

Correct and Fast Persistency Guarantees

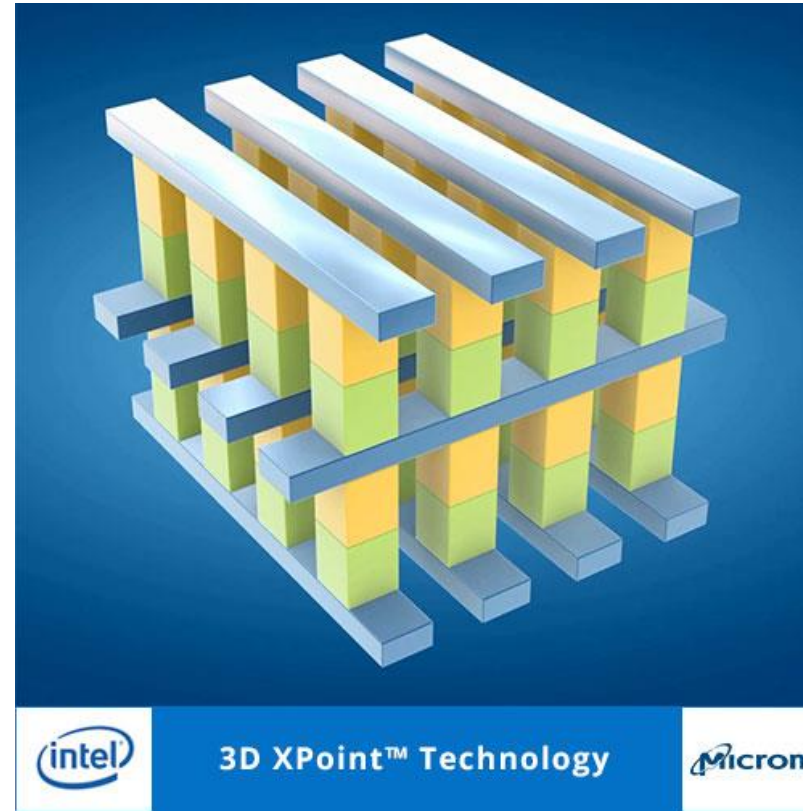
Sara Mahdizadeh Shahri
Aasheesh Kolli



PennState

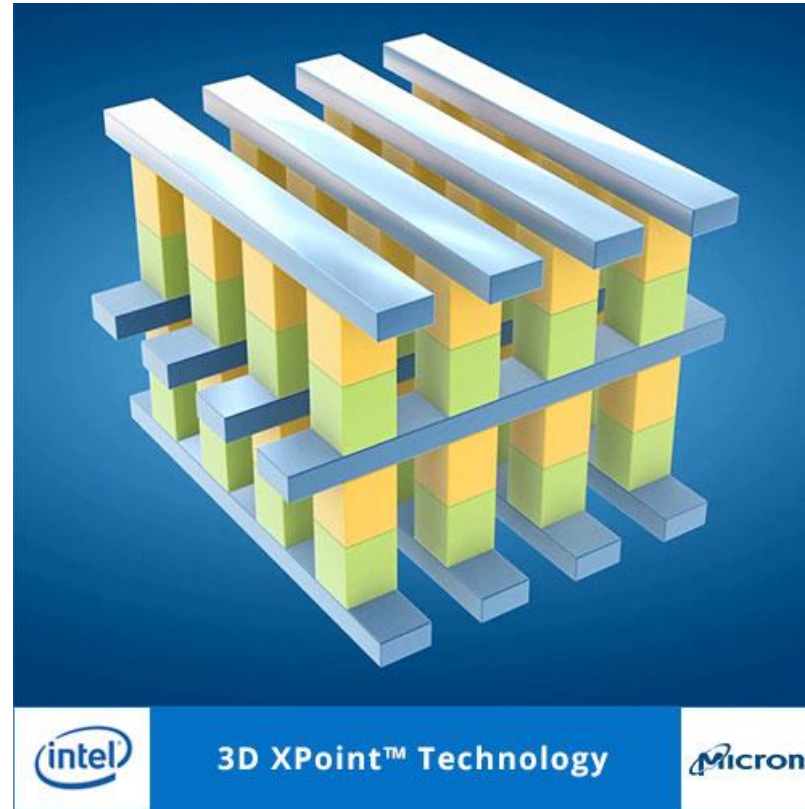
YArch Workshop - February 2019

Why Persistent Memory?



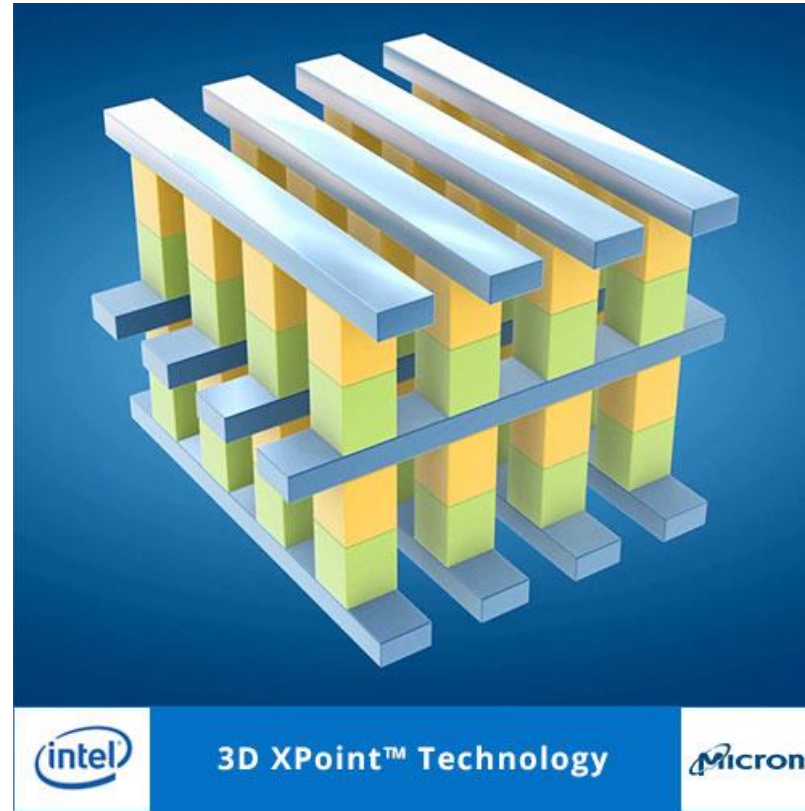
Why Persistent Memory?

Non-volatility



Why Persistent Memory?

Non-volatility

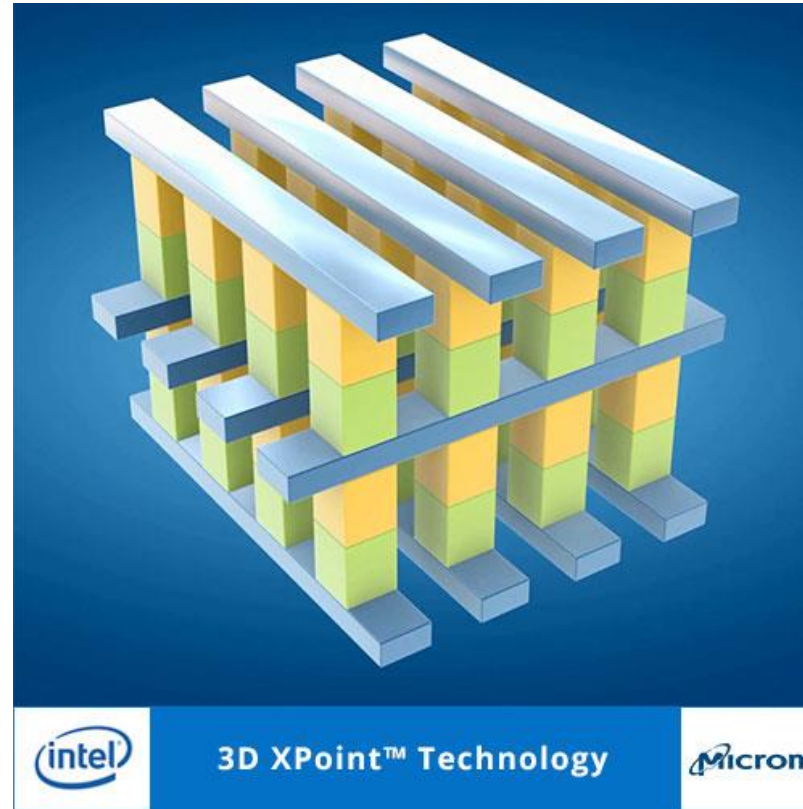


Density

Why Persistent Memory?

Byte-addressability

Non-volatility

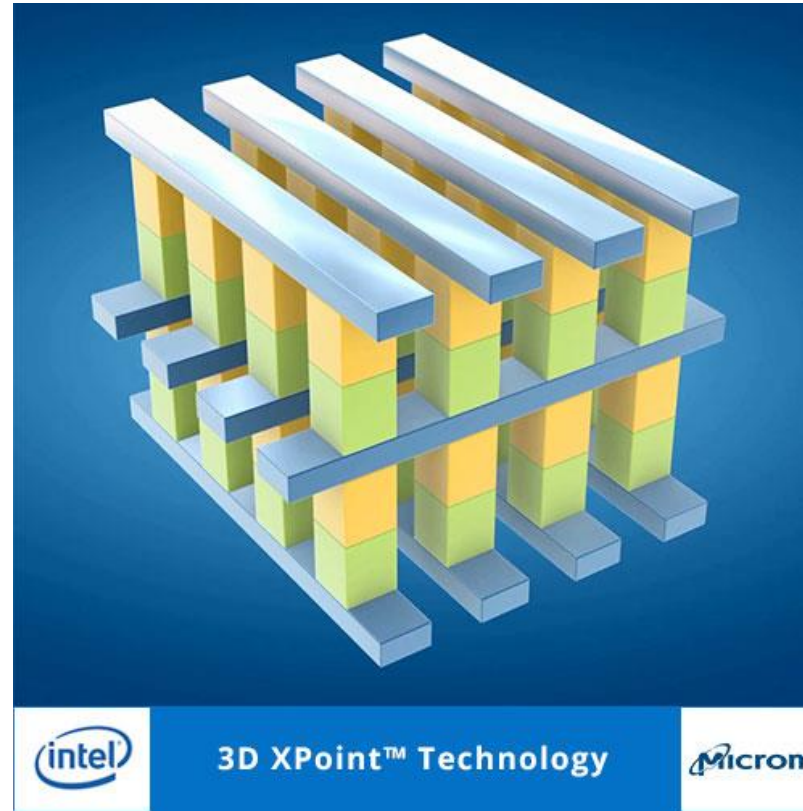


Density

Why Persistent Memory?

Byte-addressability

Non-volatility

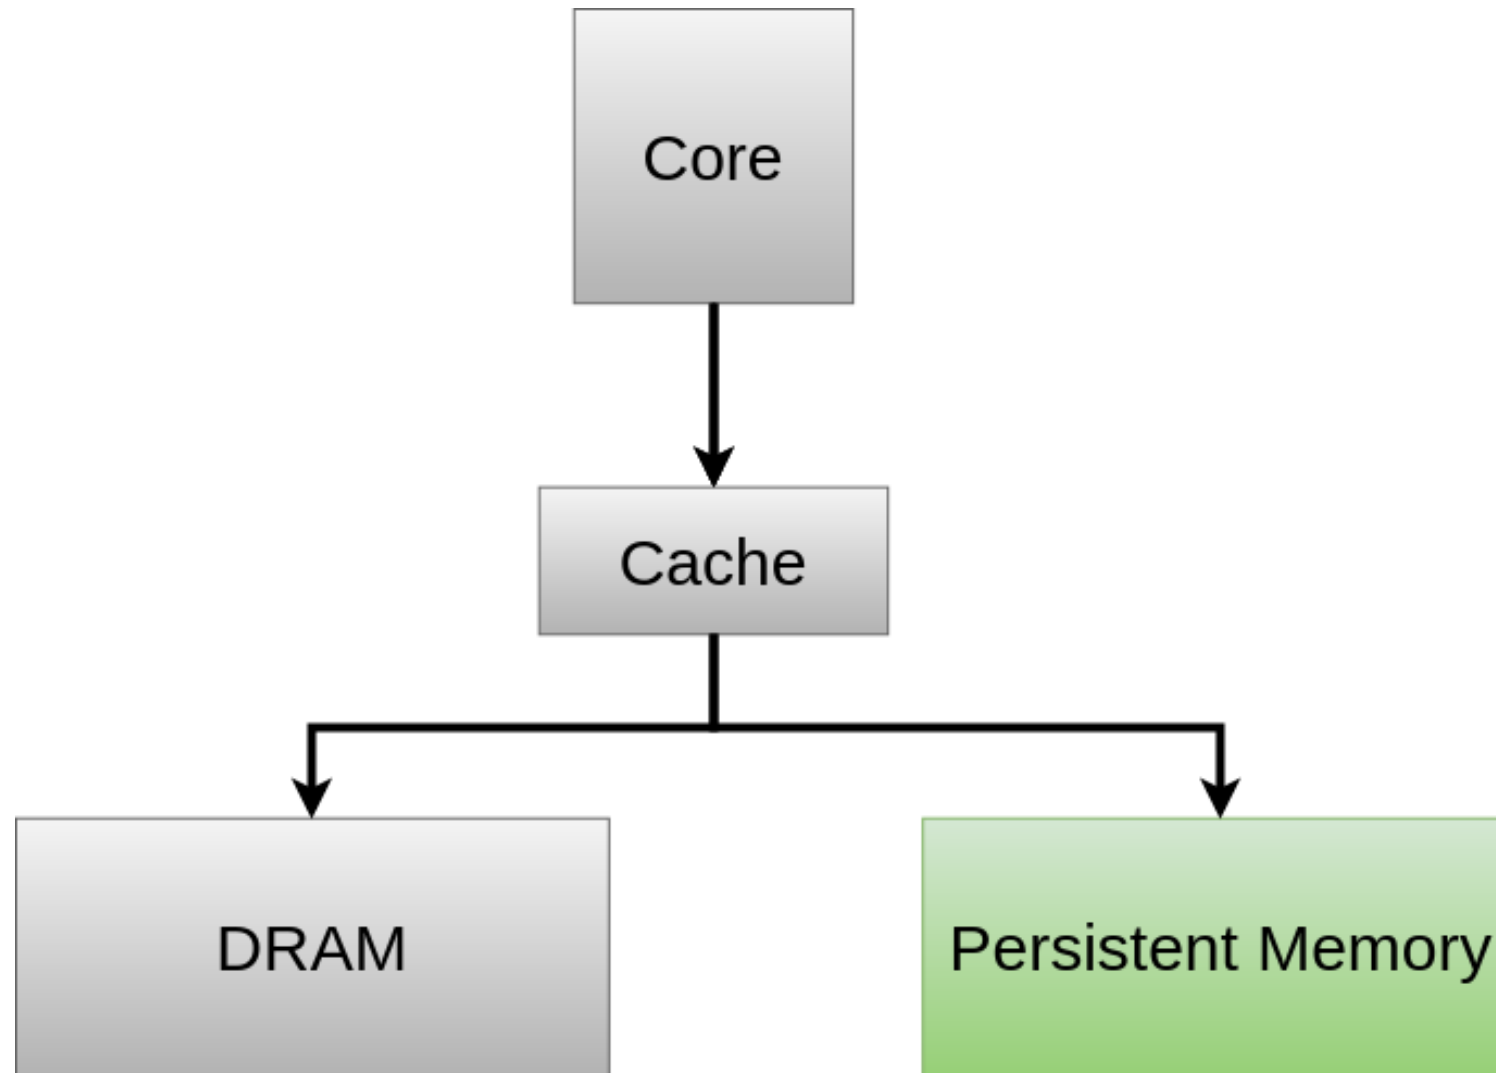


Performance

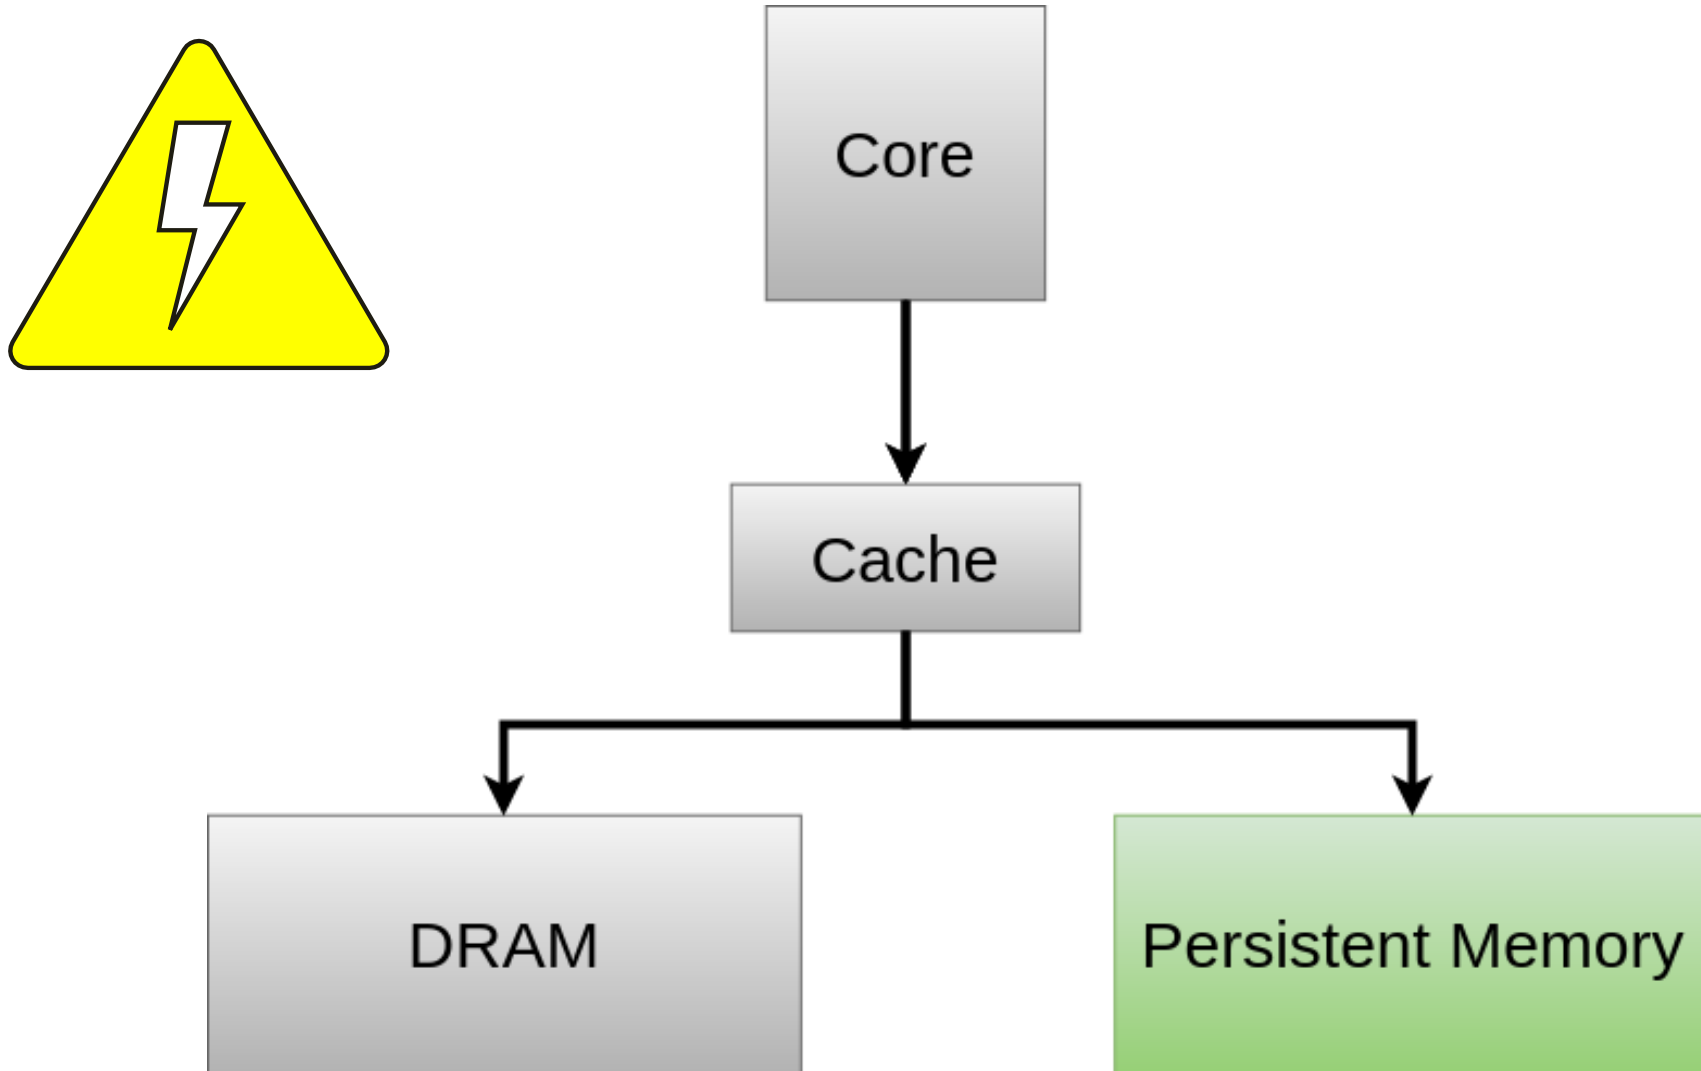


Density

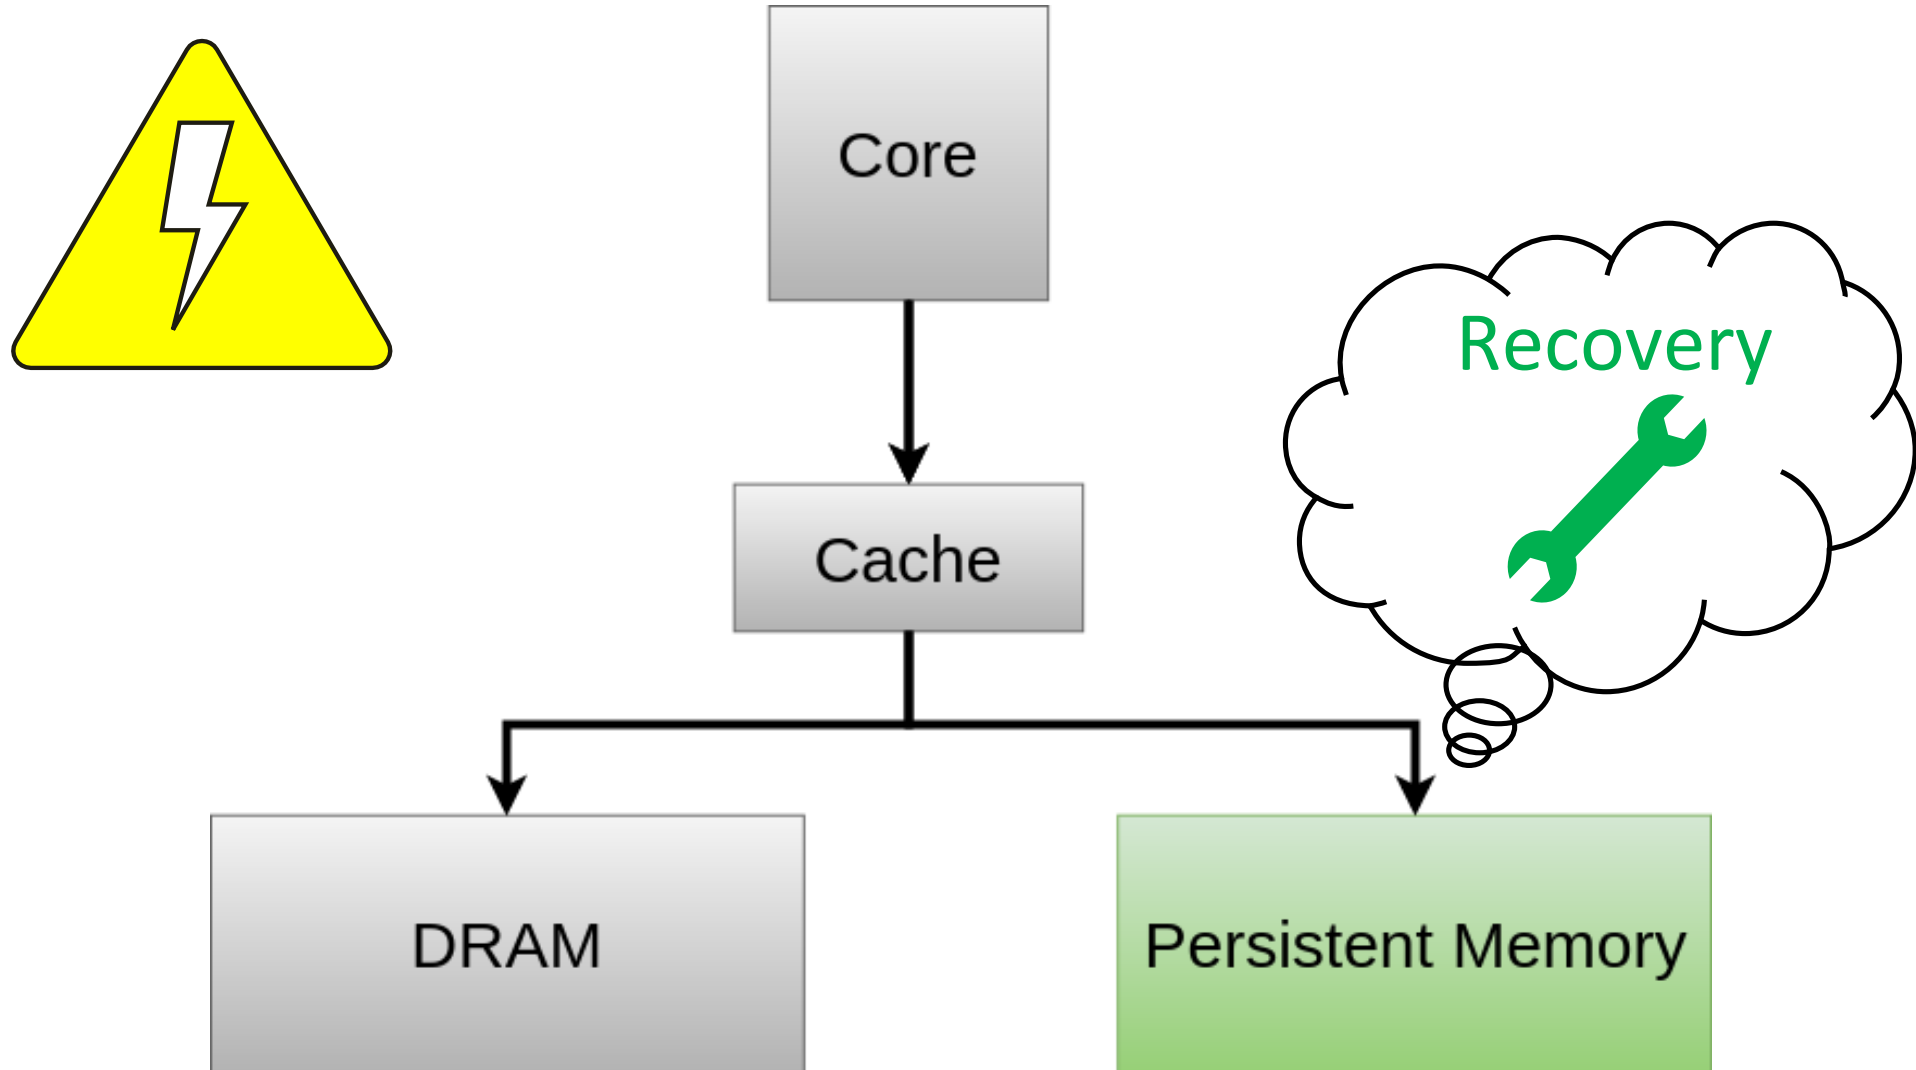
Memory Hierarchy



Memory Hierarchy

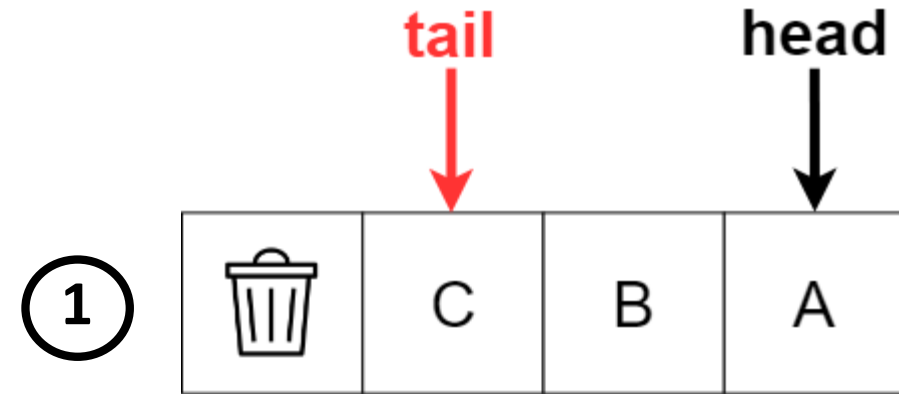


Memory Hierarchy



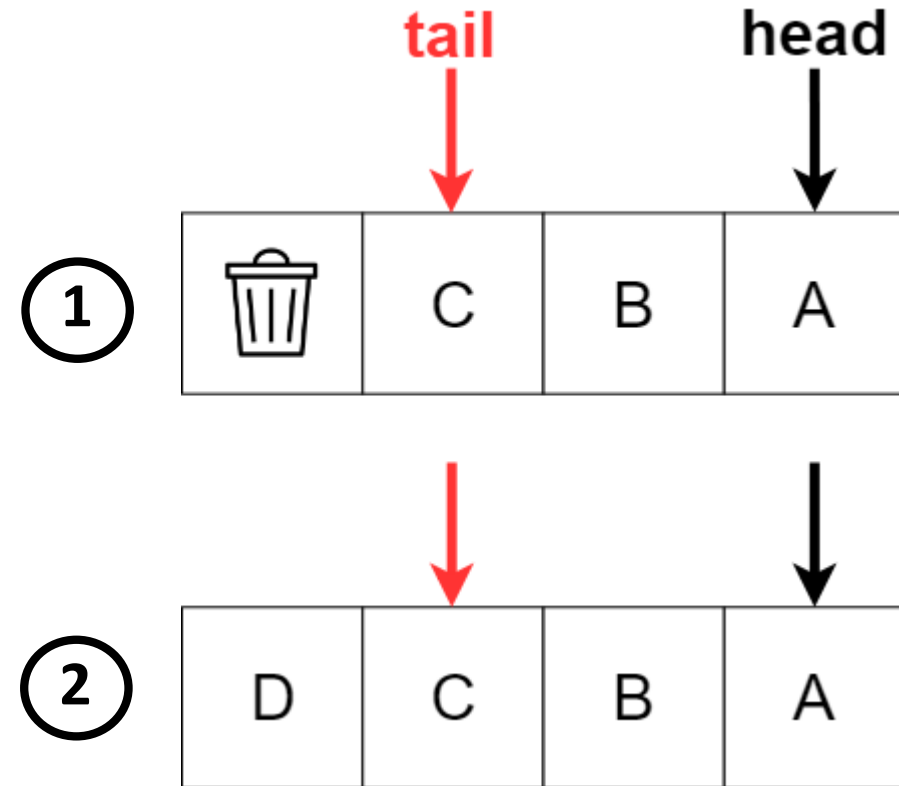
Correct Recovery

Use Case: Insertion in Queue



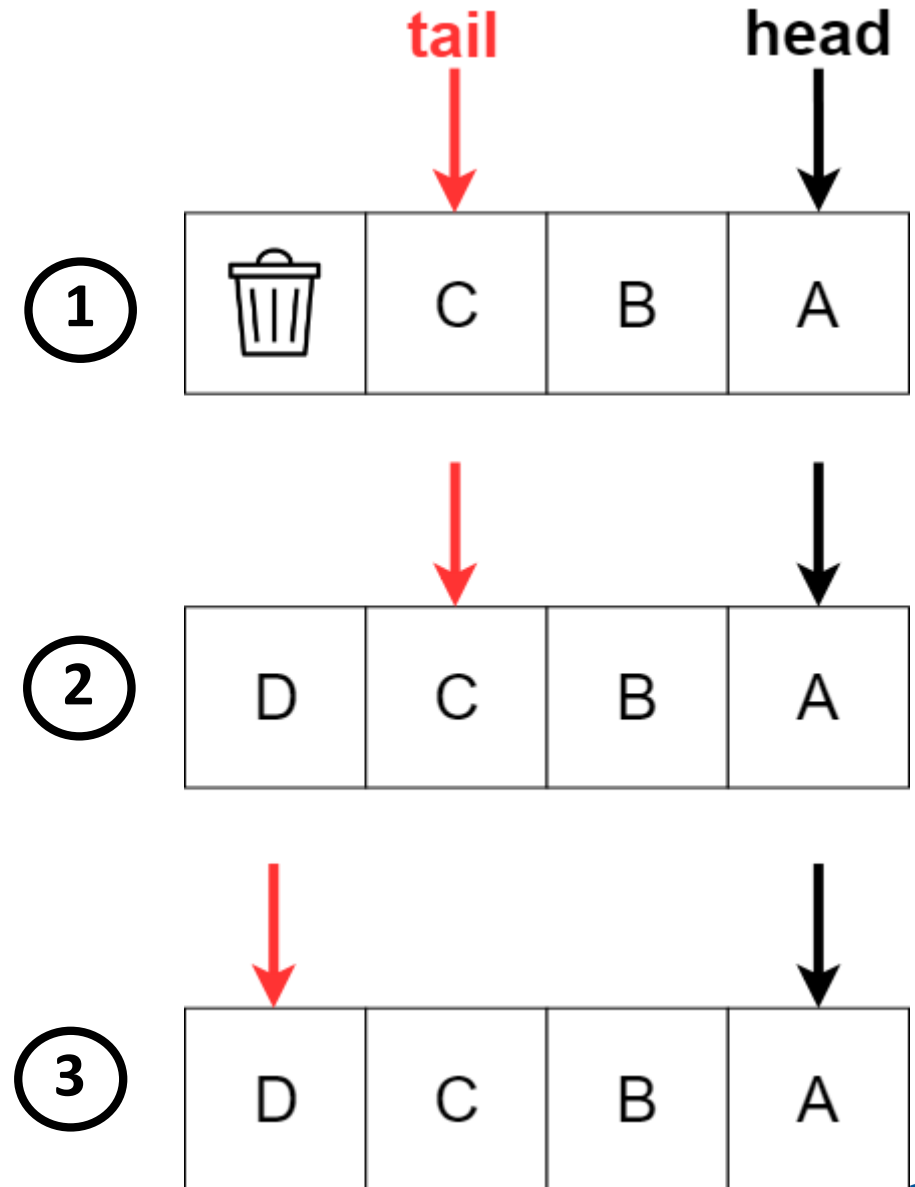
Correct Recovery

Use Case: Insertion in Queue



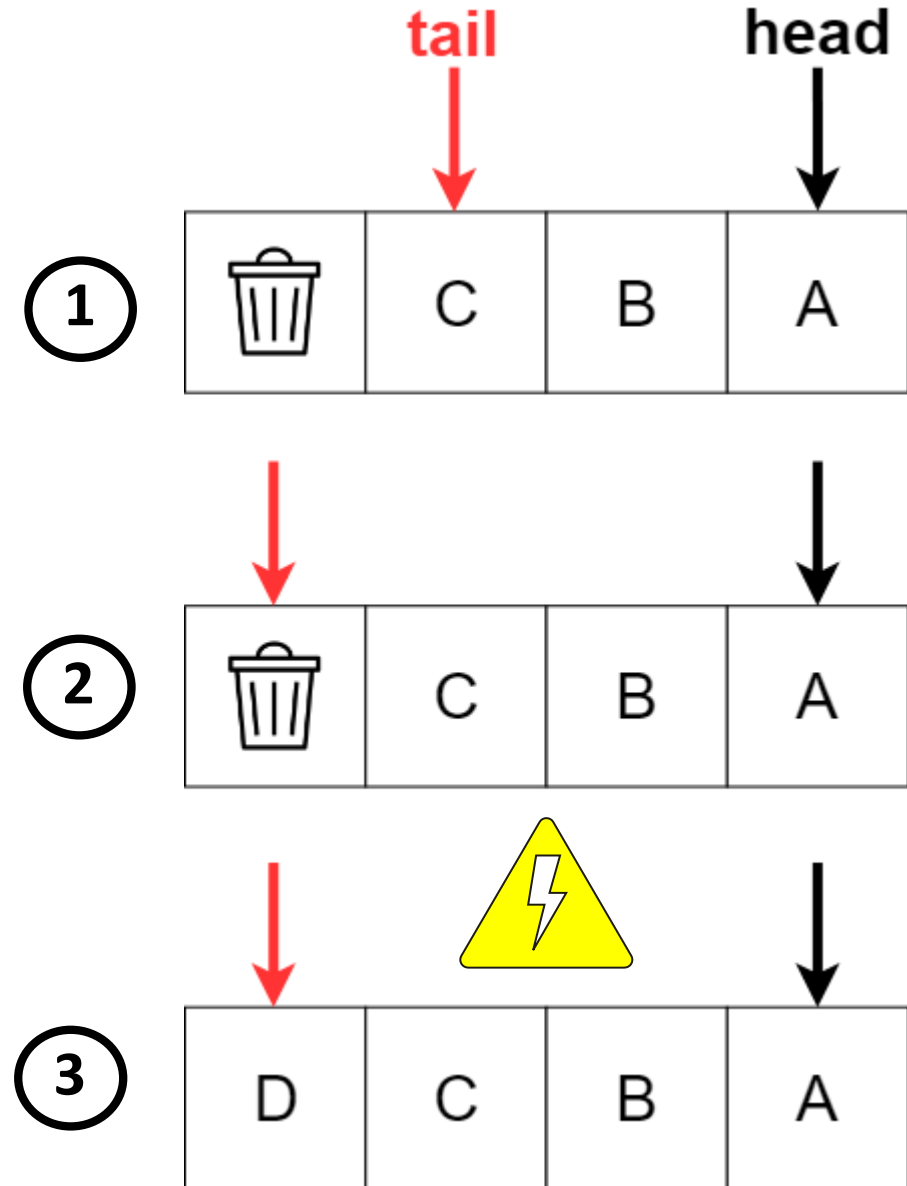
Correct Recovery

Use Case: Insertion in Queue



Correct Recovery

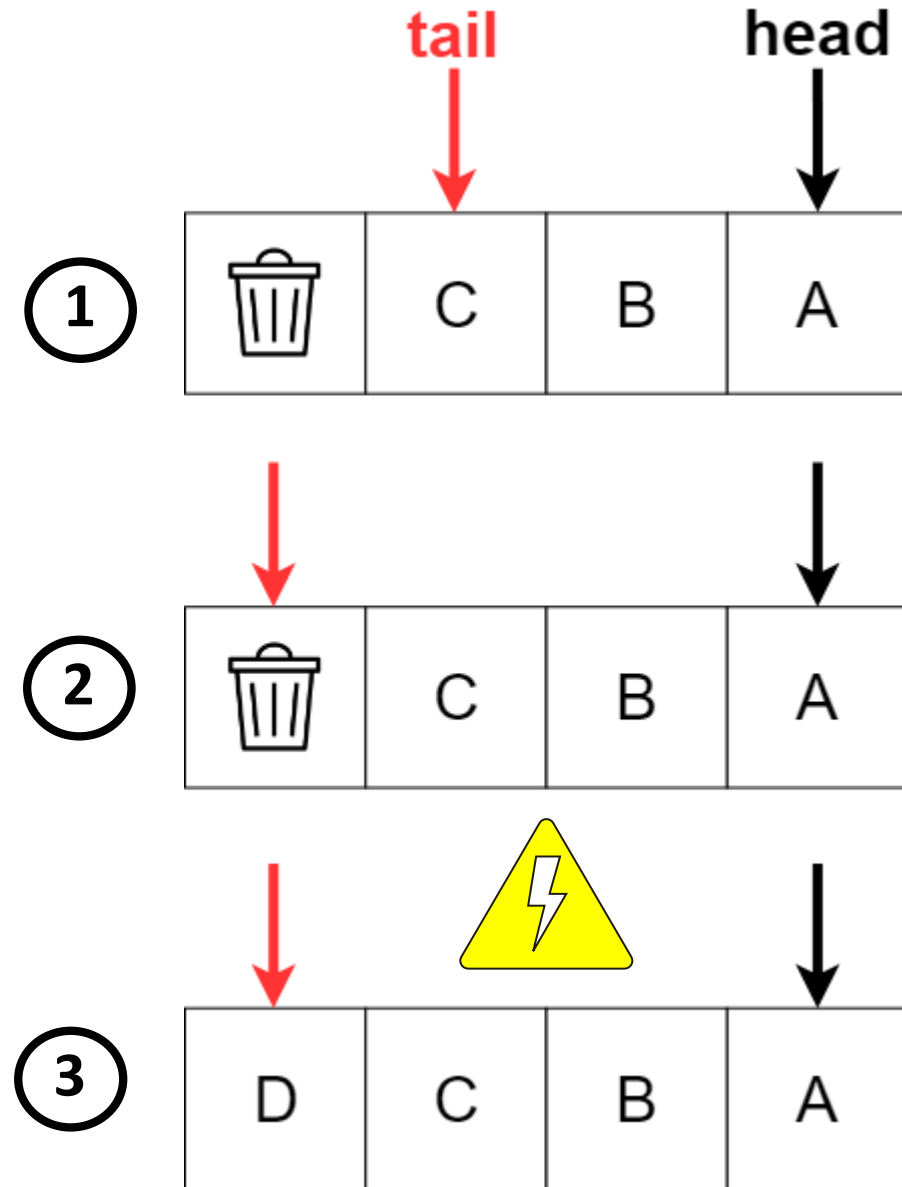
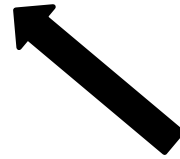
Use Case: Insertion in Queue



Correct Recovery

Use Case: Insertion in Queue

Order of Updates

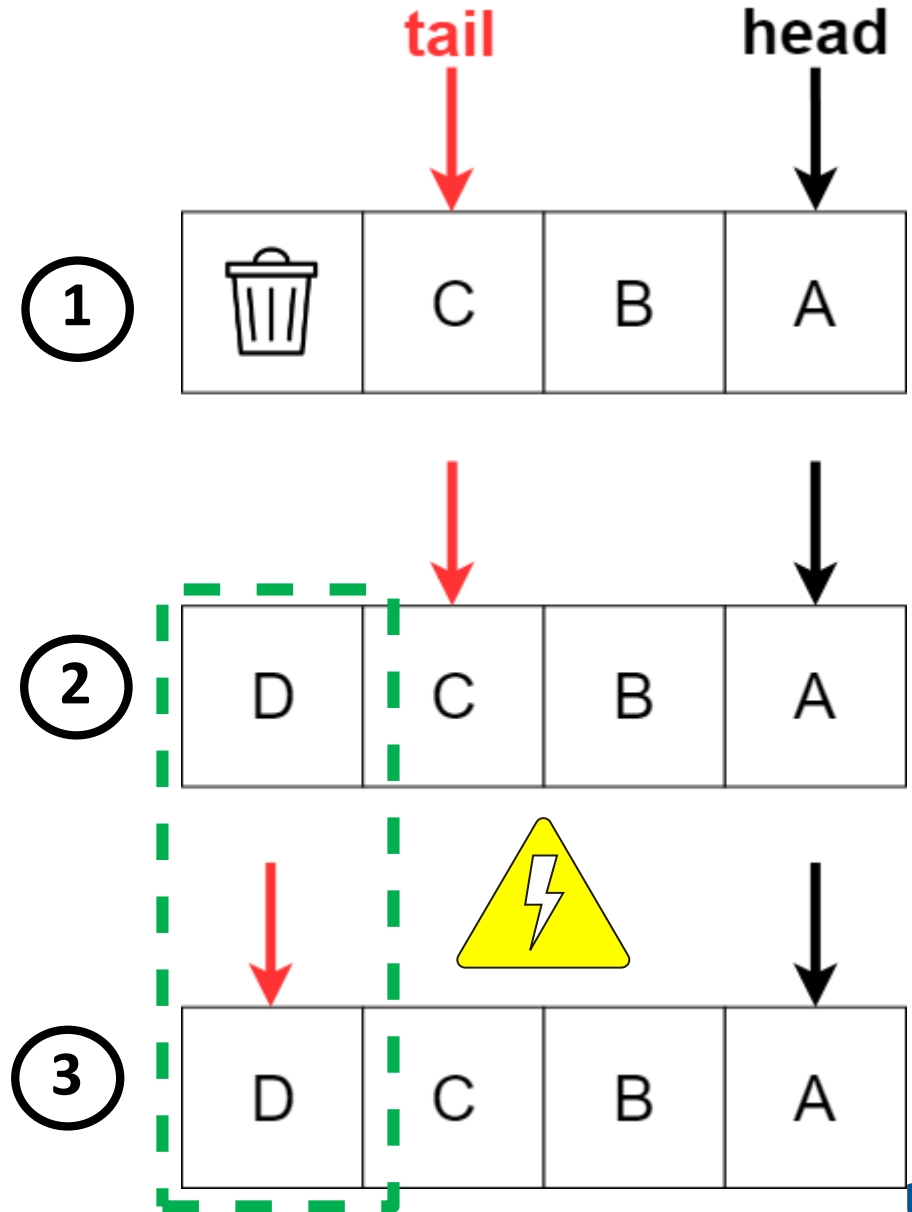


Correct Recovery

Use Case: Insertion in Queue

Order of Updates

Failure Atomicity

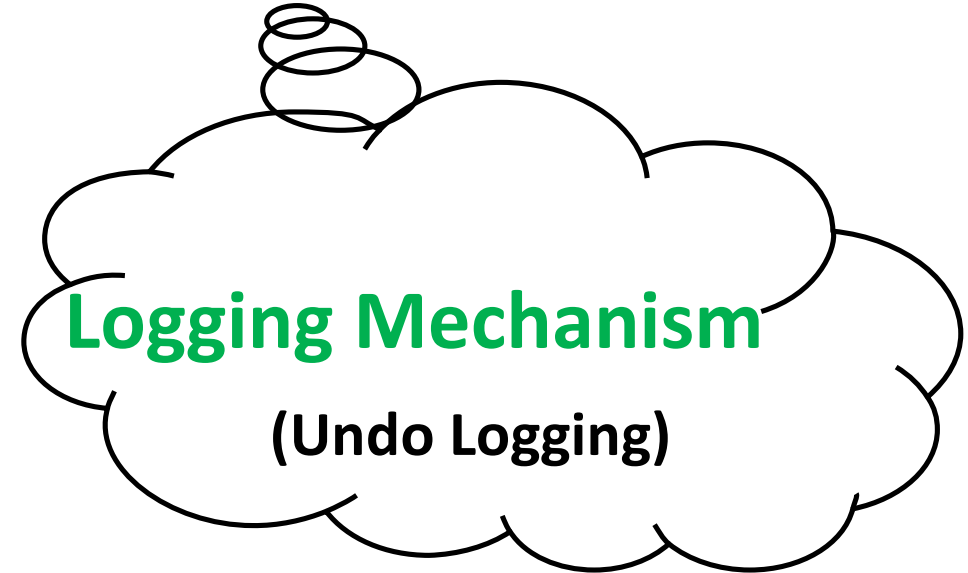


Failure Atomicity Semantics

- Providing Failure Atomicity for Granularity Beyond Persist Access

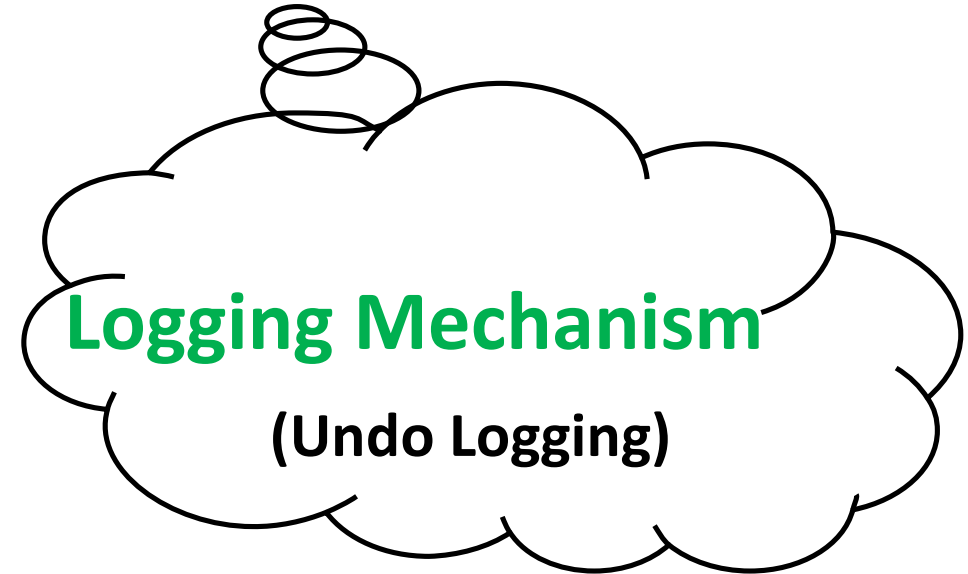
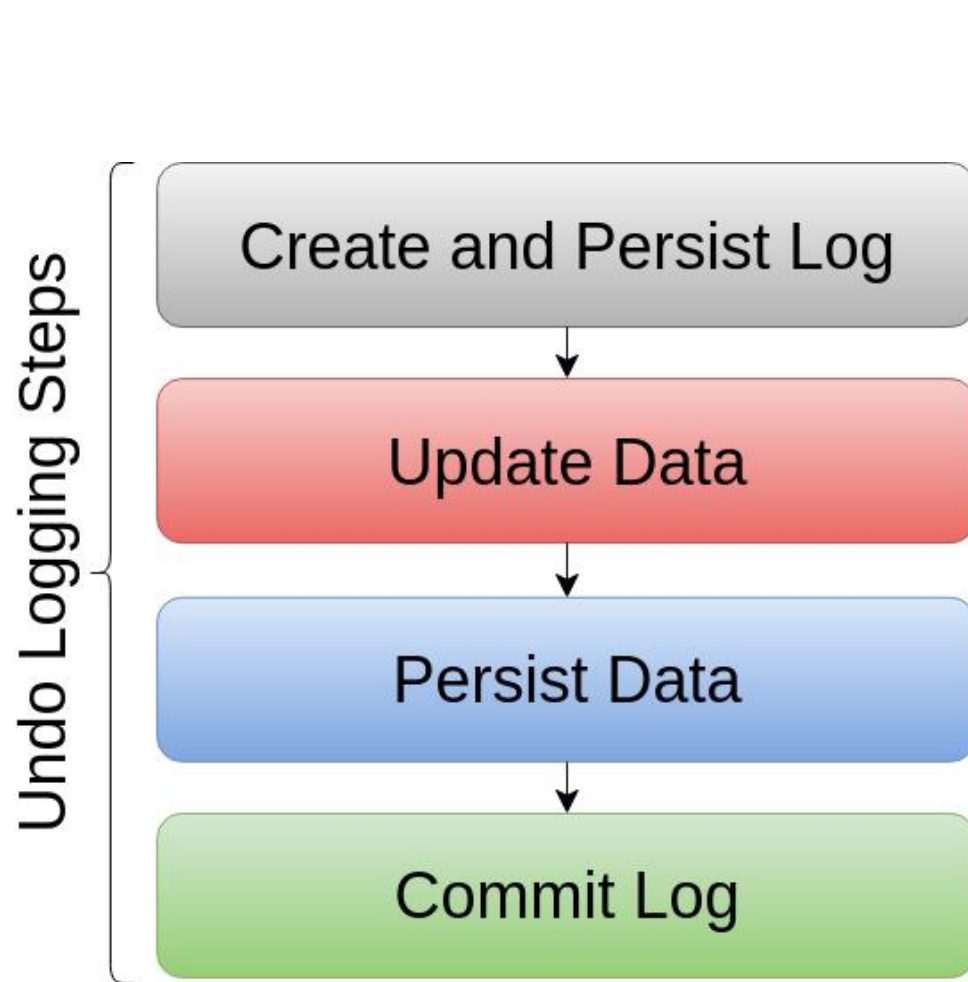
Failure Atomicity Semantics

- Providing Failure Atomicity for Granularity Beyond Persist Access



Failure Atomicity Semantics

- Providing Failure Atomicity for Granularity Beyond Persist Access



Semantics for Delivering Persistency

Challenge #1: Correctness

Prior Work

Atlas [Chakrabarti' 14]
Coupled/Decoupled SFR [Gogte' 18]

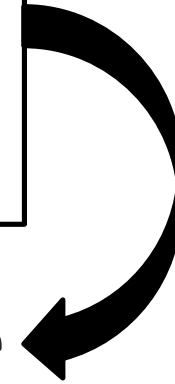
Semantics for Delivering Persistency

Challenge #1: Correctness

Prior Work

Atlas [Chakrabarti' 14]
Coupled/Decoupled SFR [Gogte' 18]

Detect Persist Updates at **Compile Time**



Semantics for Delivering Persistency

Challenge #1: Correctness

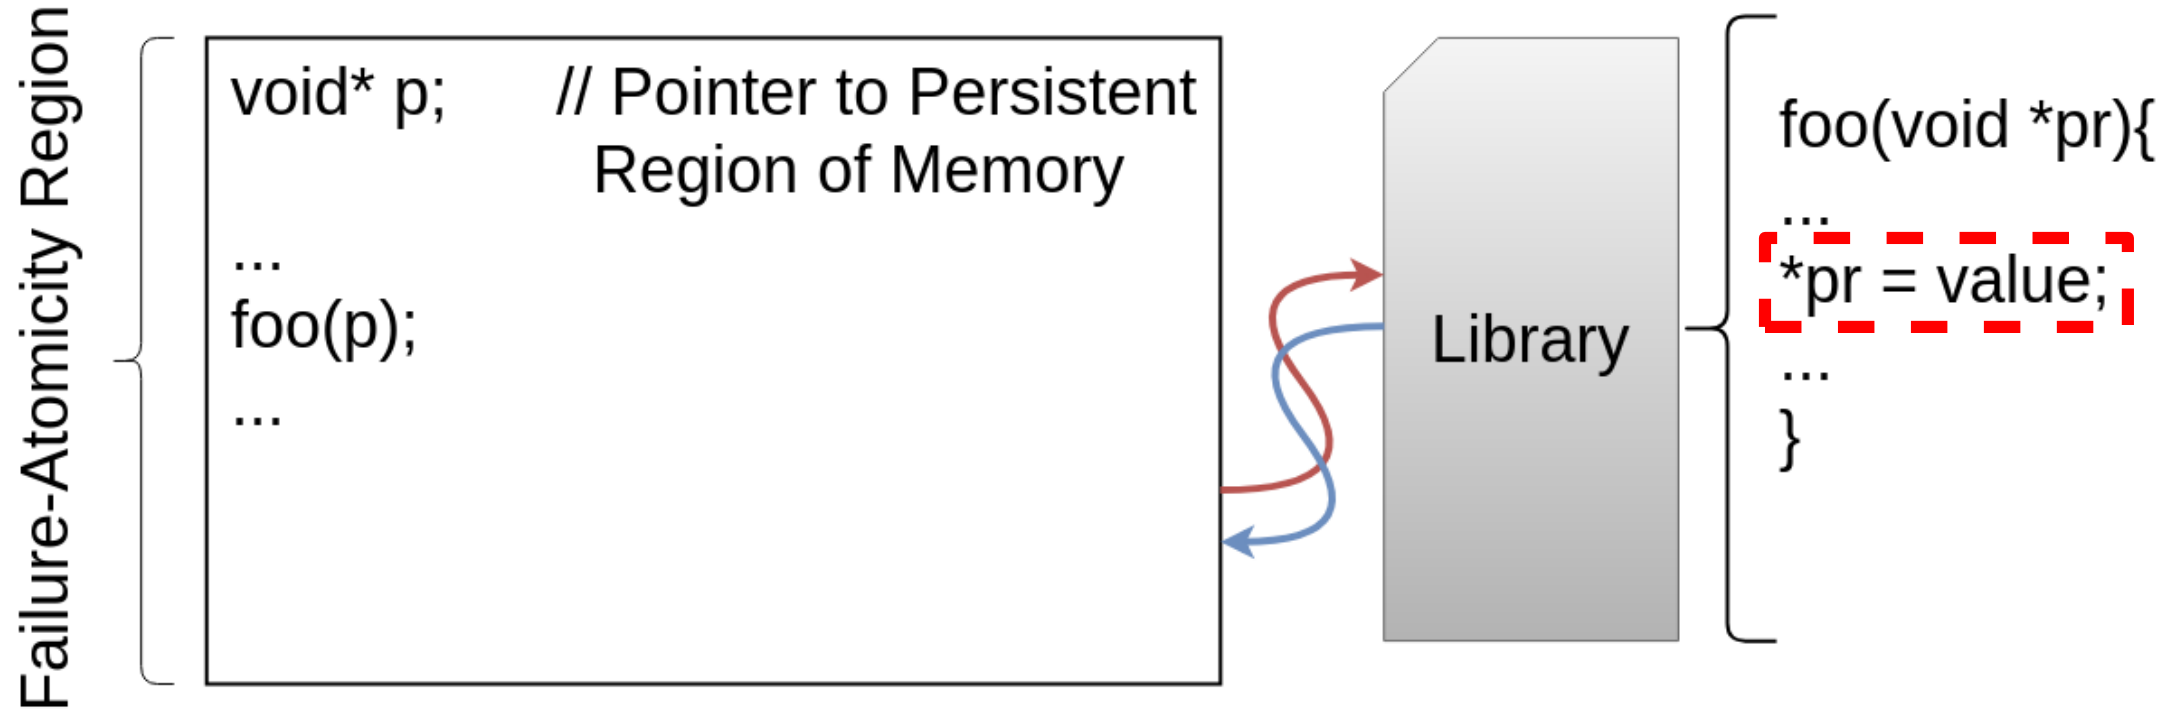
Prior Work

Atlas [Chakrabarti' 14]
Coupled/Decoupled SFR [Gogte' 18]

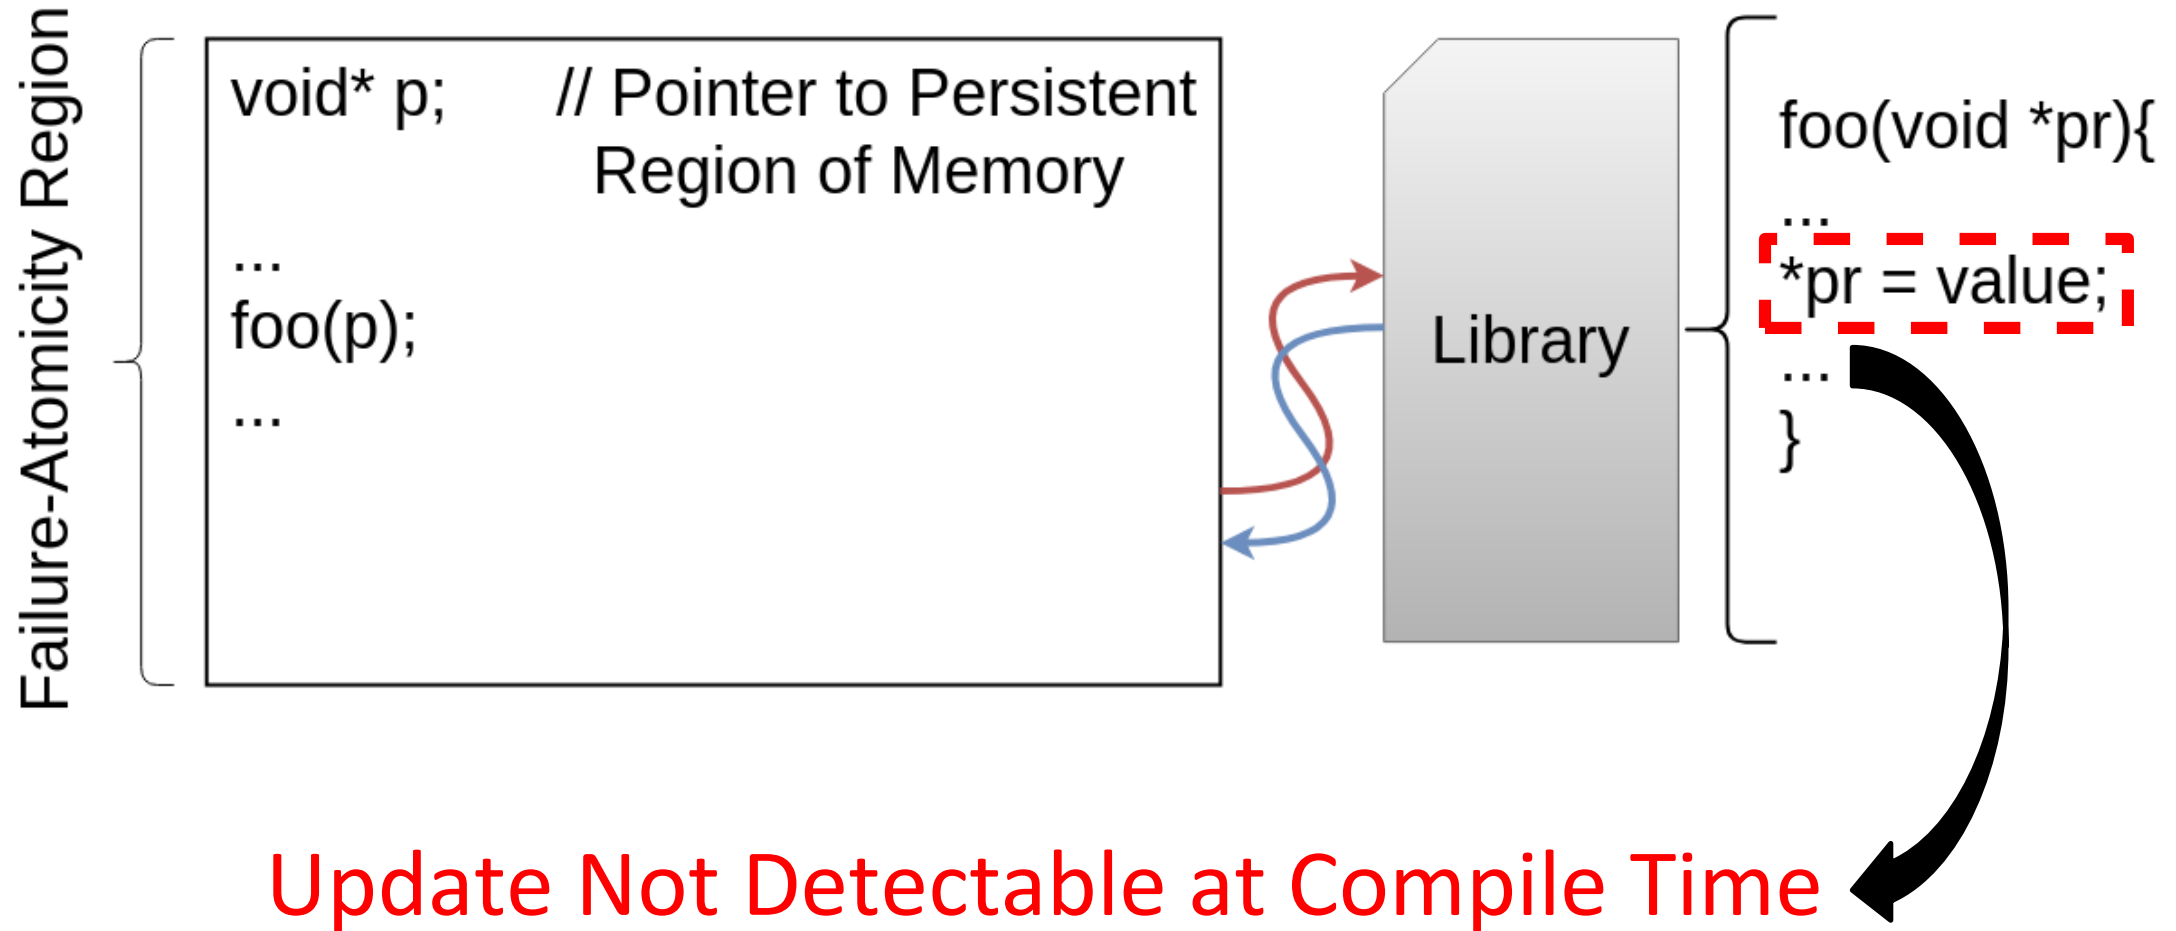
Detect Persist Updates at **Compile Time**

What Can Go Wrong?

Challenge #1: One Possible Scenario



Challenge #1: One Possible Scenario



Semantics for Delivering Persistency

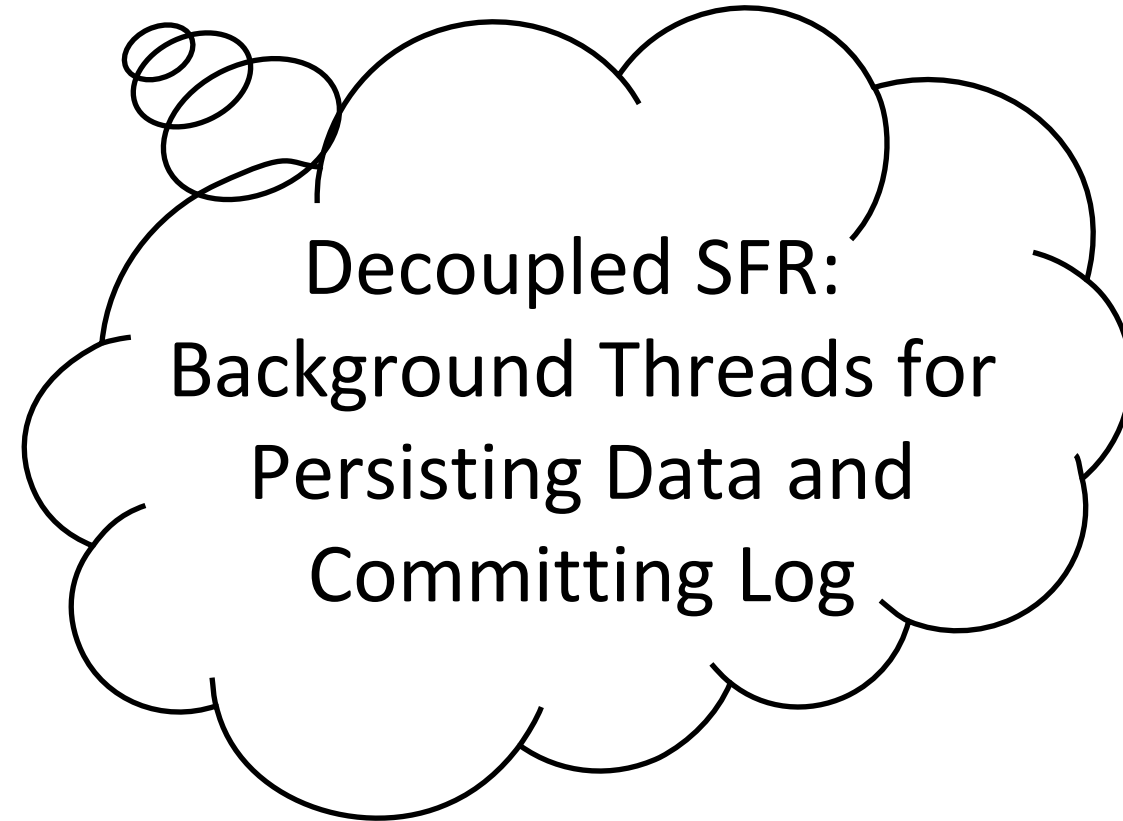
Challenge #2: Performance

High Logging Overheads (Atlas and Coupled SFR)

Semantics for Delivering Persistency

Challenge #2: Performance

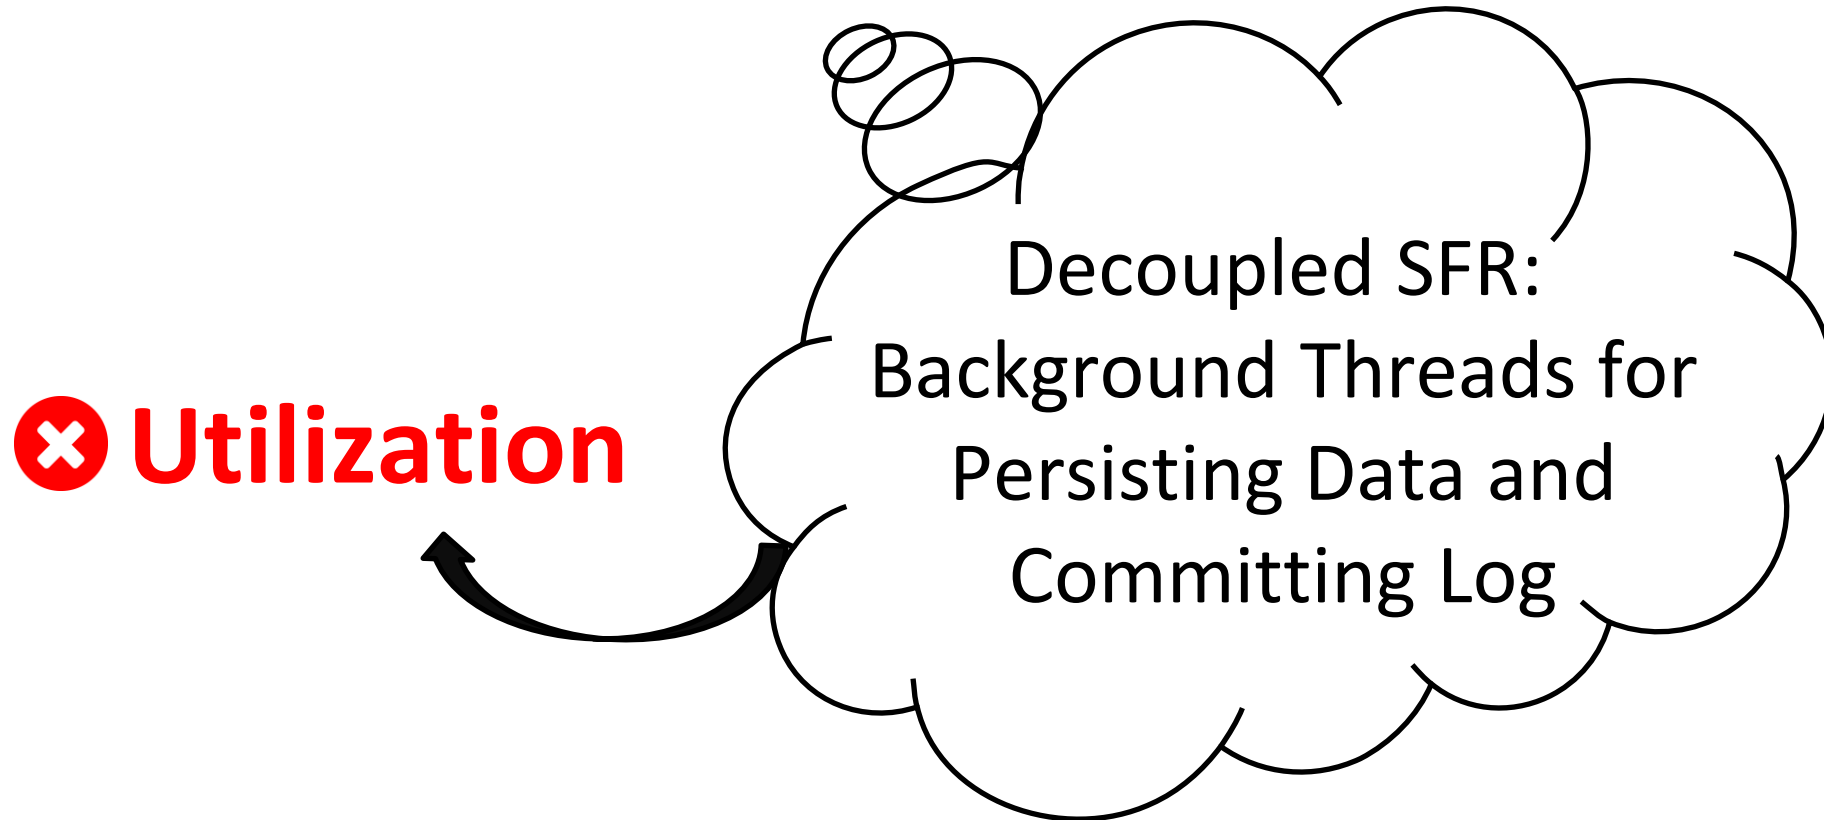
High Logging Overheads (Atlas and Coupled SFR)



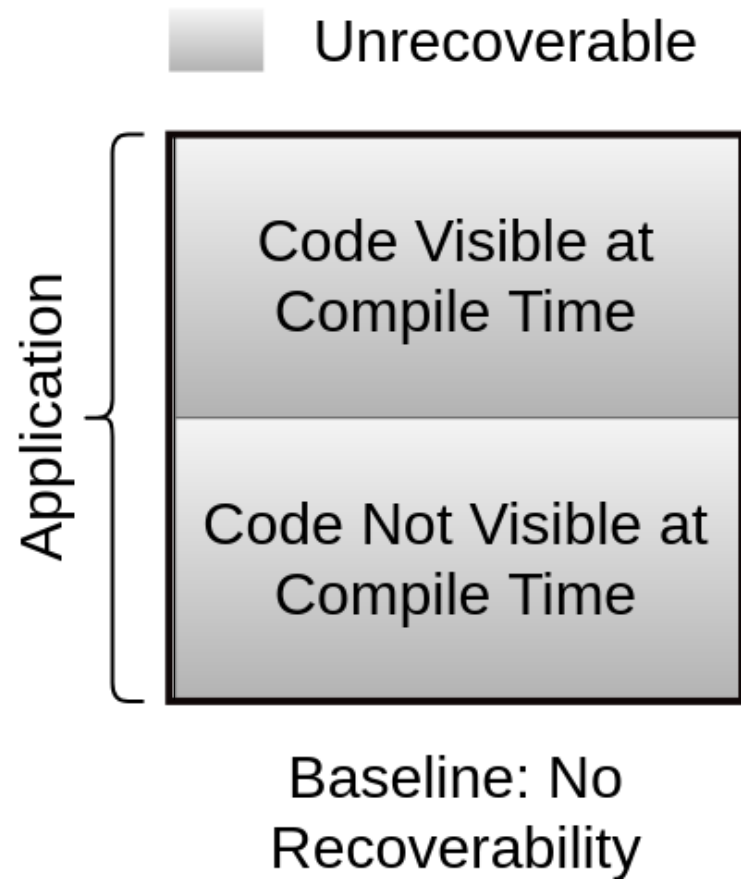
Semantics for Delivering Persistency

Challenge #2: Performance

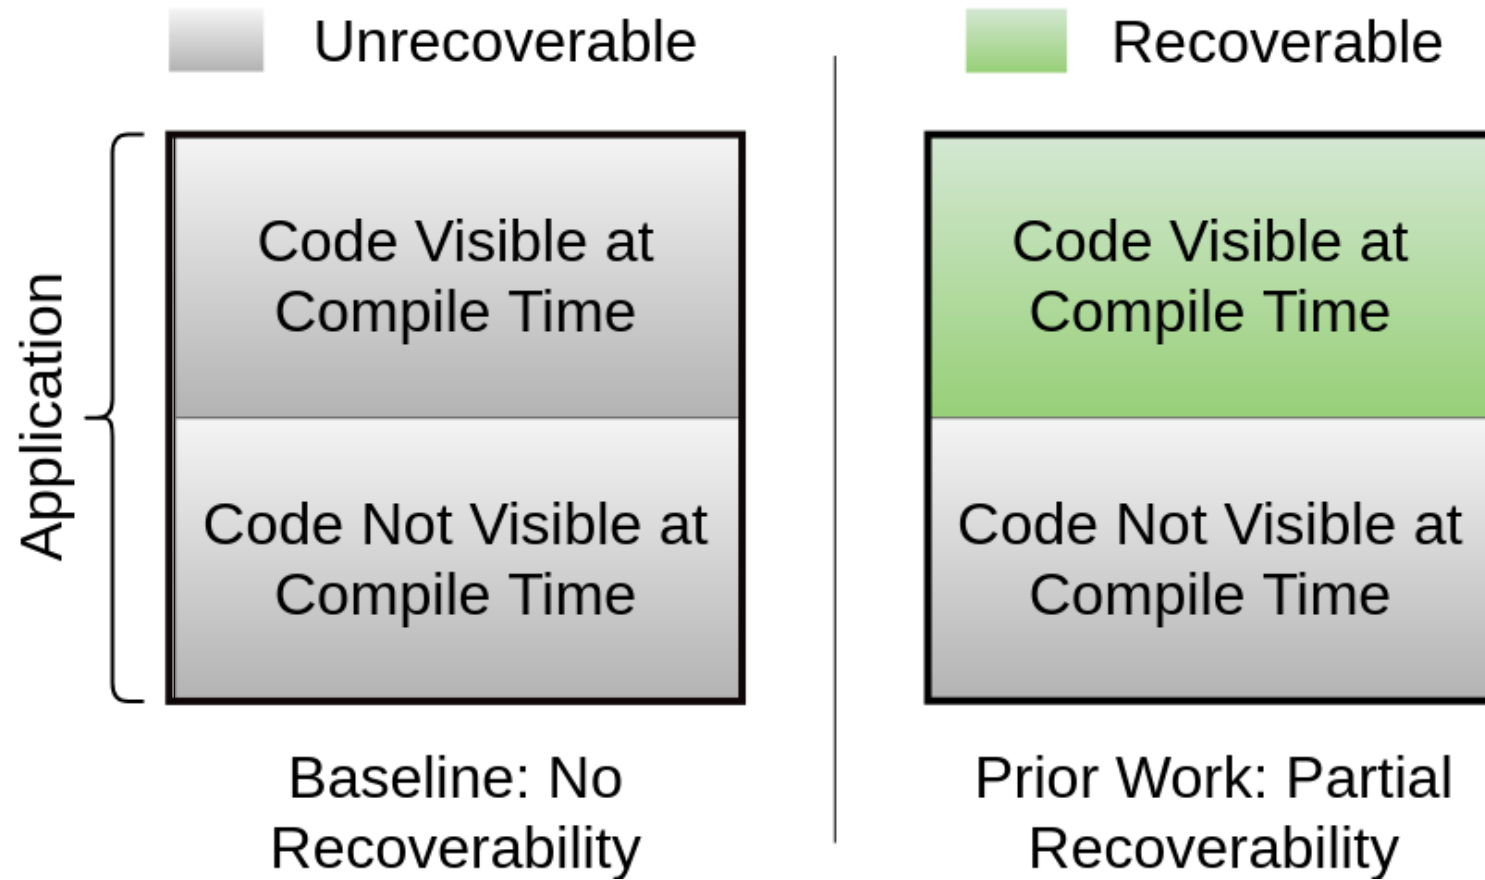
High Logging Overheads (Atlas and Coupled SFR)



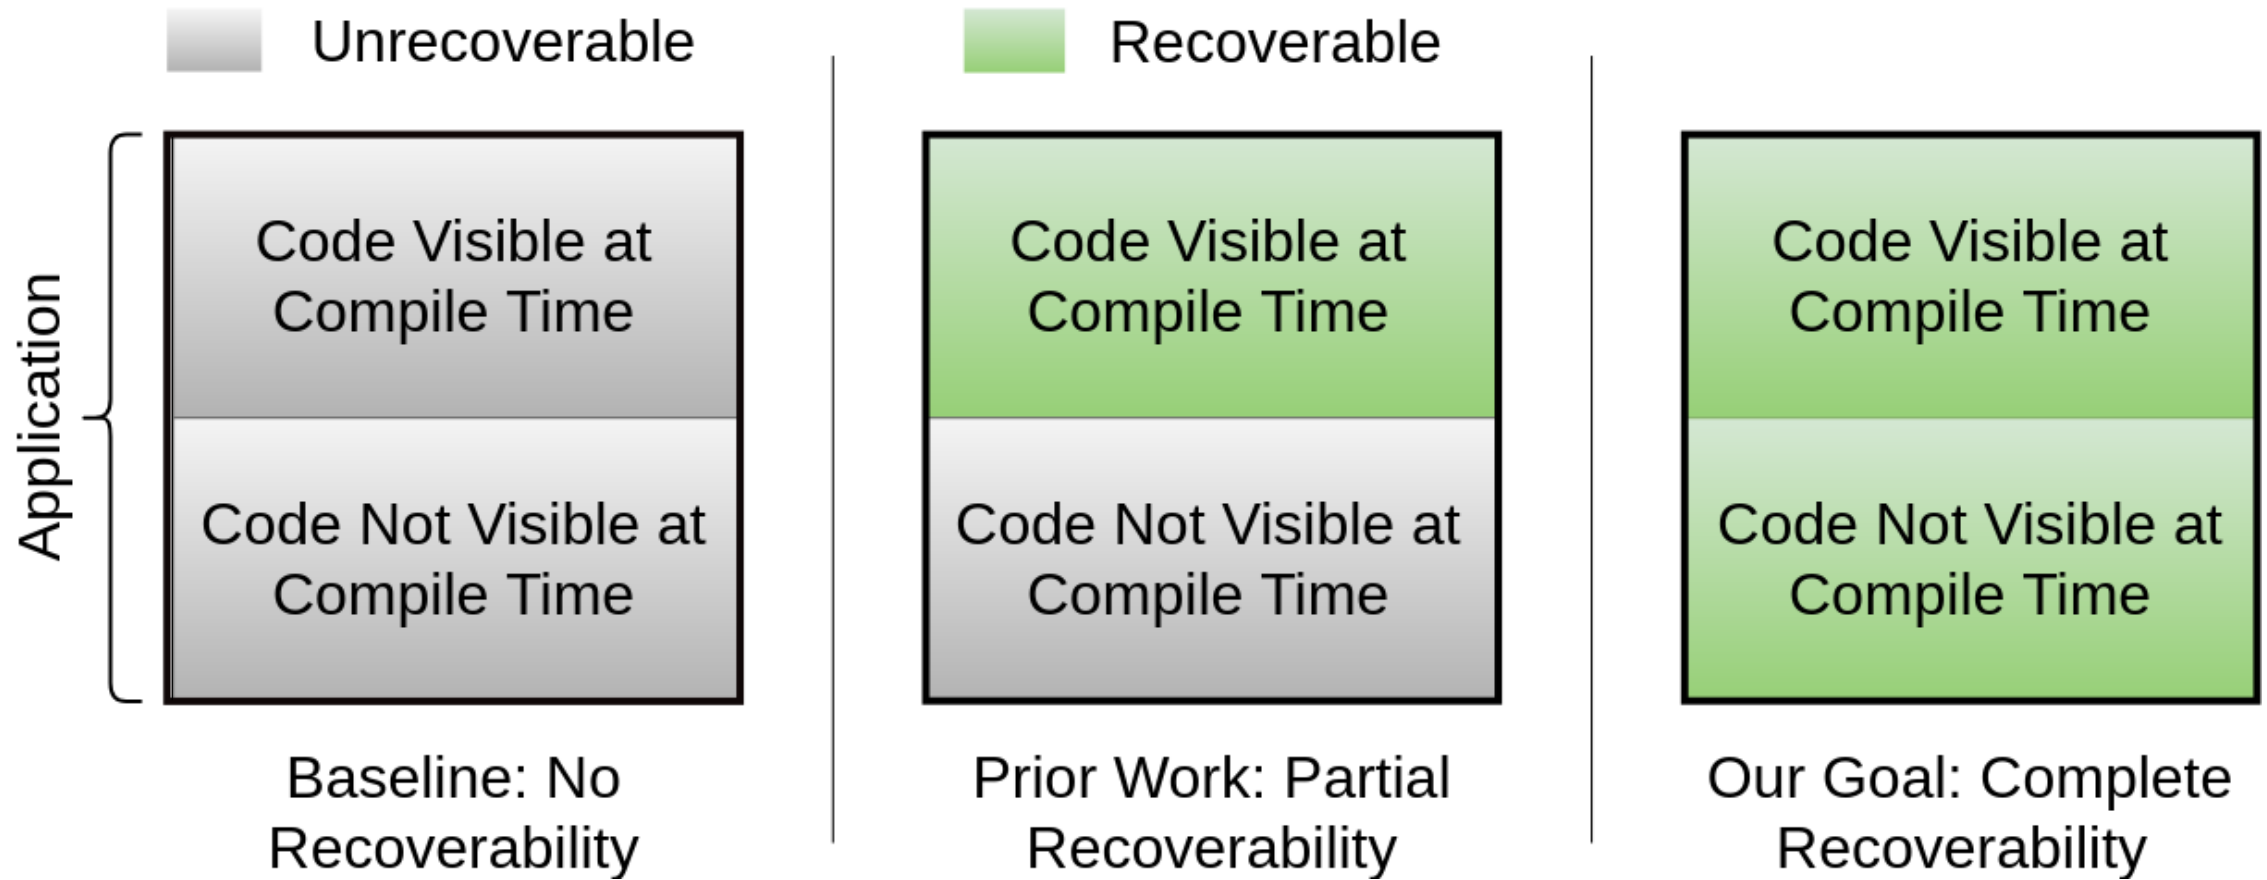
Idea #1: Persistency for the Whole Program



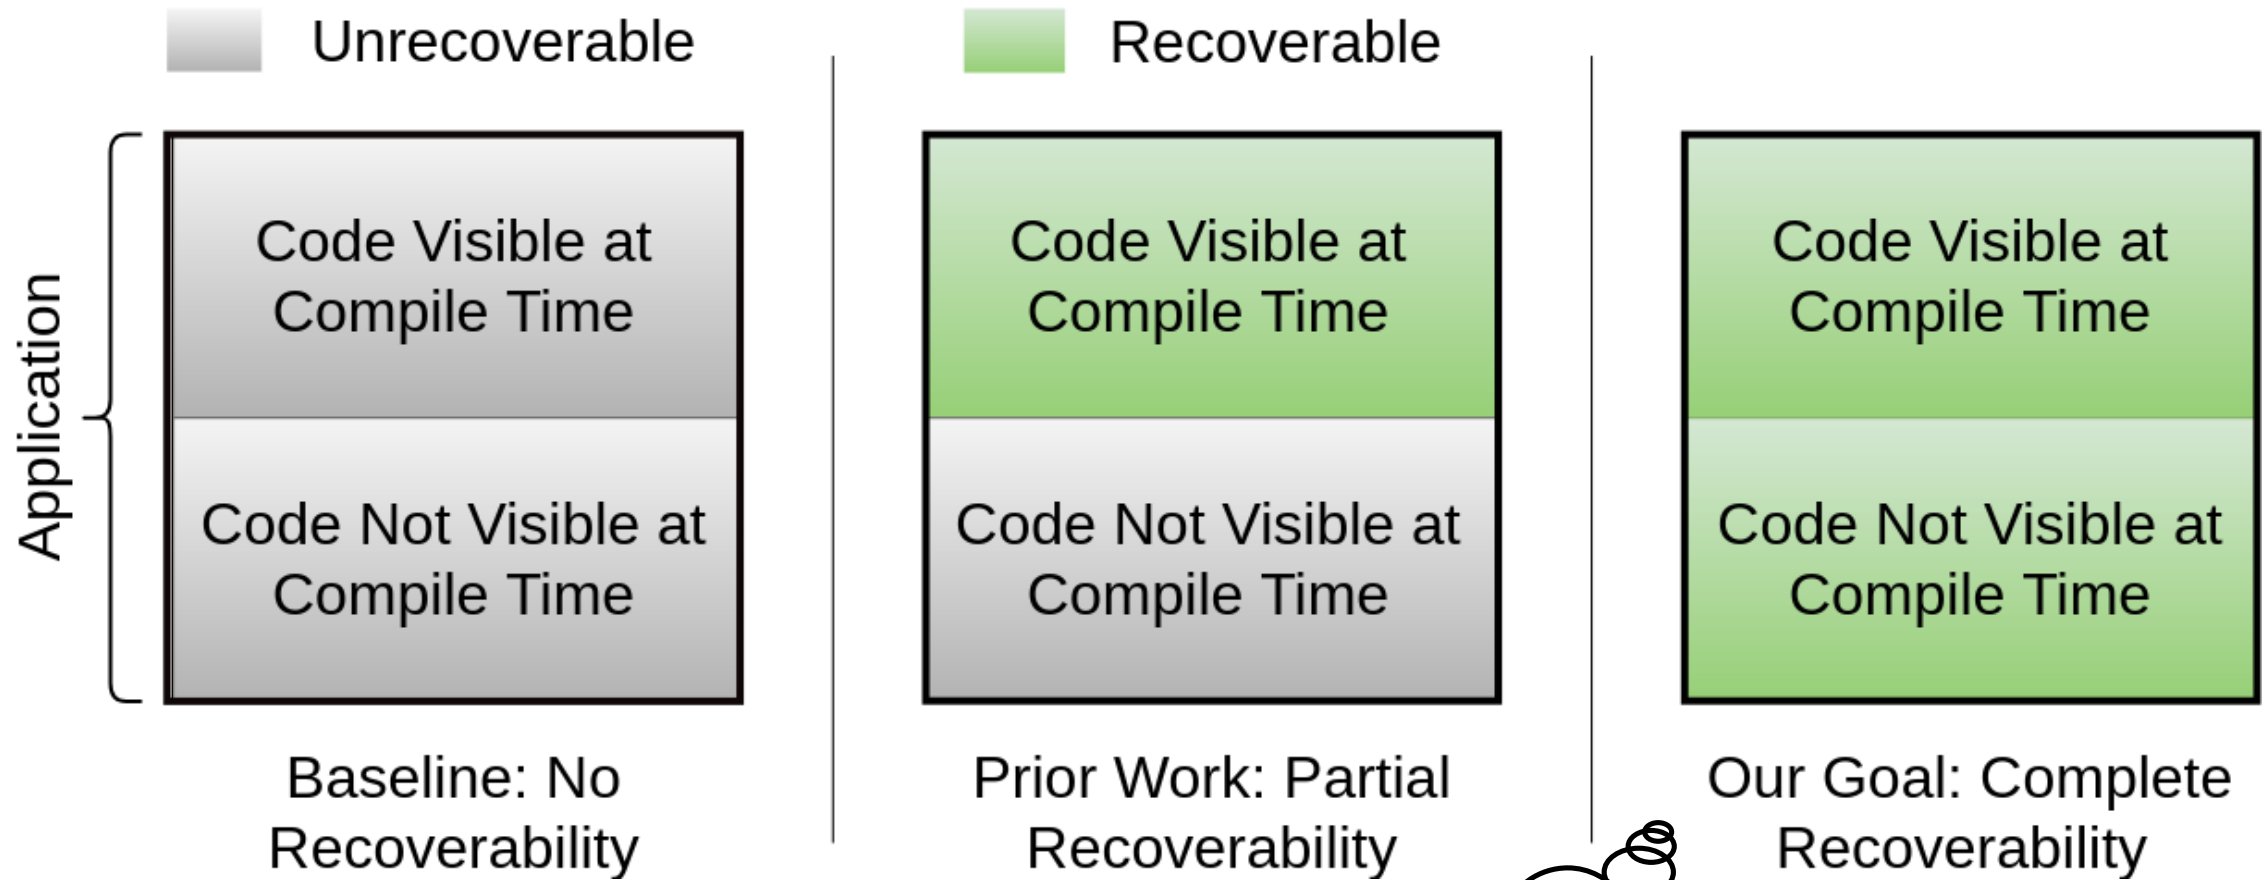
Idea #1: Persistency for the Whole Program



Idea #1: Persistency for the Whole Program

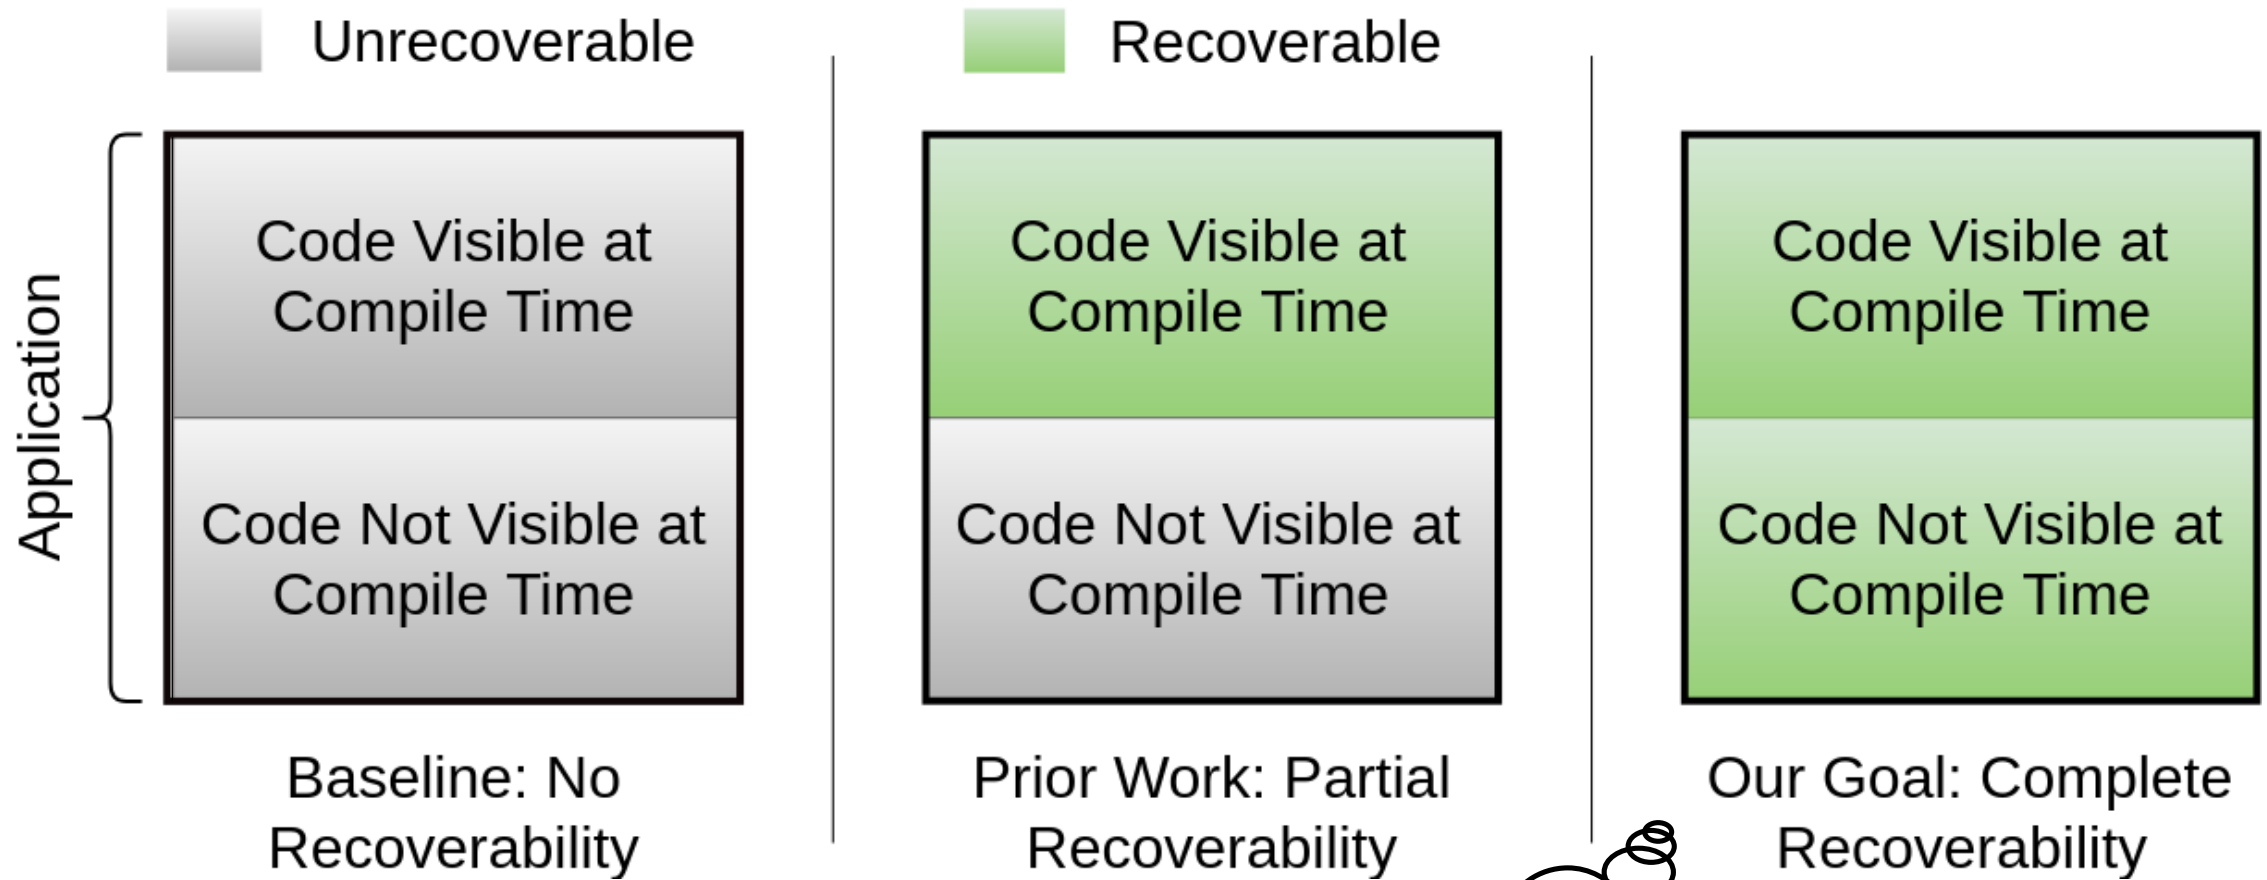


Idea #1: Persistency for the Whole Program

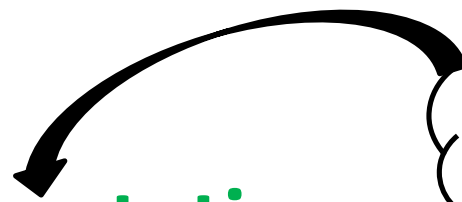
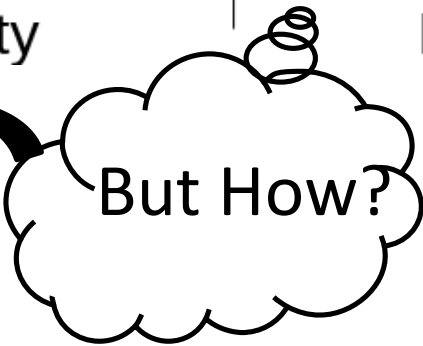


But How?

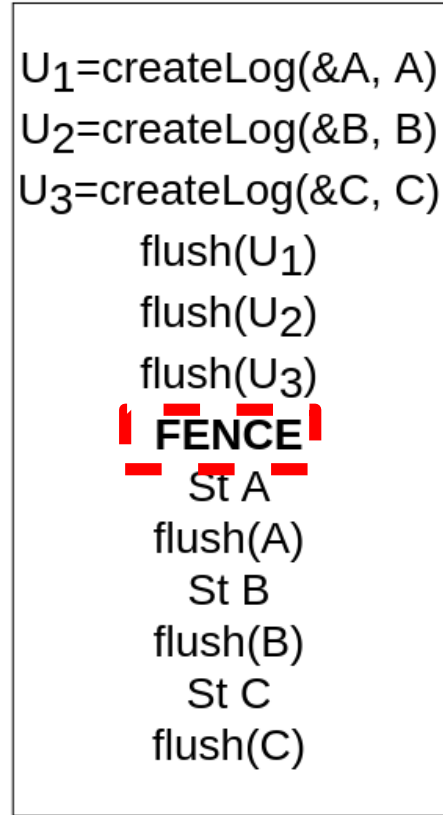
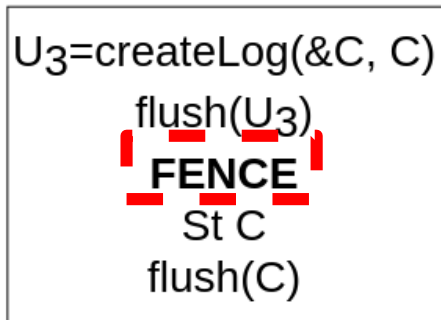
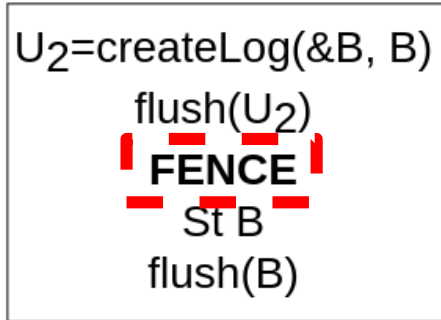
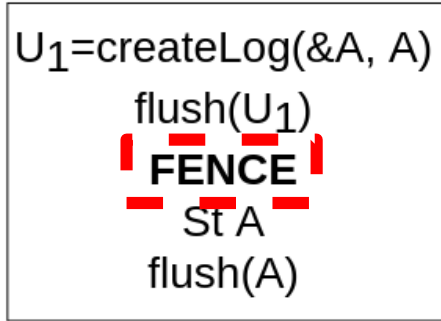
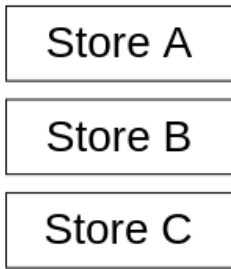
Idea #1: Persistency for the Whole Program



Binary Instrumentation



Idea #2: More Optimization in Logging



Baseline: No Logging

Prior Work: Serial Logging

Our Goal: Batch Logging

Idea #2: More Optimization in Logging

Store A
Store B
Store C

```
U1=createLog(&A, A)
flush(U1)
FENCE
St A
flush(A)
```

```
U2=createLog(&B, B)
flush(U2)
FENCE
St B
flush(B)
```

```
U3=createLog(&C, C)
flush(U3)
FENCE
St C
flush(C)
```

Baseline: No Logging

Prior Work: Serial Logging

```
U1=createLog(&A, A)
U2=createLog(&B, B)
U3=createLog(&C, C)
flush(U1)
flush(U2)
flush(U3)
FENCE
St A
flush(A)
St B
flush(B)
St C
flush(C)
```

Our Goal: Batch Logging



Can Hardware Help Us?

Evaluation Methodology

Tools:

- LLVM [Lattner' 04]
- DynamoRIO [Bruening' 12]

Benchmarks:

- WHISPER [Nalli' 17]
- Real-word Persistence Applications



Correct and Fast Persistency Guarantees

Sara Mahdizadeh Shahri
Aasheesh Kolli



PennState

YArch Workshop - February 2019

Backup Slides

Prior work

Atlas

Ensure Atomicity Within
Outermost Critical Section

Coupled SFR

Ensure atomicity within
Synchronization Free Regions

