drawback: Not Robust

If a thread doesn't exit its active state, reclamation is indefinitely blocked.

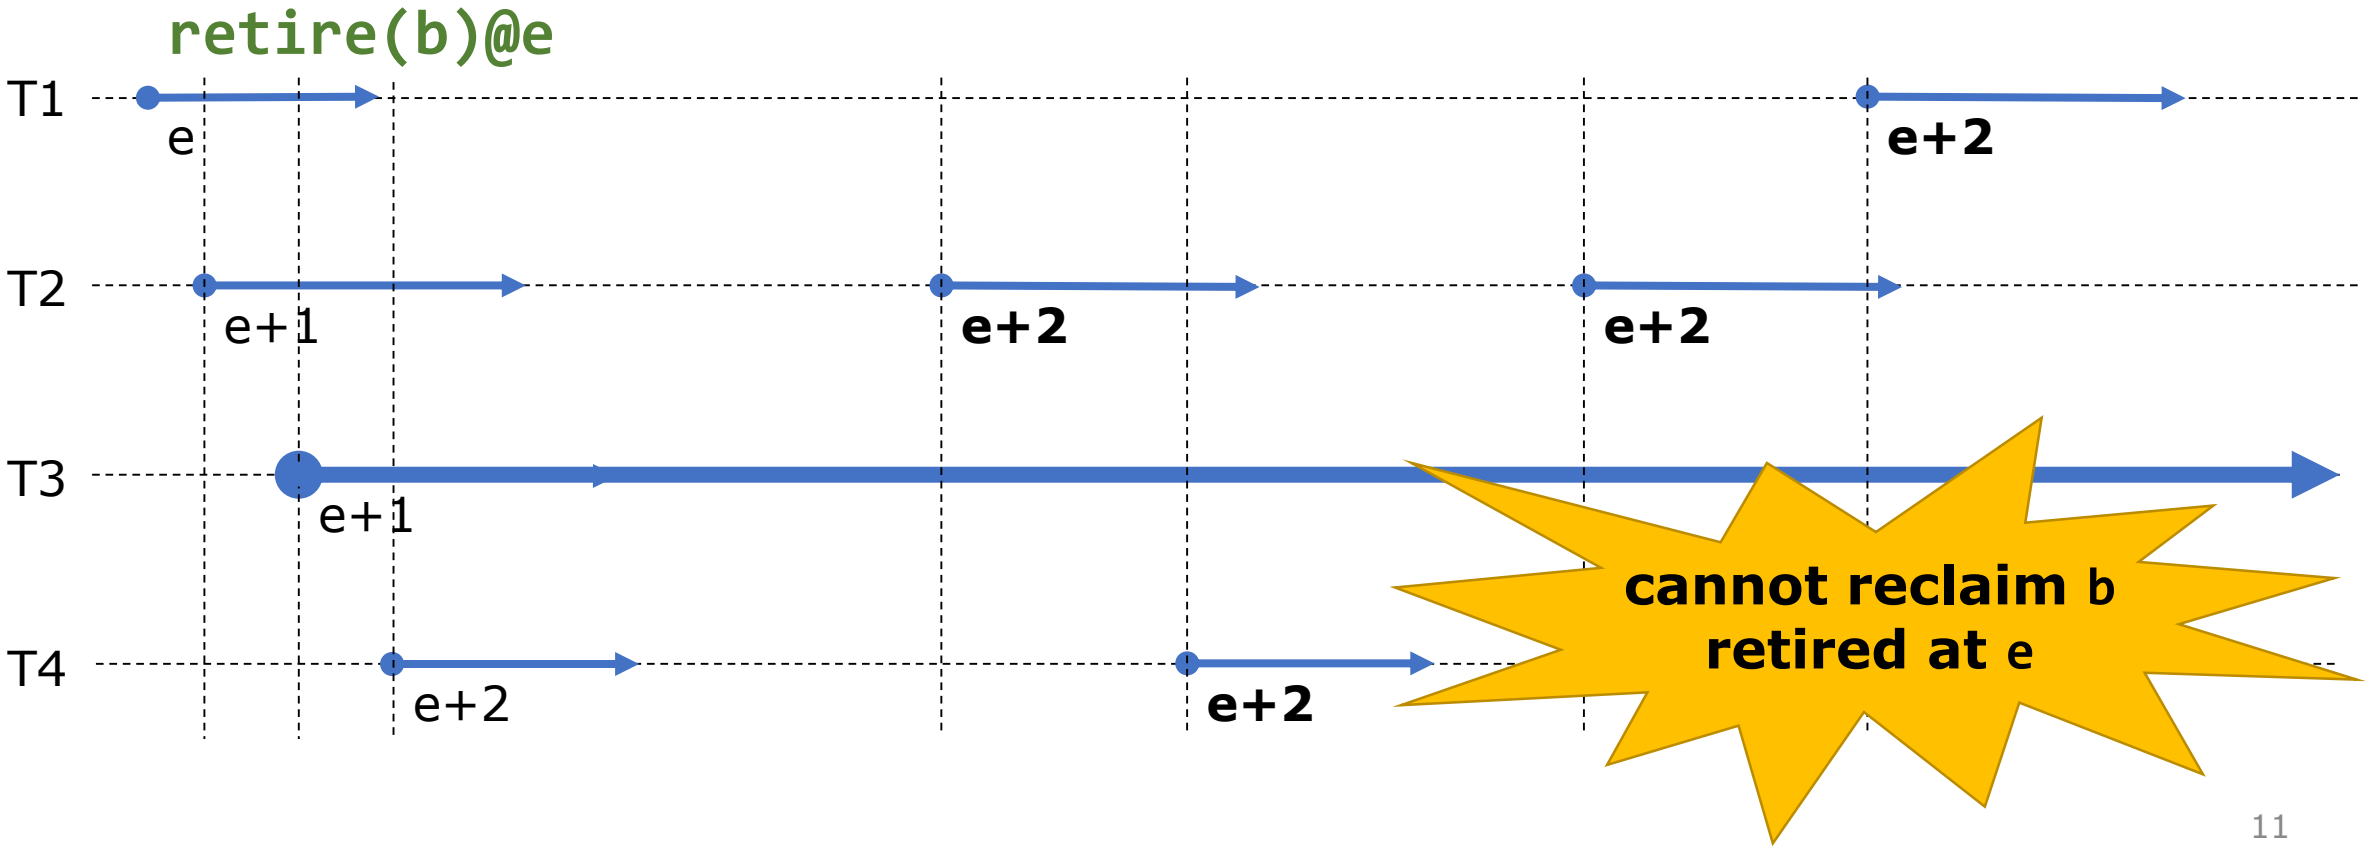

SMR Algorithms Literature

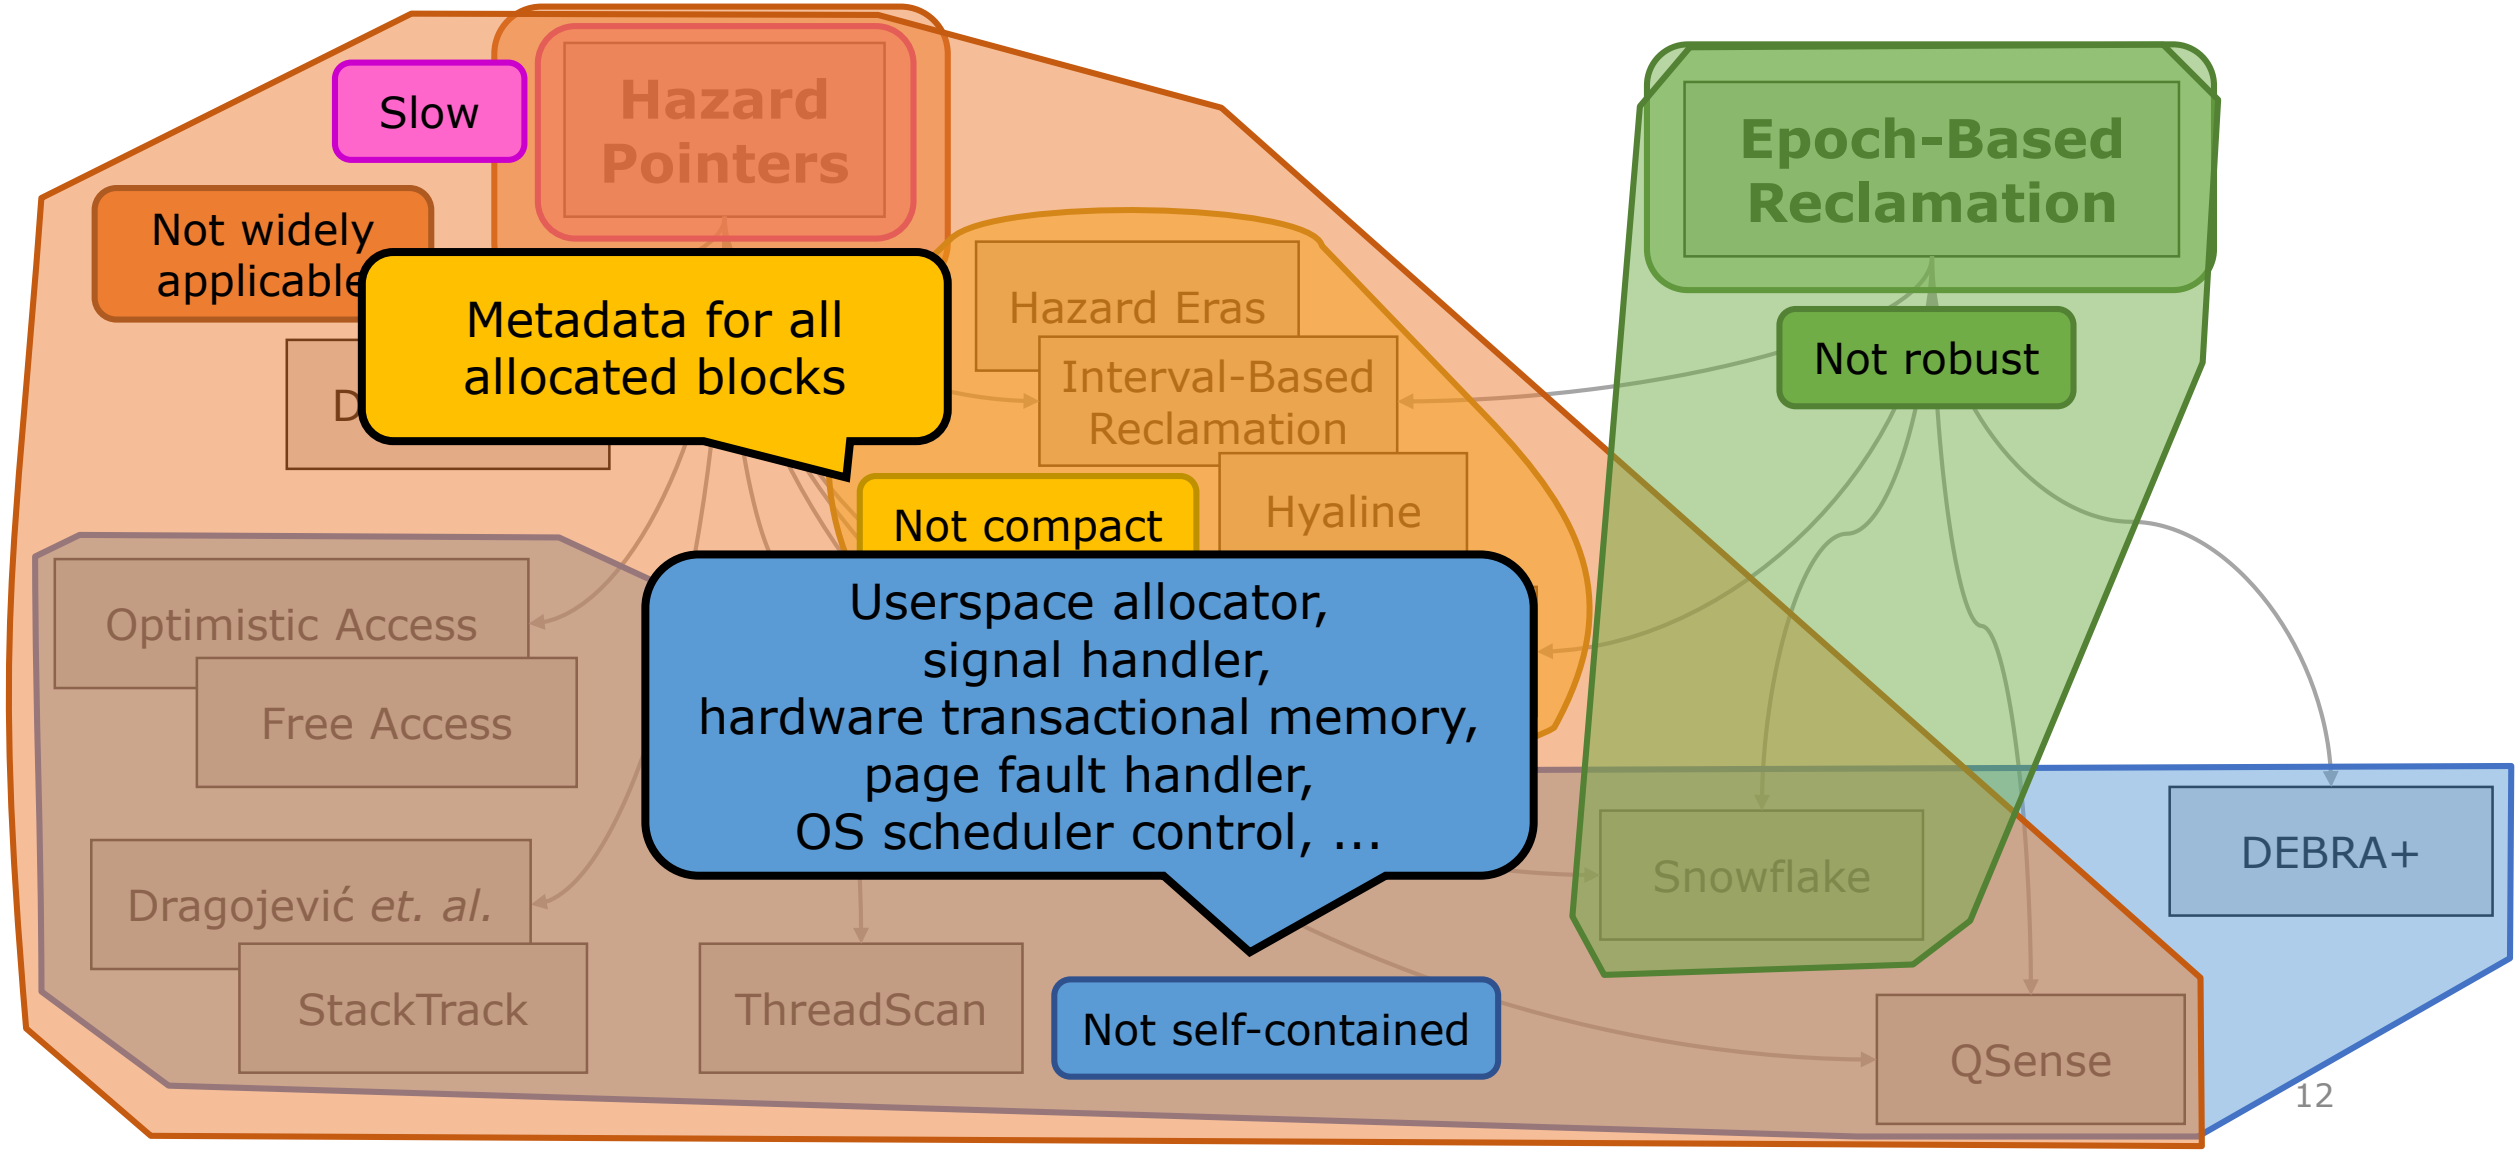

Pointer-and-Epoch-Based Reclamation

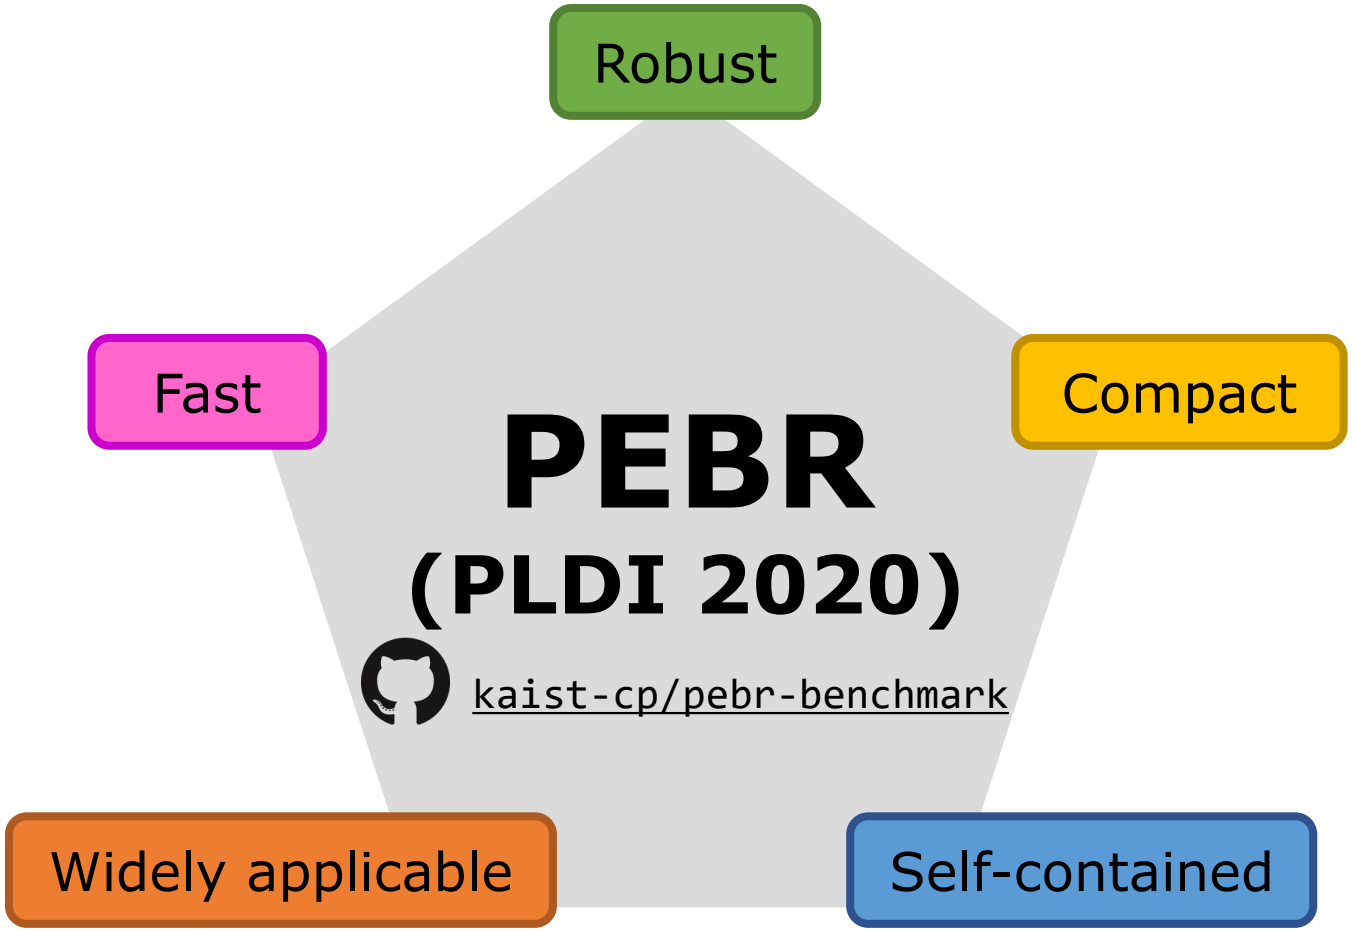

1. Hybrid of EBR and HP using **ejection**.
2. **Widely applicable** API even in the presence of ejection
3. **Robust**, **self-contained** and **compact** ejection algorithm

Making EBR Robust with Ejection

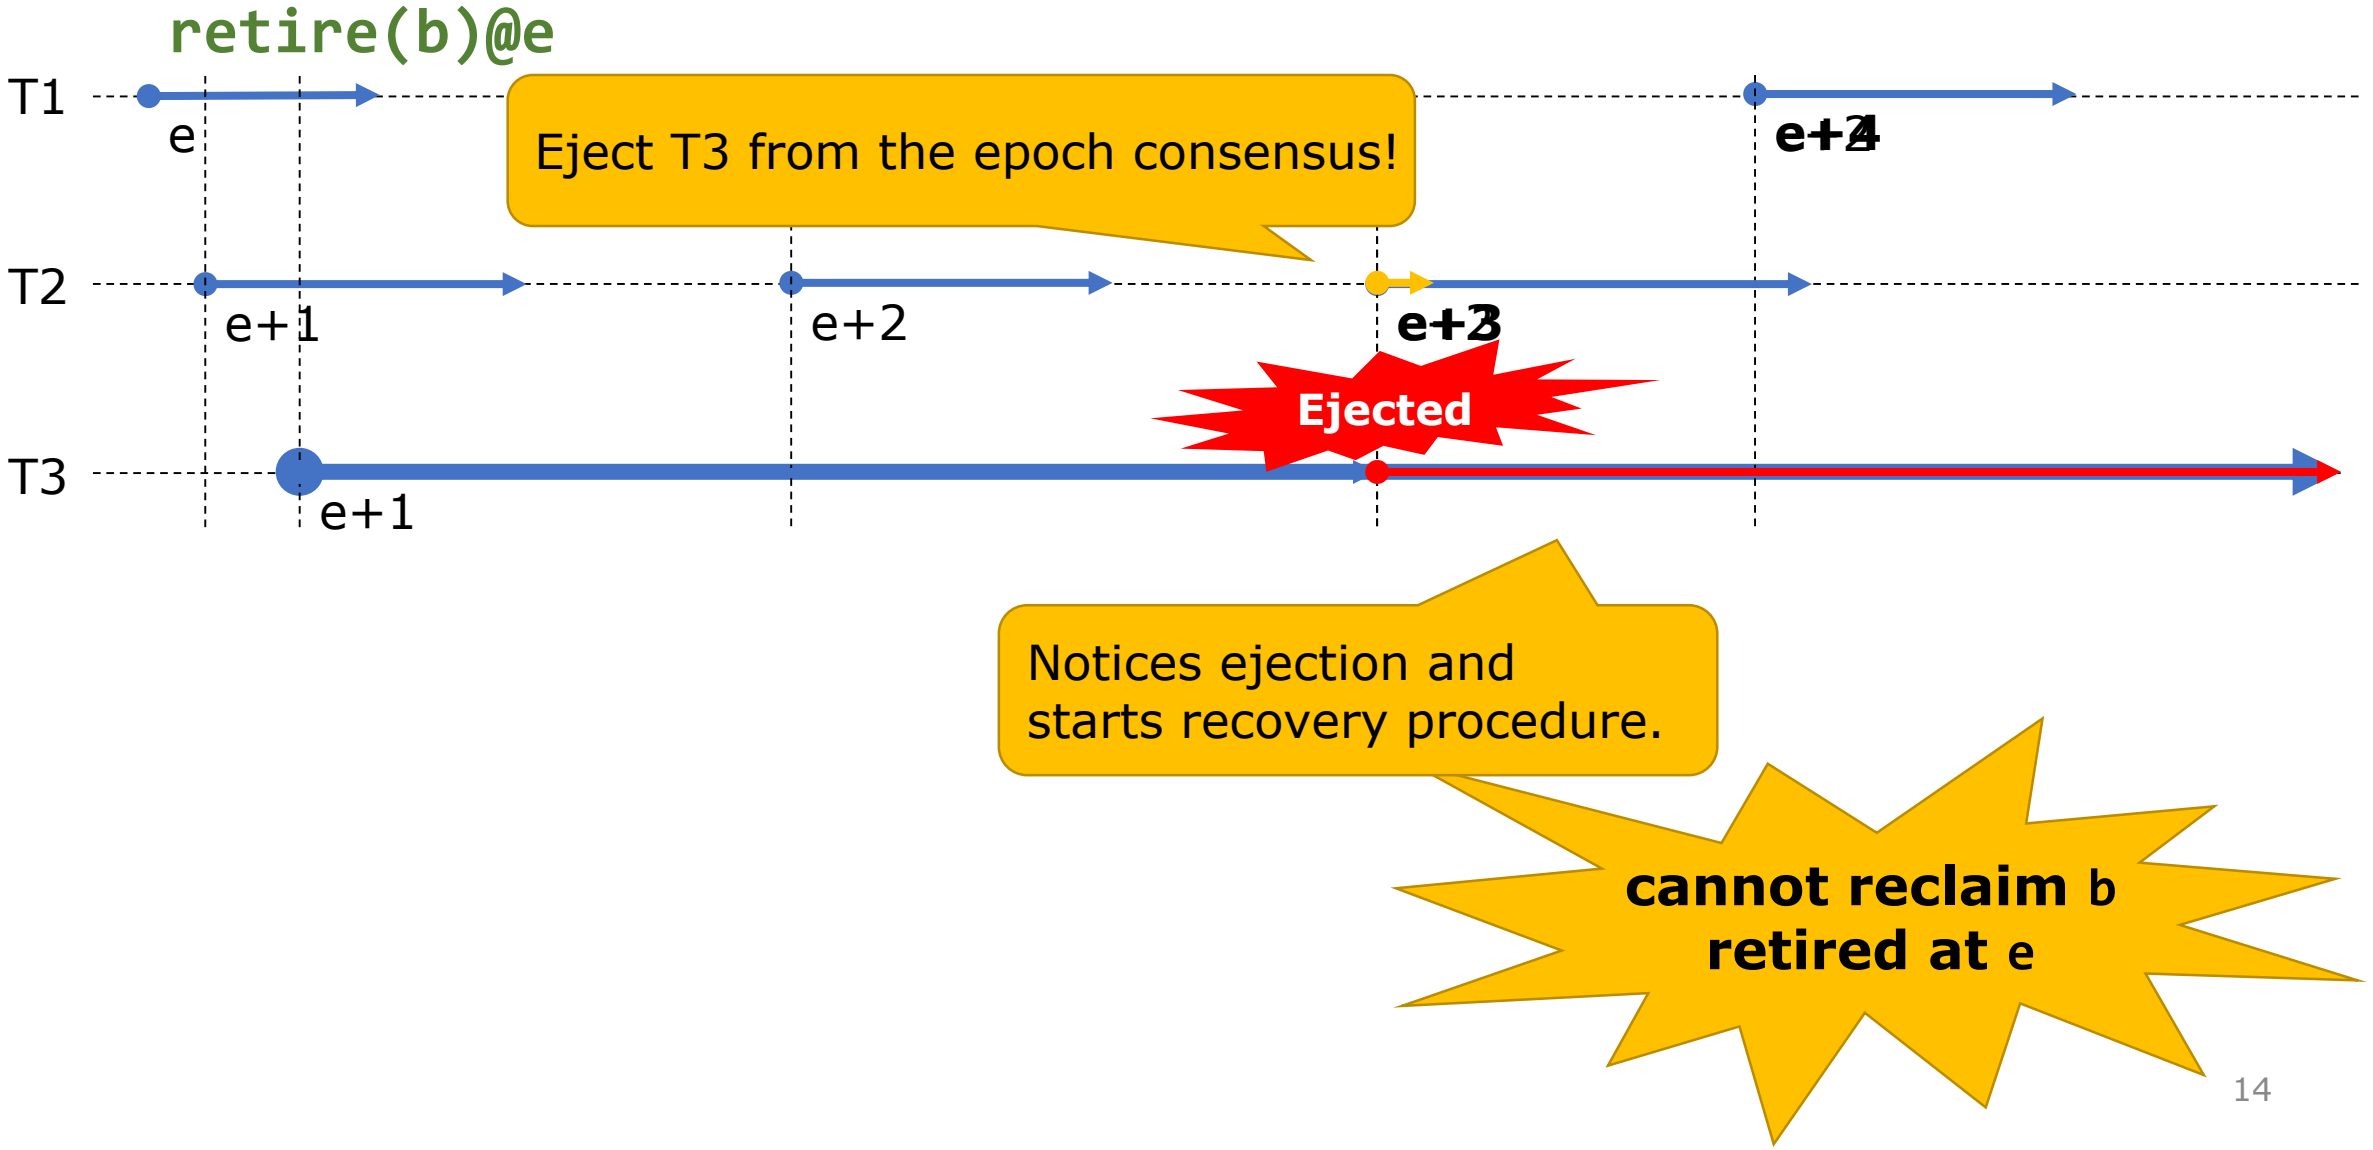

PEBR in a Nutshell

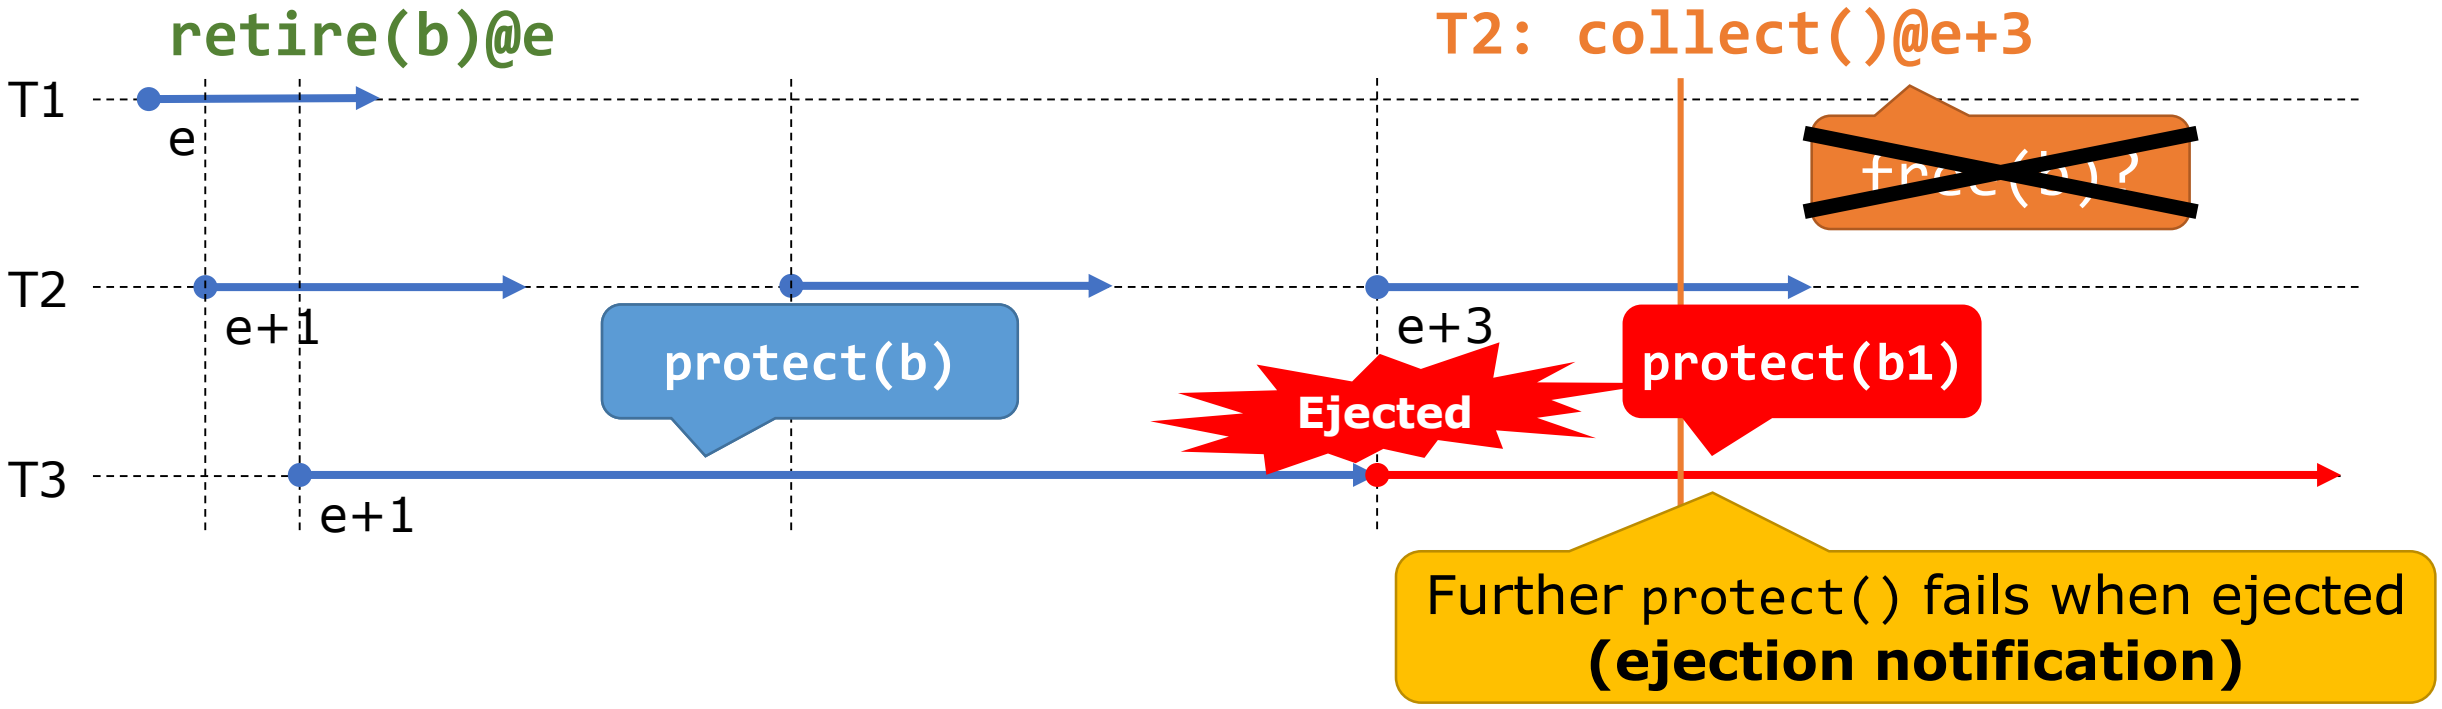

Q1: How do I protect b from being free()'ed?

A1: protect(b) inside an active state.
If failed, then you're ejected.

Q2: When is it safe to free(b)?

A2: At epoch **e+3**, if isn't not protect()'ed.

SMR Algorithms Literature: Ejection

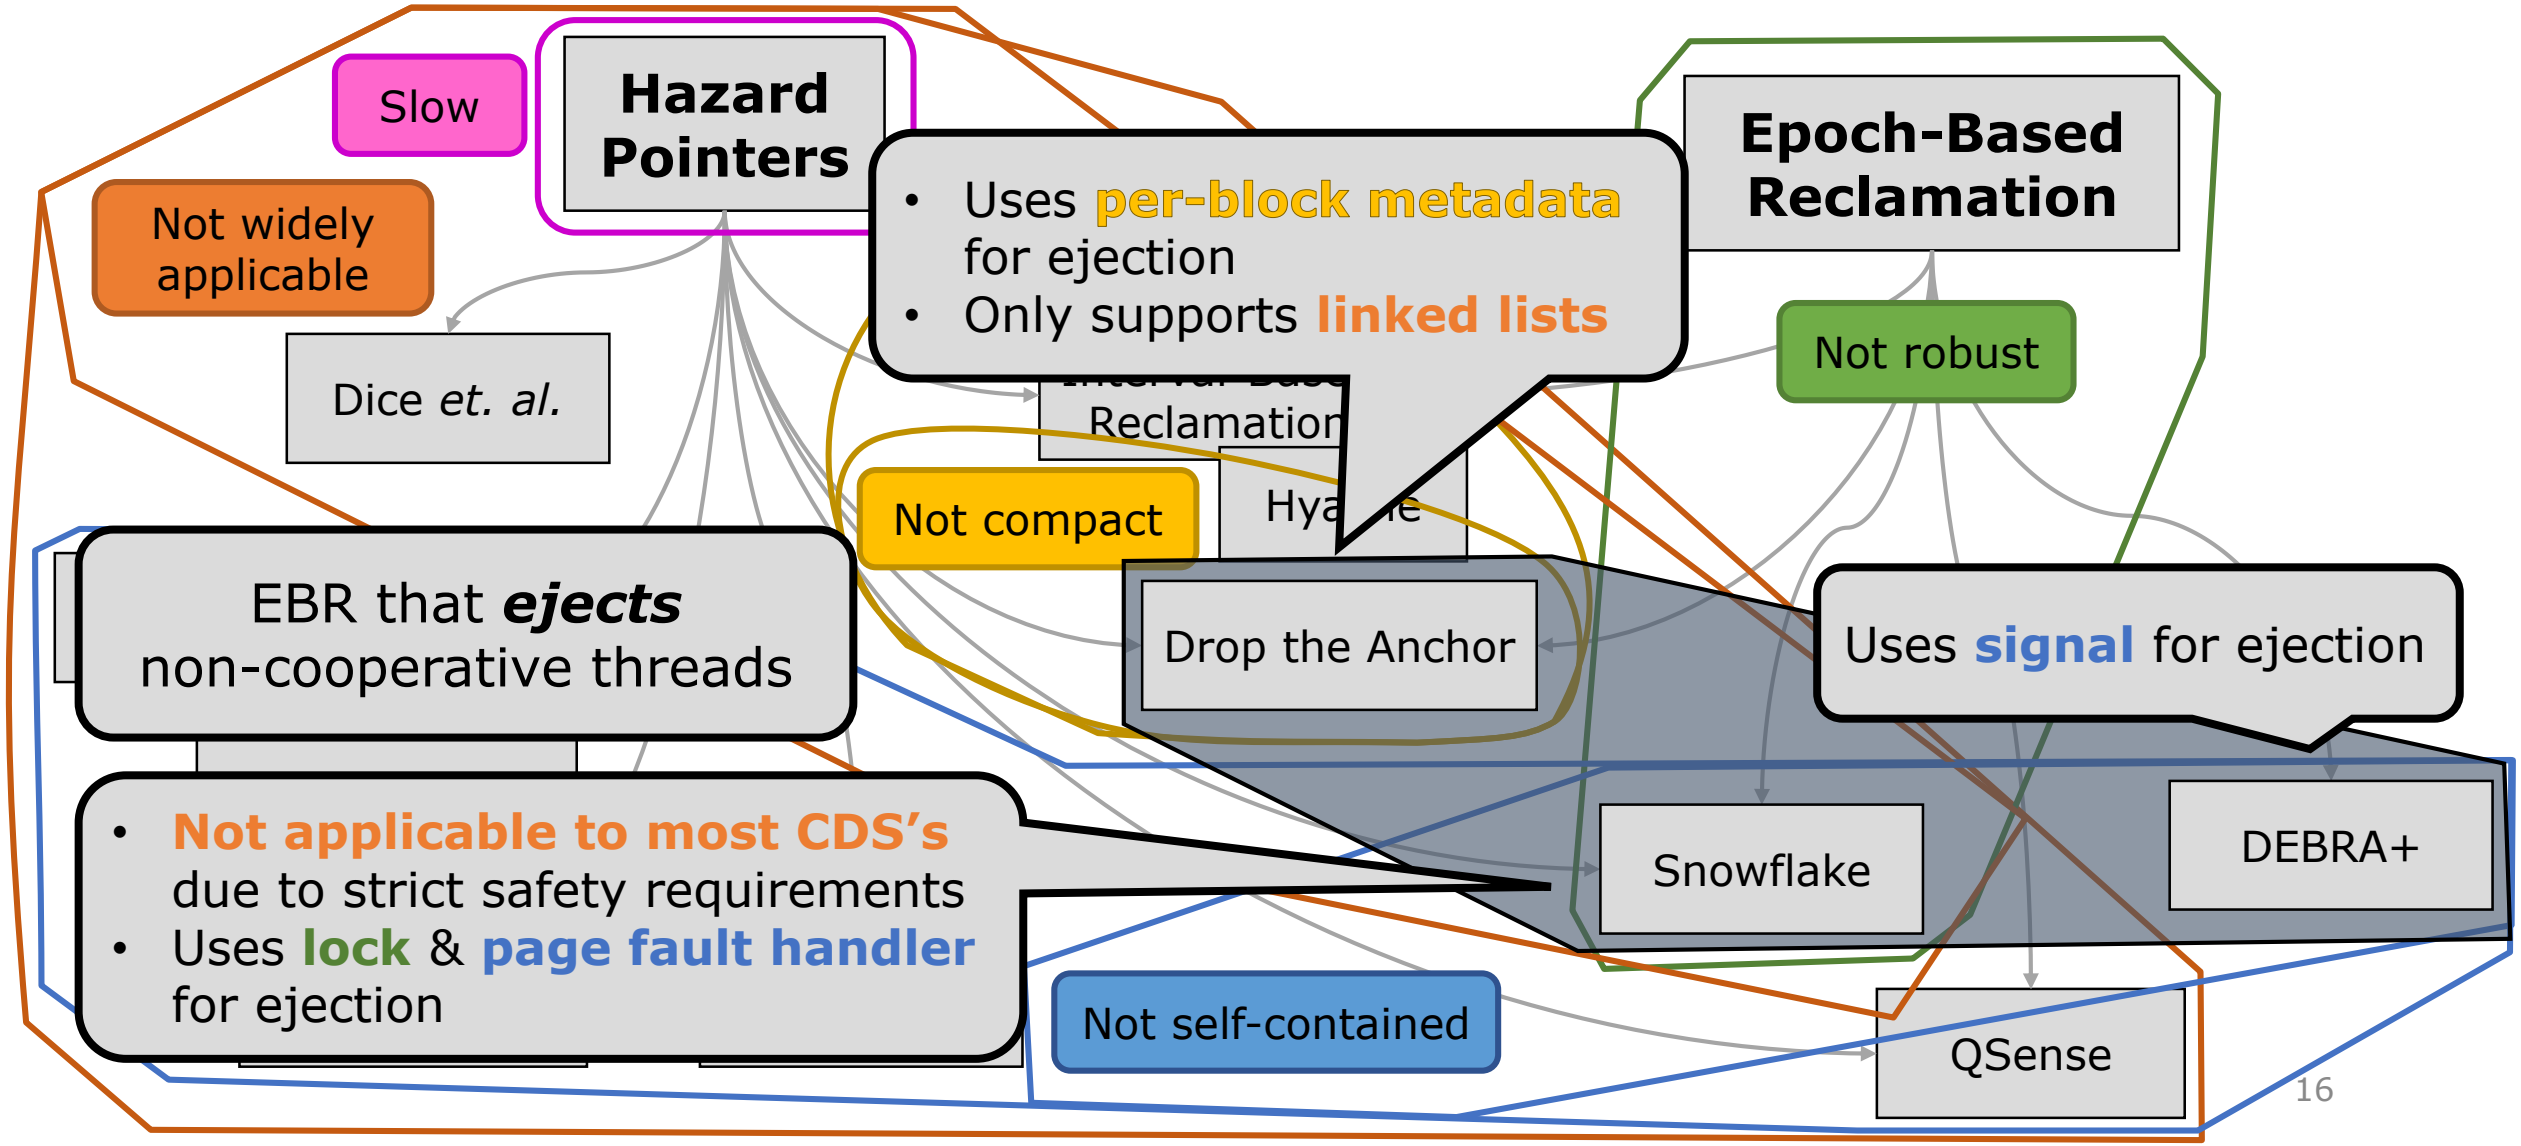

Pointer-and-Epoch-Based Reclamation

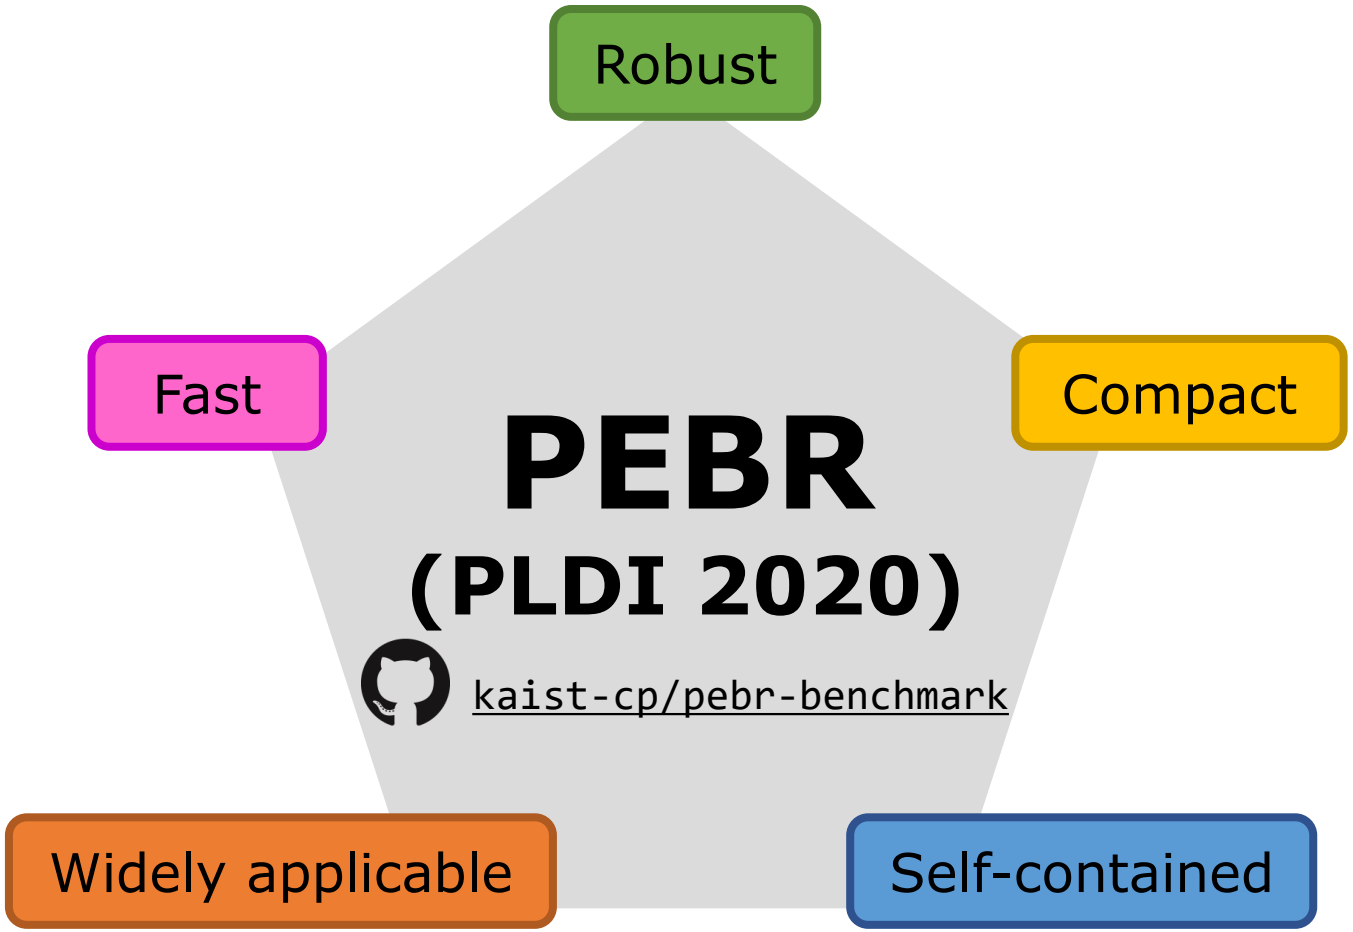

1. Hybrid of EBR and HP using **ejection**.
2. **Widely applicable** API even in the presence of ejection
3. **Robust**, **self-contained** and **compact** ejection algorithm

Wide Applicability by Ejection Notification

PEBR prevents **unsafe `protect()`** with **ejection notification**.

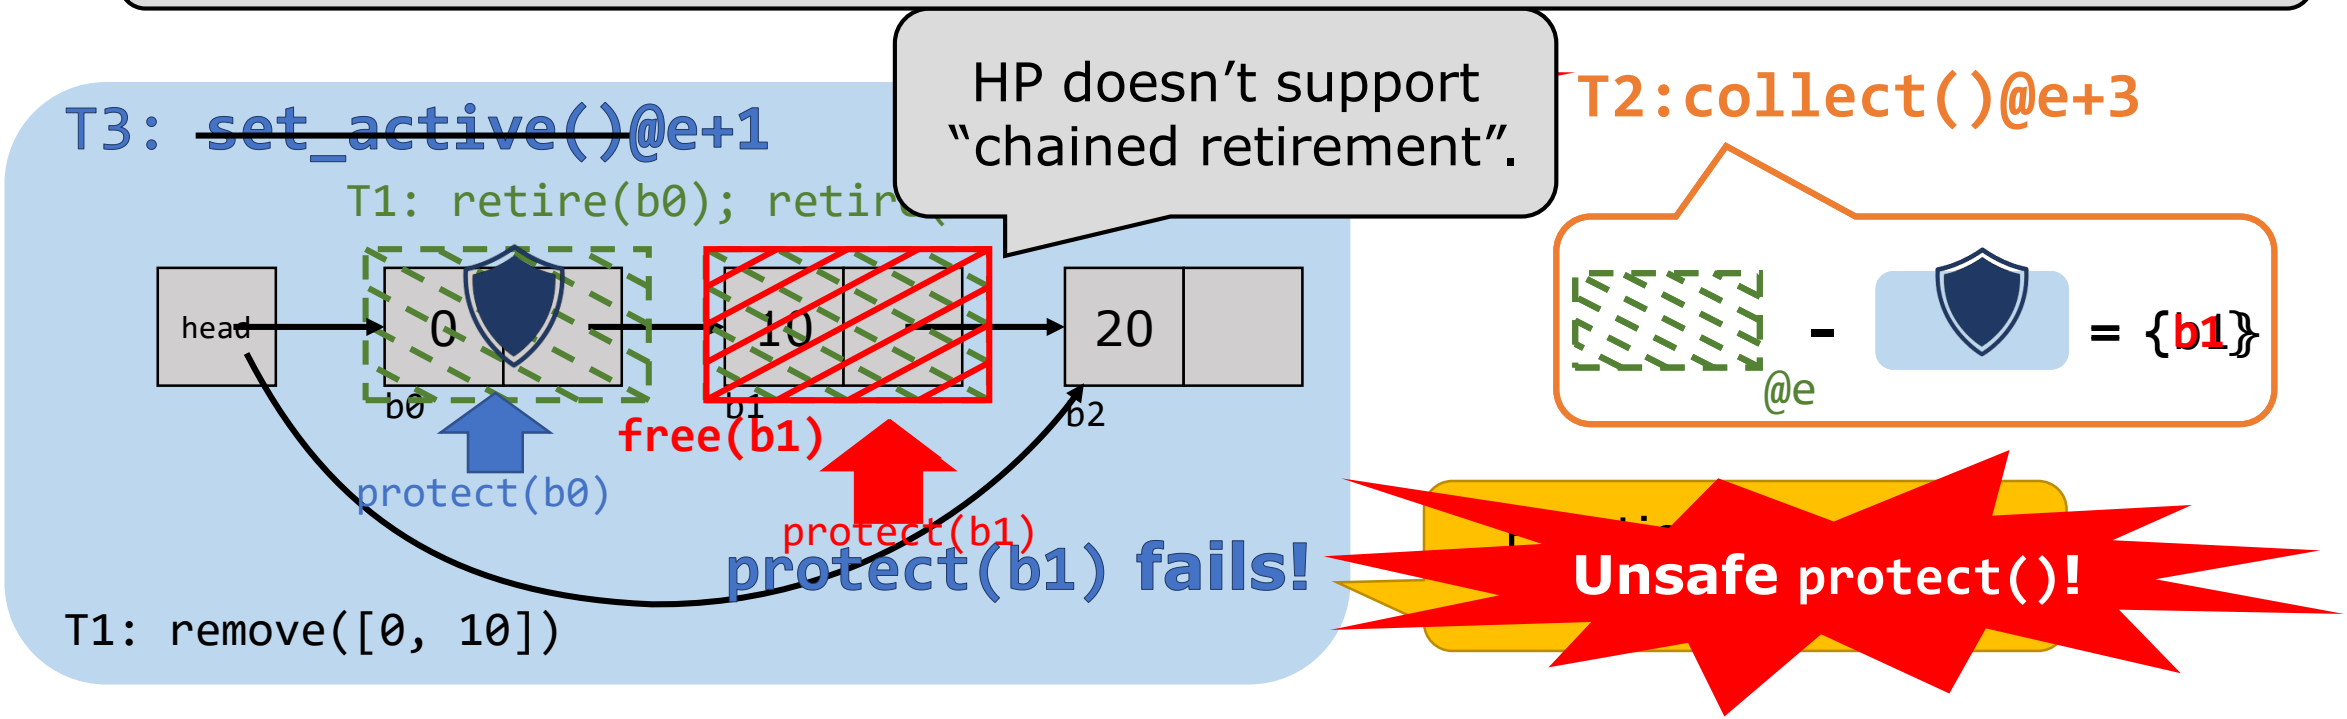

Safety Requirement:
When ejected, don't start traversal from the previously protected blocks.

Wide Applicability in Practice

Example: Porting EBR-based data structure to PEBR

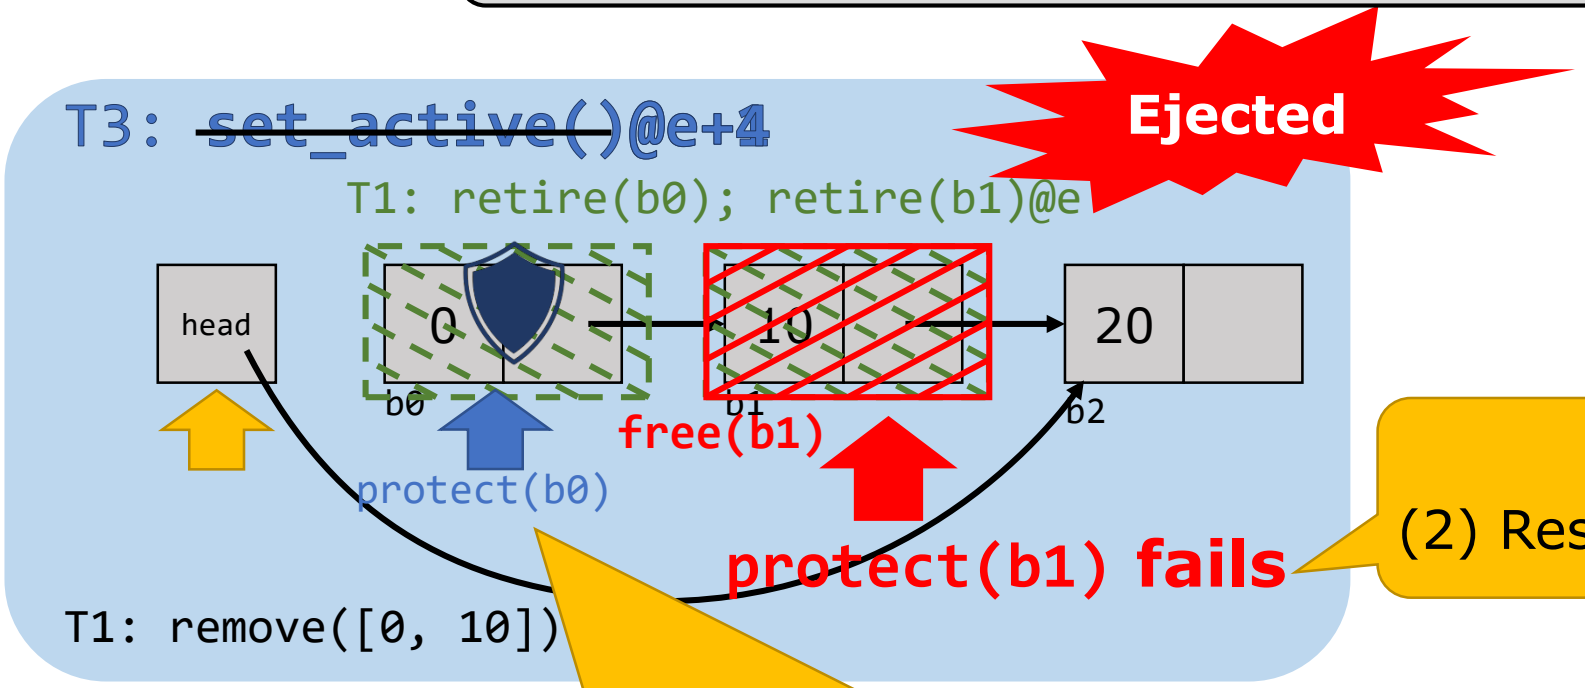

Ejection detected!
(2) Restart the operation from **head**!

(1) `protect()` before each dereference

Safety Req
When ejected, don't start traversal from the previously protected blocks.

Pointer-and-Epoch-Based Reclamation

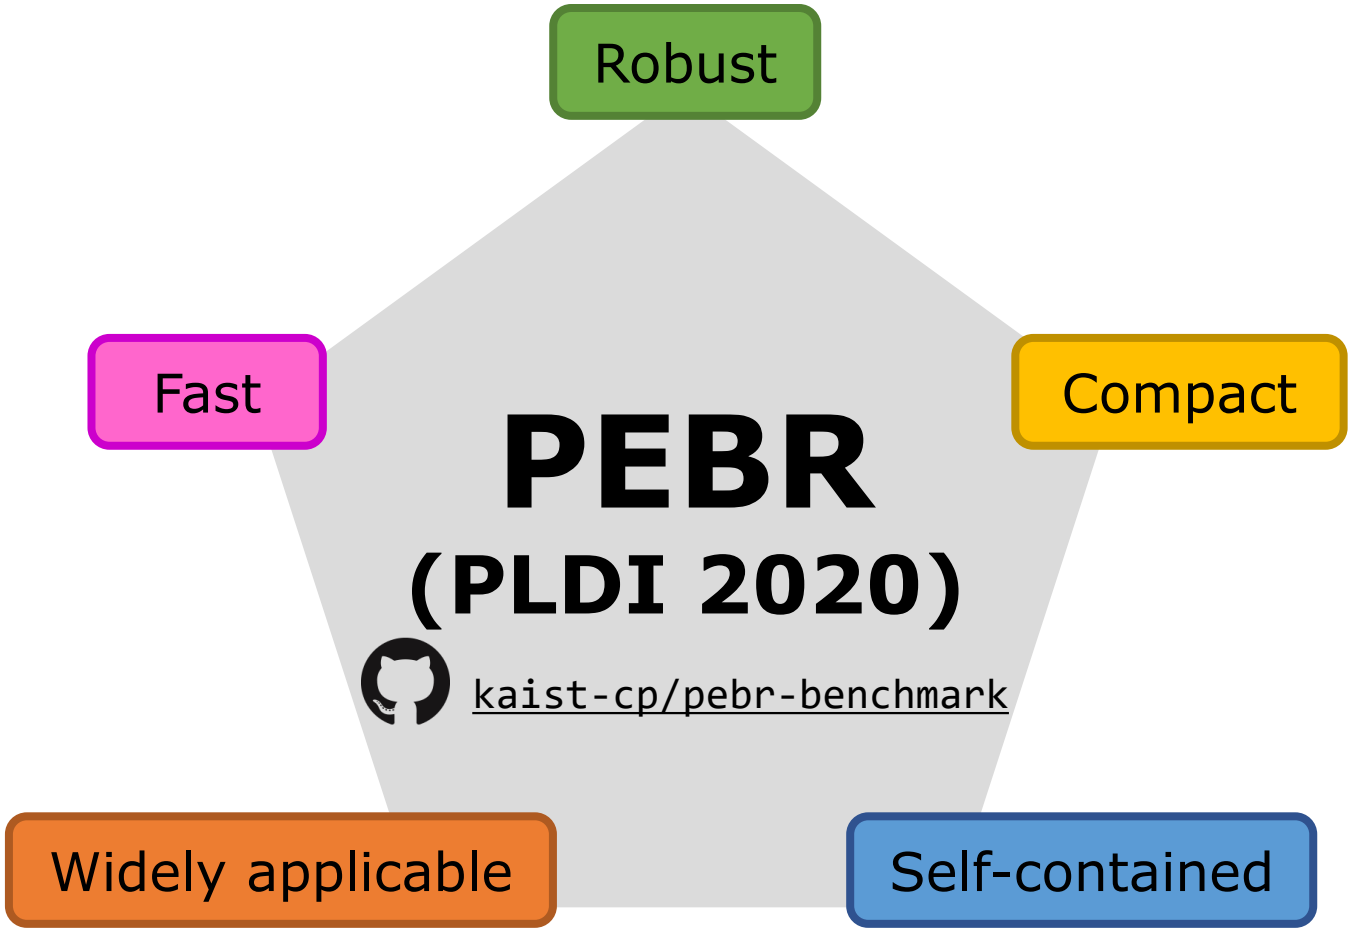

1. Hybrid of EBR and HP using **ejection**.
2. **Widely applicable** API even in the presence of ejection
3. **Robust**, **self-contained** and **compact** ejection algorithm

PEBR is Fast and Robust

Throughput

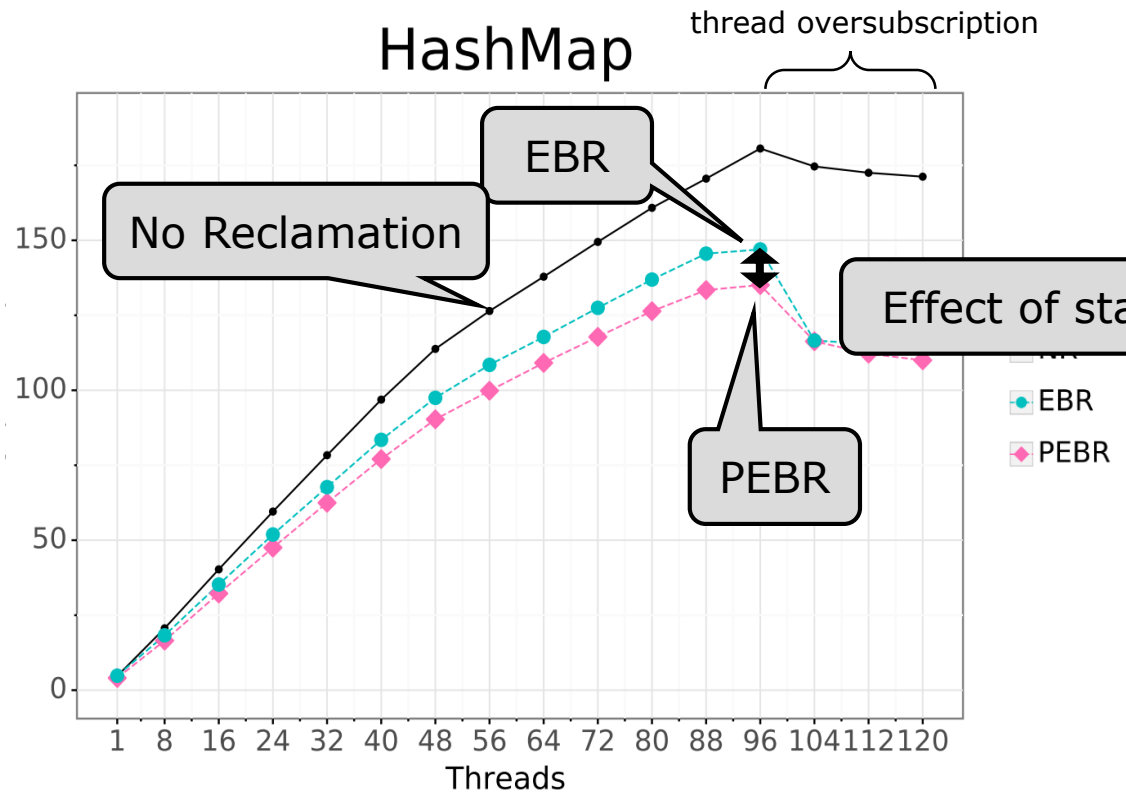

Peak Memory Usage

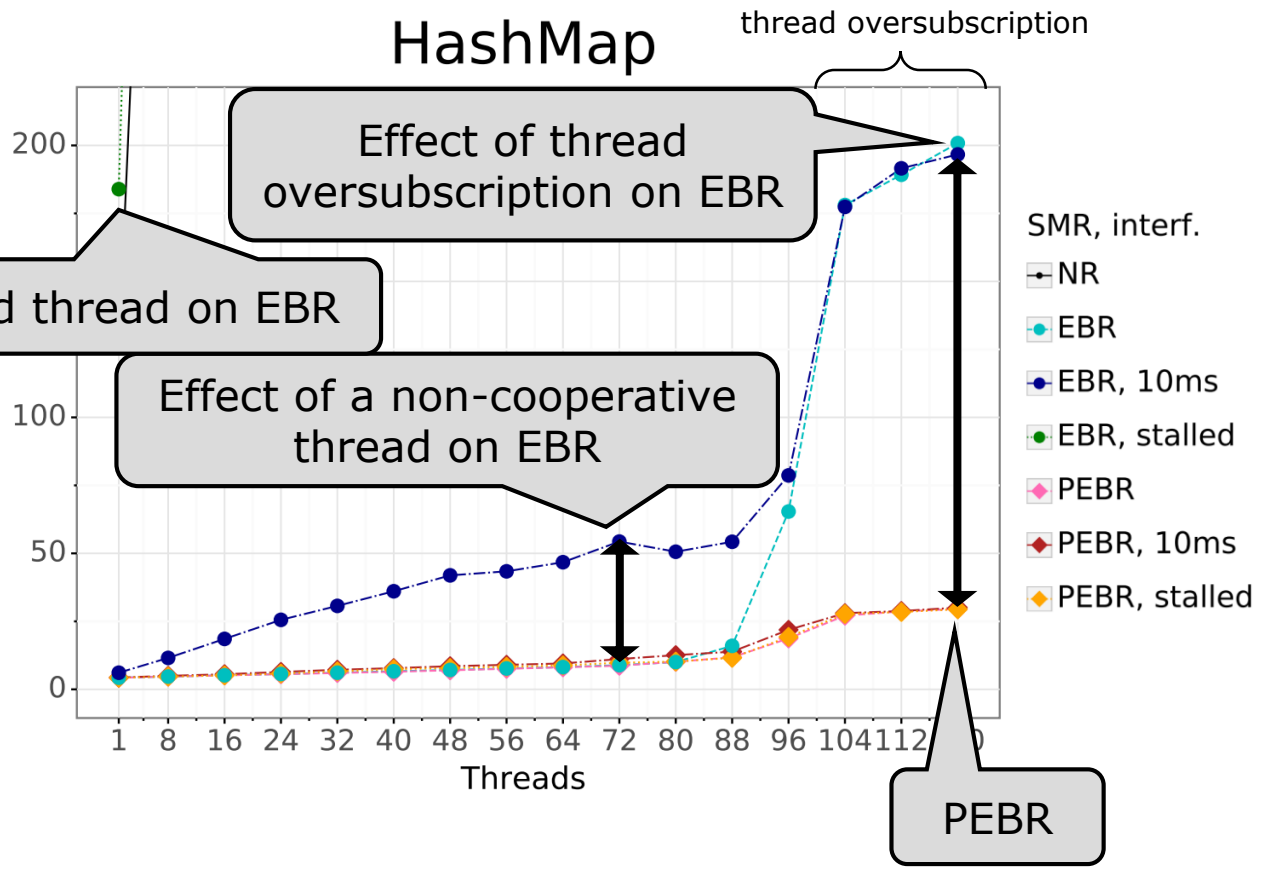

(On machine with 96 hardware threads)

What Else Is in the Paper?

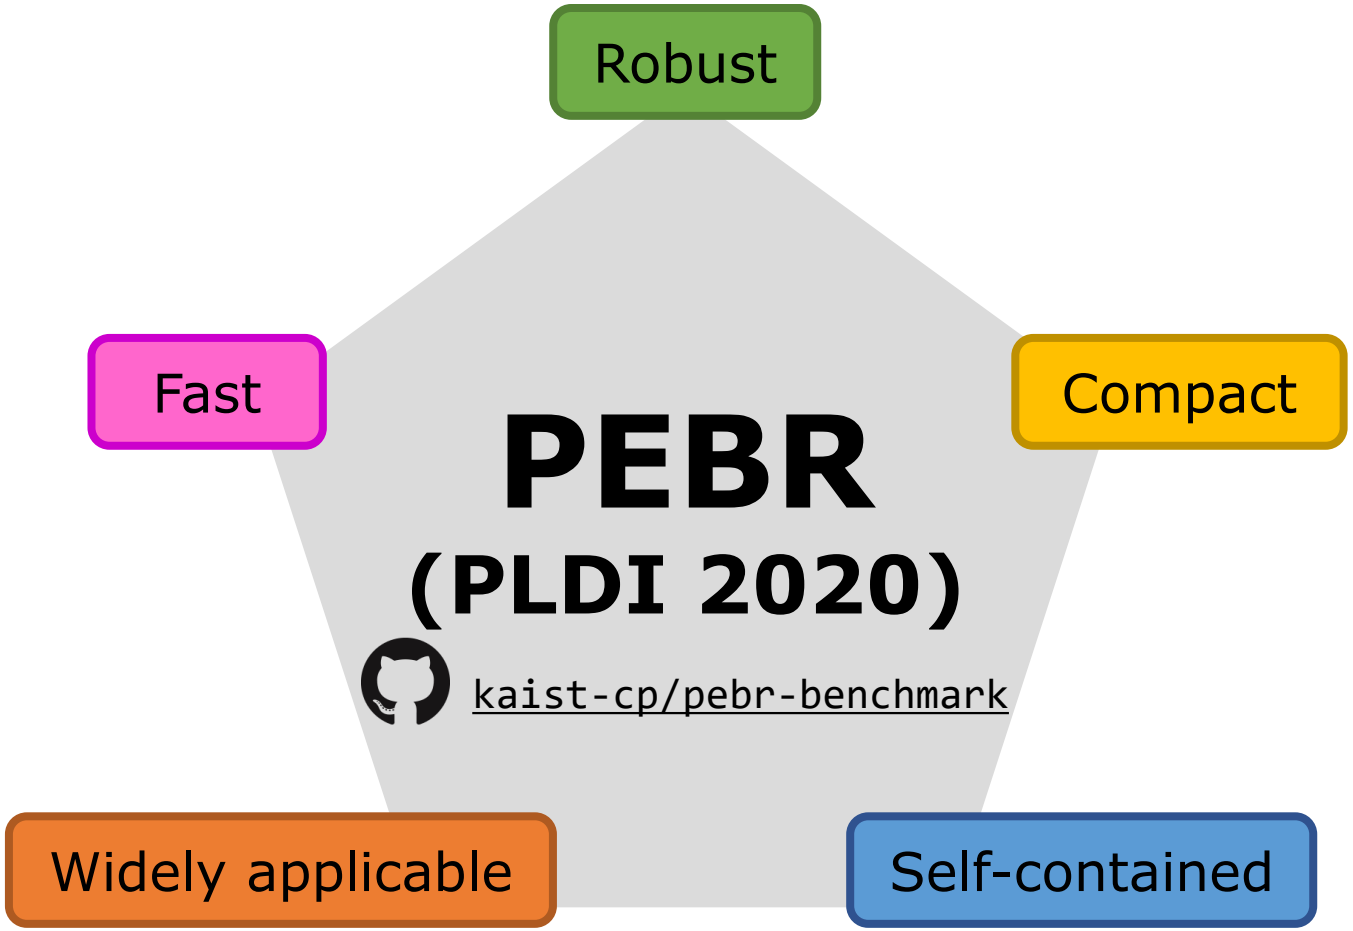

- Full algorithm
- Safety proof
- Full benchmark results
- Rust API that statically enforces some safety requirements