

# Reducing Solid-State Drive Read Latency by Optimizing Read-Retry

**Jisung Park**<sup>1</sup>, Myungsuk Kim<sup>2</sup>, Myoungjun Chun<sup>2</sup>,  
Lois Orosa<sup>1</sup>, Jihong Kim<sup>2</sup>, and Onur Mutlu<sup>1</sup>

<sup>1</sup> **SAFARI**  
**ETH** zürich



**ASPLOS 2021 (Session 17: Solid State Drives)**

# Executive Summary

---

- **Problem:** Long read latency in modern SSDs due to read-retry
  - ❑ Multiple retry steps required to read an erroneous page
  - ❑ Read-retry frequently occurs in modern NAND flash memory
- **Goal:** Reduce the latency of each read-retry operation
- **Key Ideas:**
  - ❑ **Pipelined Read-Retry (PR<sup>2</sup>):** Concurrently perform consecutive retry steps using the CACHE READ command
  - ❑ **Adaptive Read-Retry (AR<sup>2</sup>):** Reduce read-timing parameters for every retry step by exploiting the reliability margin provided by strong ECC
- **Evaluation Results:** Our proposal improves SSD response time by
  - ❑ Up to 51% (35% on average) compared to a high-end SSD
  - ❑ Up to 32% (17% on average) compared to a state-of-the-art baseline

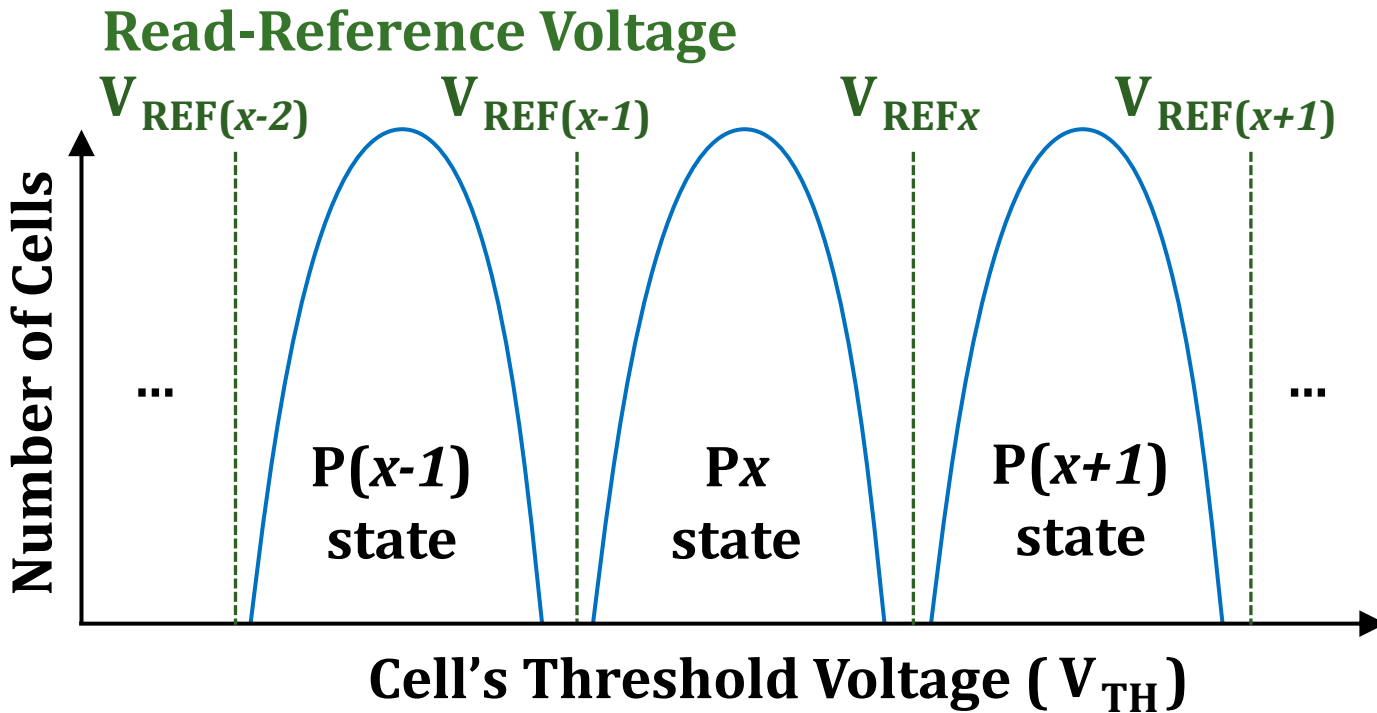
# Talk Outline

---

- Read-Retry in Modern NAND Flash-Based SSDs
- PR<sup>2</sup>: Pipelined Read-Retry
- AR<sup>2</sup>: Adaptive Read-Retry
- Evaluation Results

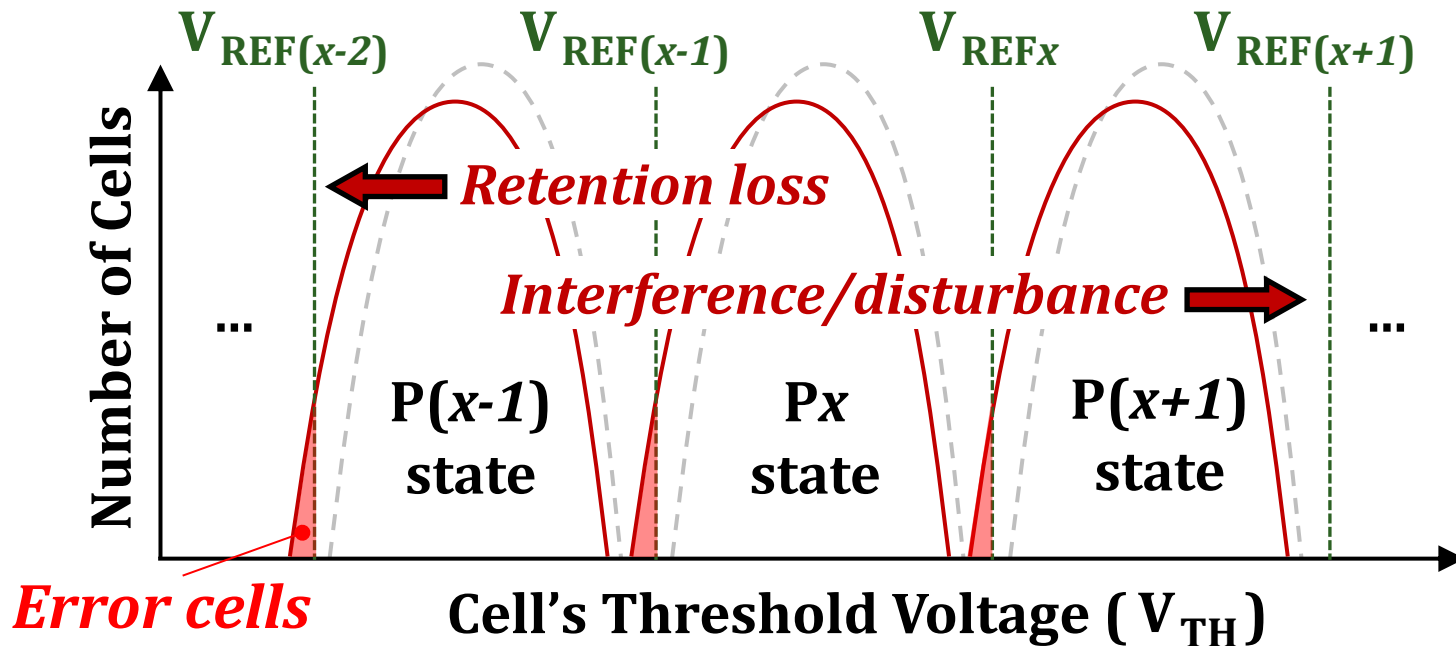
# Errors in NAND Flash Memory

- NAND flash memory stores data by using **cells'  $V_{TH}$  values**



# Errors in NAND Flash Memory


- Various sources **shift and widen** programmed  $V_{TH}$  states
  - Retention loss, program interference, read disturbance, etc.



# of error cells > ECC correction capability  
→ **Uncorrectable errors** in stored data

# Read-Retry Operation

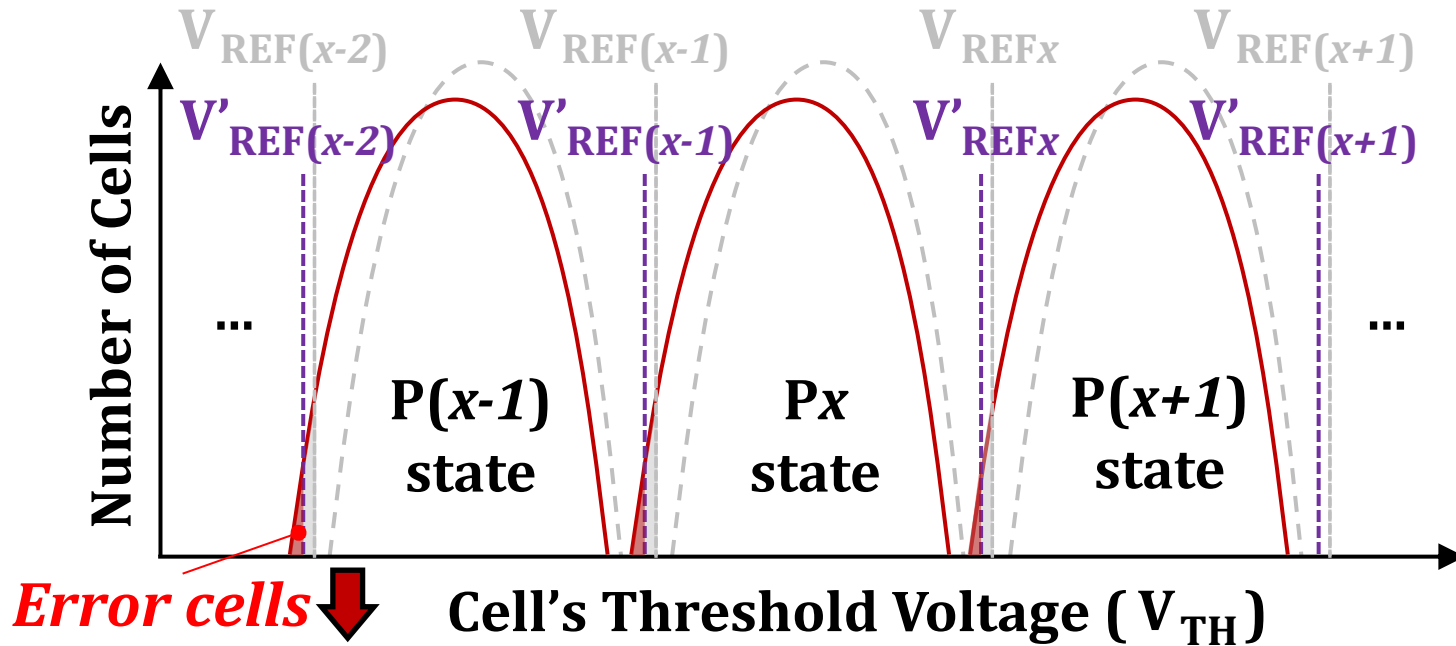
- Reads the page **again** with **adjusted**  $V_{REF}$  values



# Read-Retry Operation

- Reads the page **again** with **adjusted**  $V_{REF}$  values

## *Read-retry: Adjusting $V_{REF}$ values*



Read using **properly-adjusted**  $V_{REF}$  values  
→ Decreases # of raw bit errors

# Read-Retry: Performance Overhead


---

$t_{\text{DMA}}$ : Data transfer

$t_{\text{R}}$ : Page sensing

$t_{\text{ECC}}$ : ECC decoding

READ A




$N_{\text{ERR}} = 32 < \text{ECC capability } C_{\text{ECC}} = 72$



# Read-Retry: Performance Overhead

ECC Capability  $C_{\text{ECC}} = 72$



Read-retry increases the read latency almost **linearly** with the **number of retry steps**

# Talk Outline

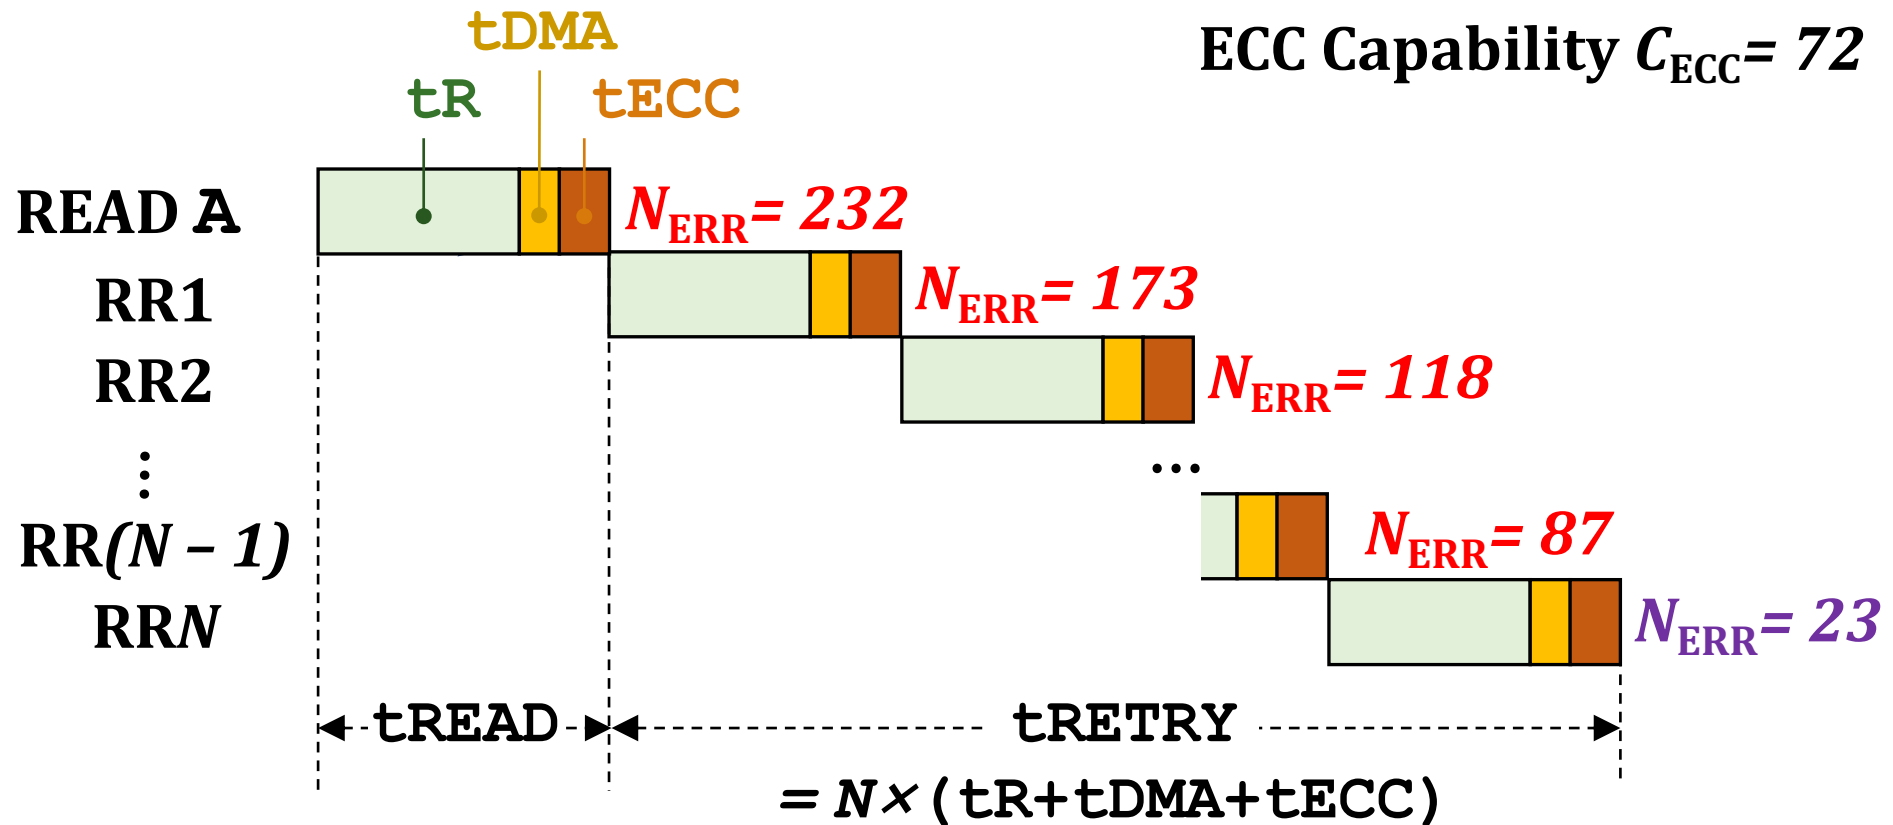
---

- Read-Retry in Modern NAND Flash-Based SSDs
- **PR<sup>2</sup>: Pipelined Read-Retry**
- AR<sup>2</sup>: Adaptive Read-Retry
- Evaluation Results

# PR<sup>2</sup>: Pipelined Read-Retry

- Key idea: Concurrently perform consecutive retry steps

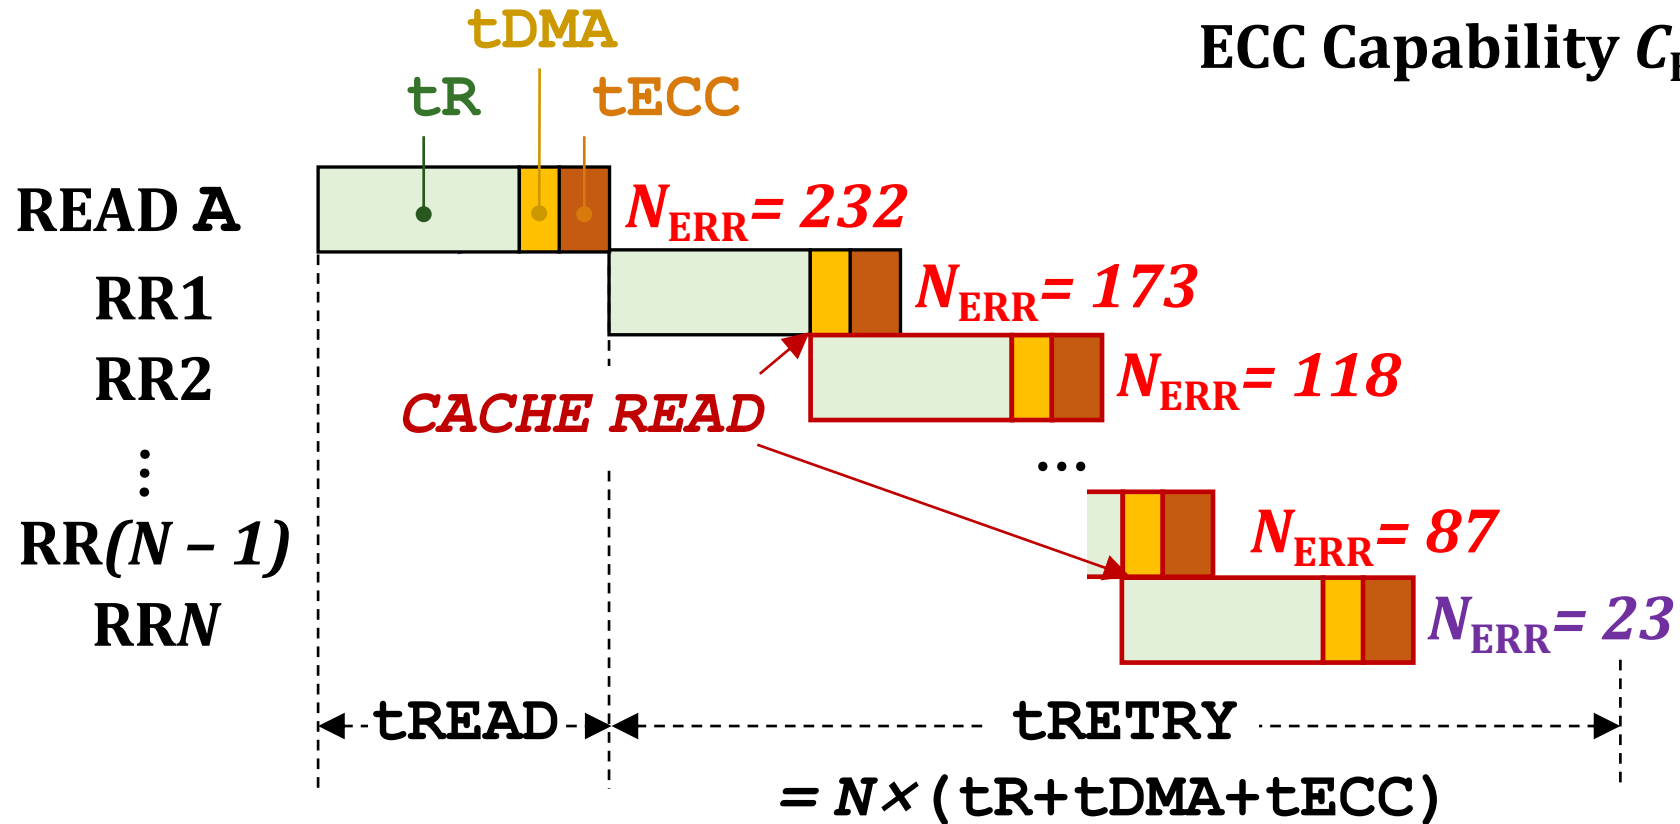
ECC Capability  $C_{\text{ECC}} = 72$



# PR<sup>2</sup>: Pipelined Read-Retry

- Key idea: **Concurrently** perform consecutive retry steps

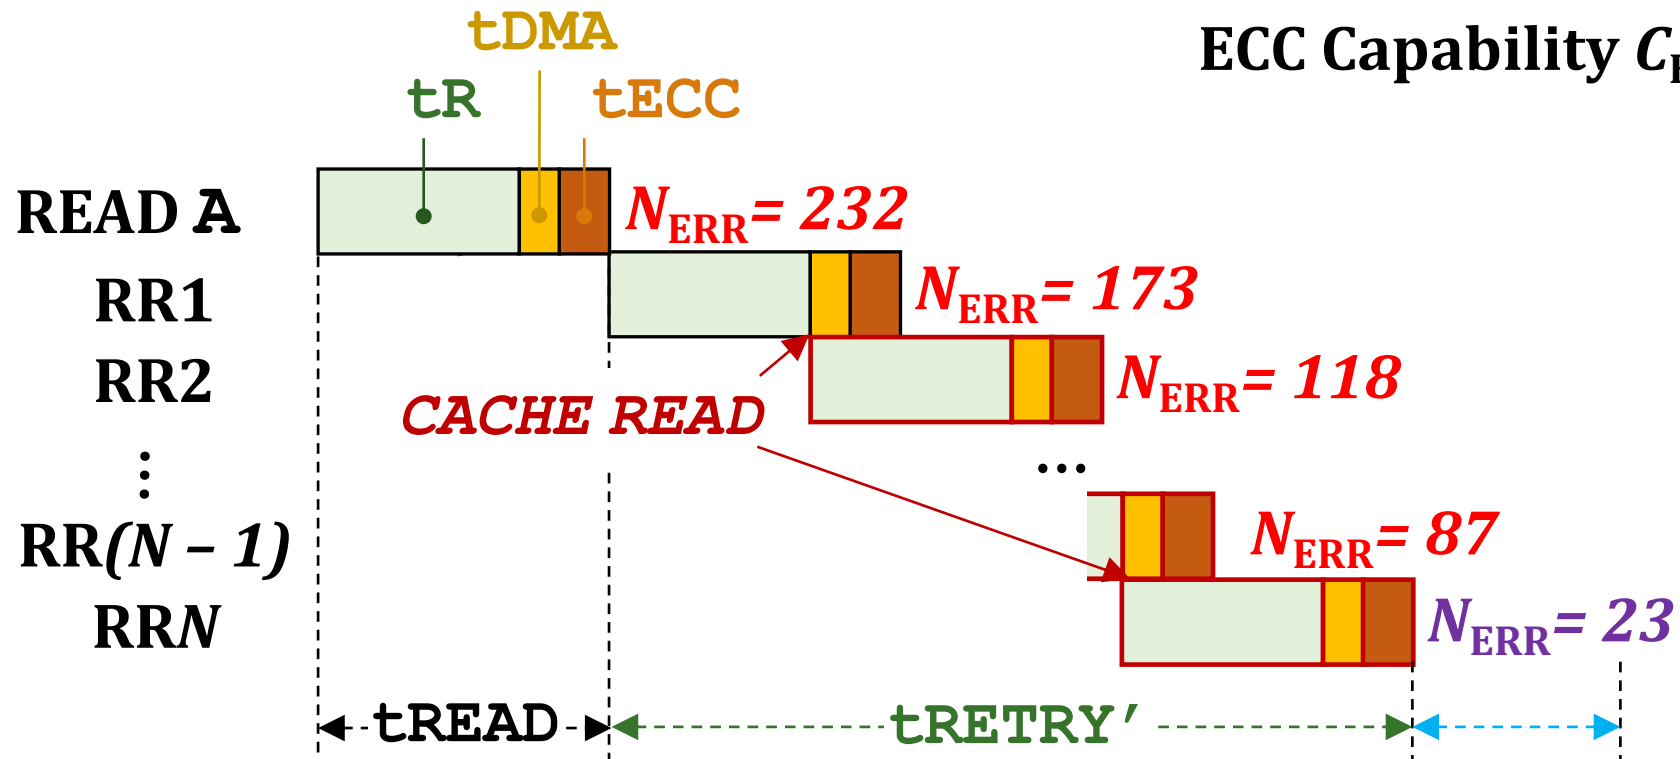
ECC Capability  $C_{\text{ECC}} = 72$



# PR<sup>2</sup>: Pipelined Read-Retry

- Key idea: **Concurrently** perform consecutive retry steps

ECC Capability  $C_{\text{ECC}} = 72$



PR<sup>2</sup>: Removes  $t_{\text{DMA}}$  &  $t_{\text{ECC}}$   
(~30% of each retry step) from the **critical path**

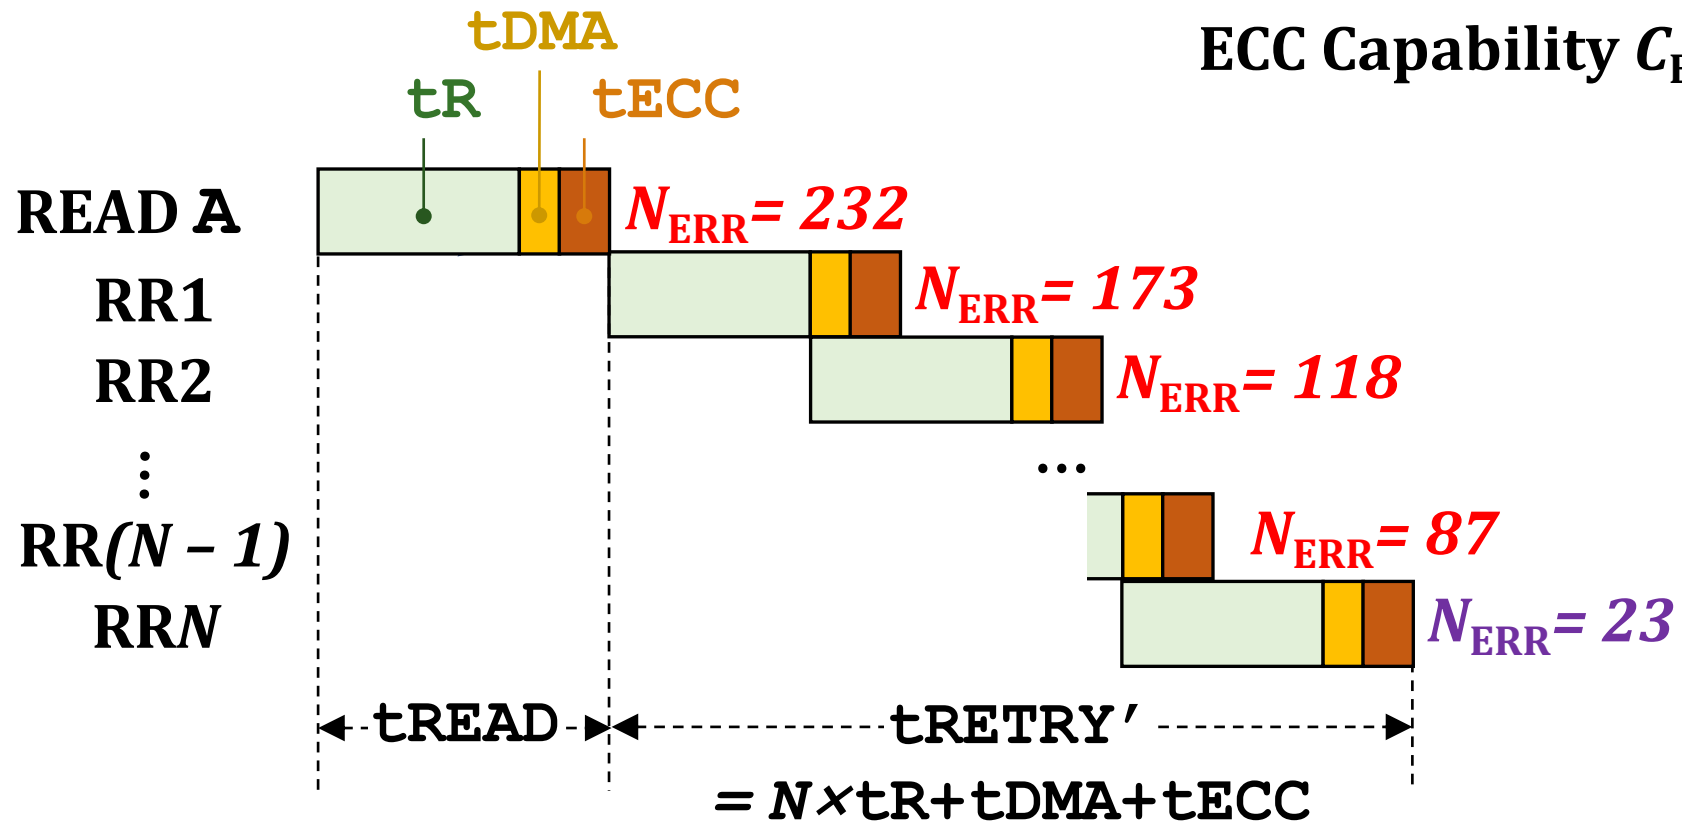
# Talk Outline

---

- Read-Retry in Modern NAND Flash-Based SSDs
- PR<sup>2</sup>: Pipelined Read-Retry
- AR<sup>2</sup>: Adaptive Read-Retry
- Evaluation Results

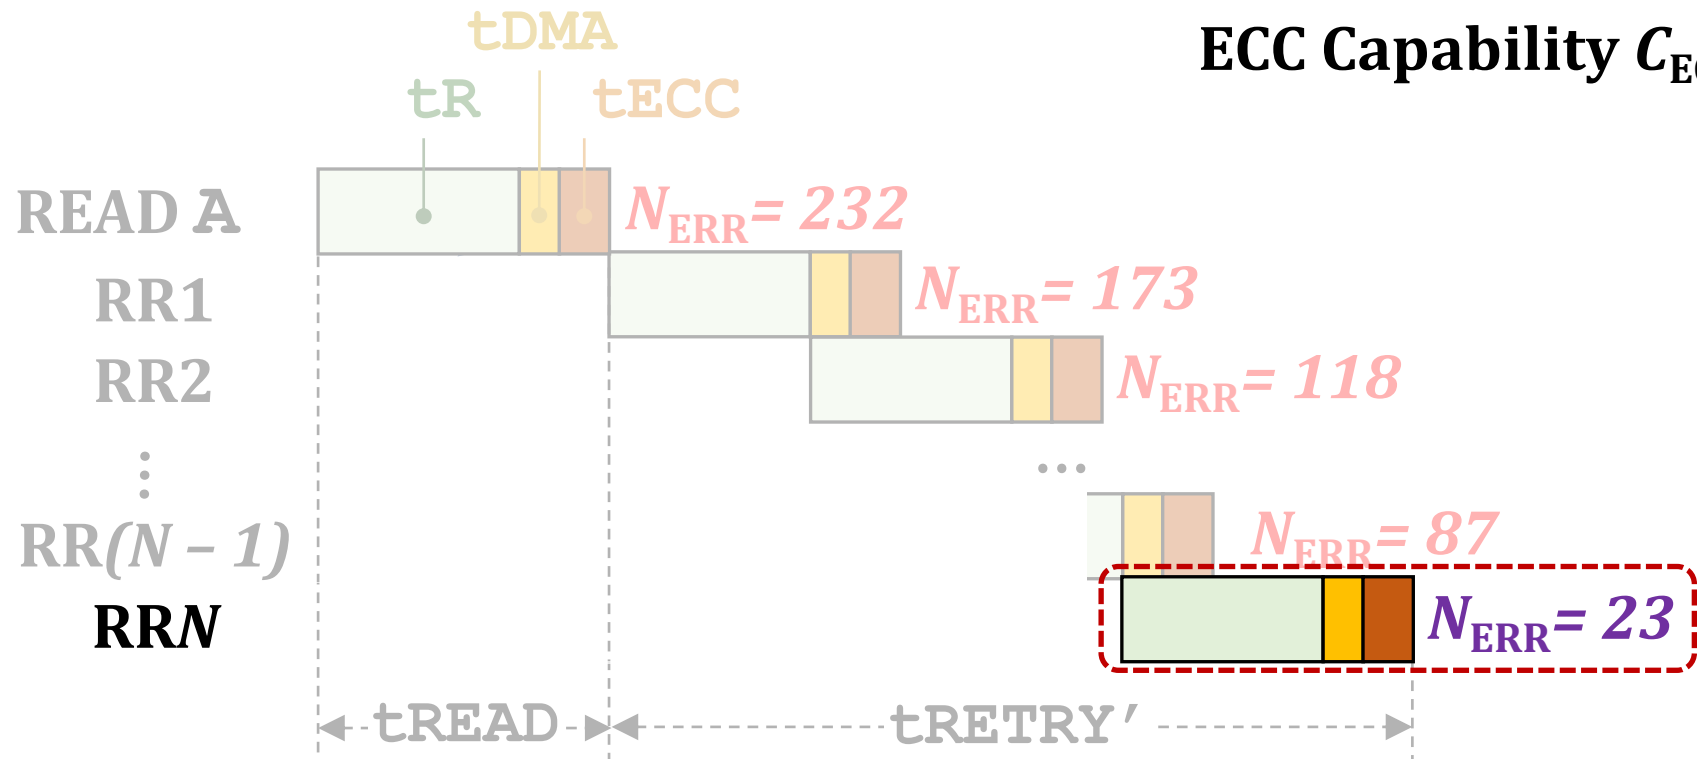
# AR<sup>2</sup>: Adaptive Read-Retry

ECC Capability  $C_{\text{ECC}} = 72$



# AR<sup>2</sup>: Adaptive Read-Retry

ECC Capability  $C_{\text{ECC}} = 72$



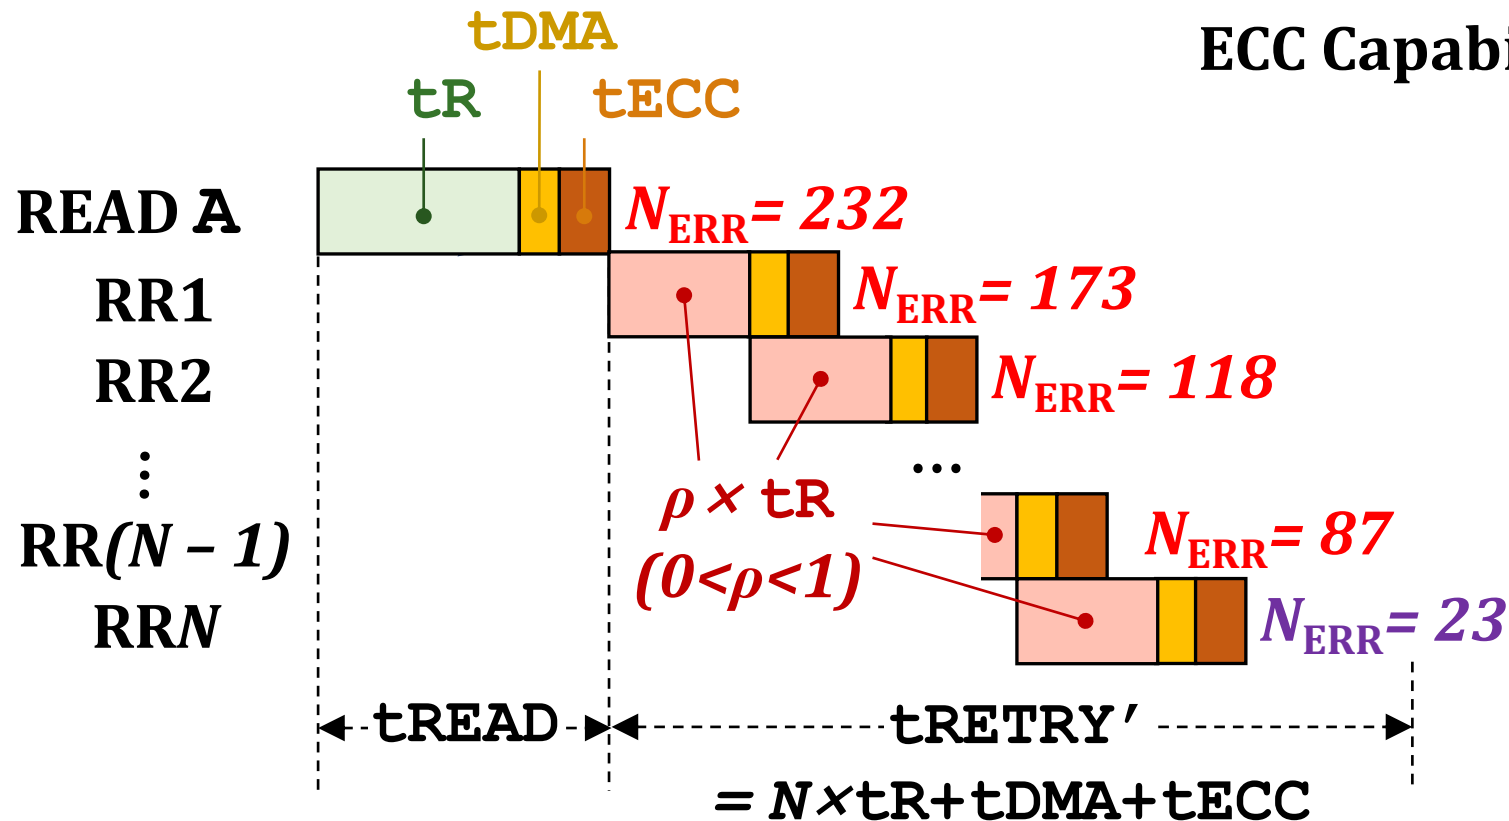
Observation: A **positive ECC margin** in the **final retry step** when read-retry succeeds



# AR<sup>2</sup>: Adaptive Read-Retry

- Key idea: Reduce read-timing parameters for every retry step

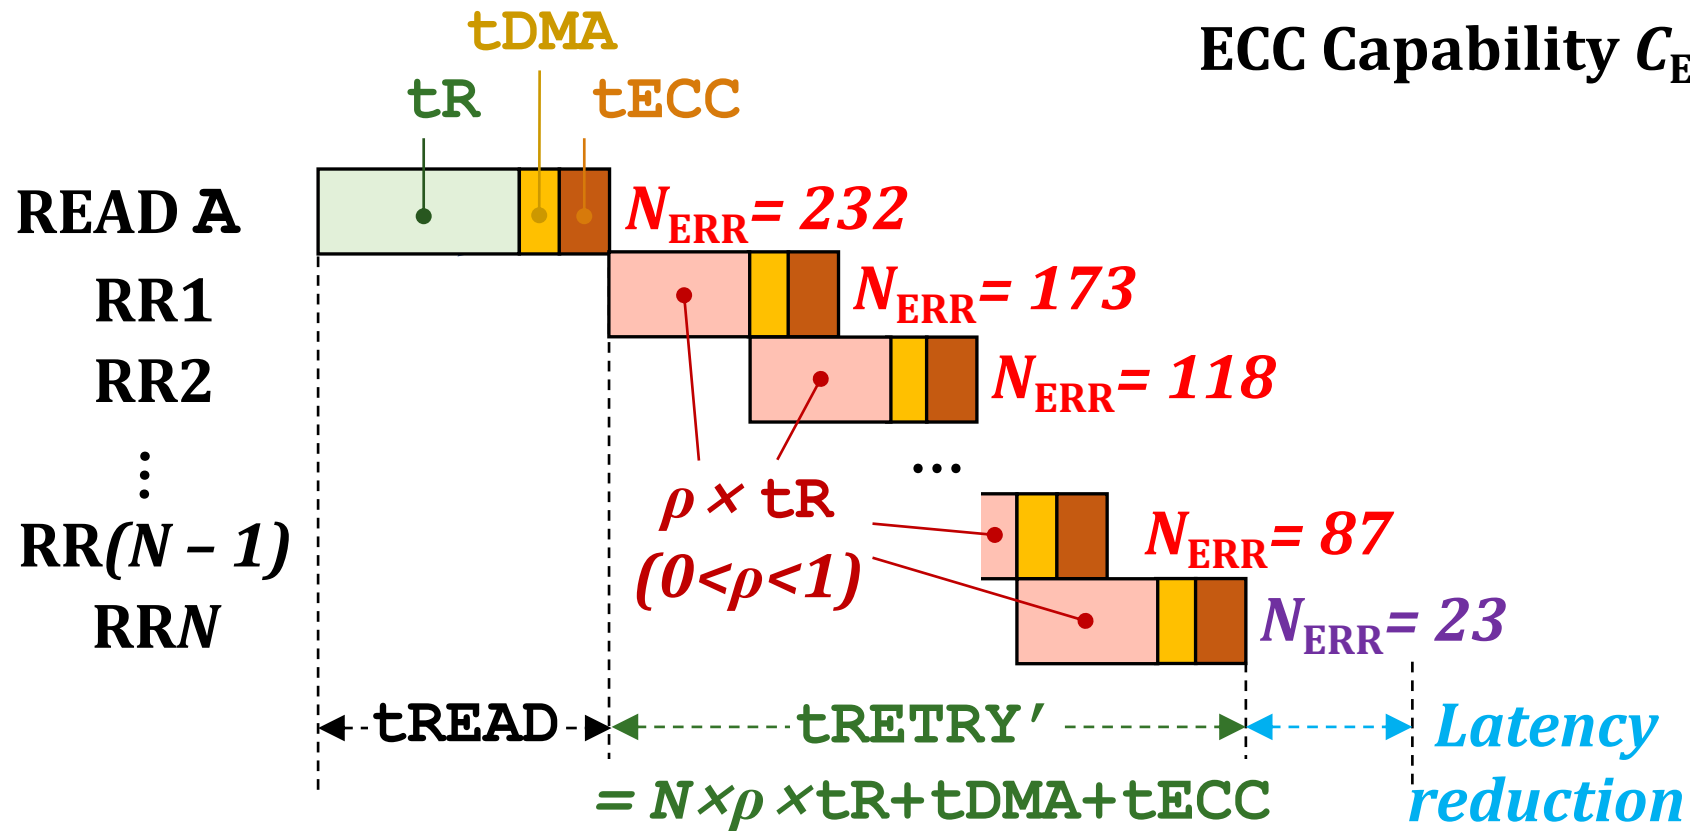
ECC Capability  $C_{\text{ECC}} = 72$



# AR<sup>2</sup>: Adaptive Read-Retry

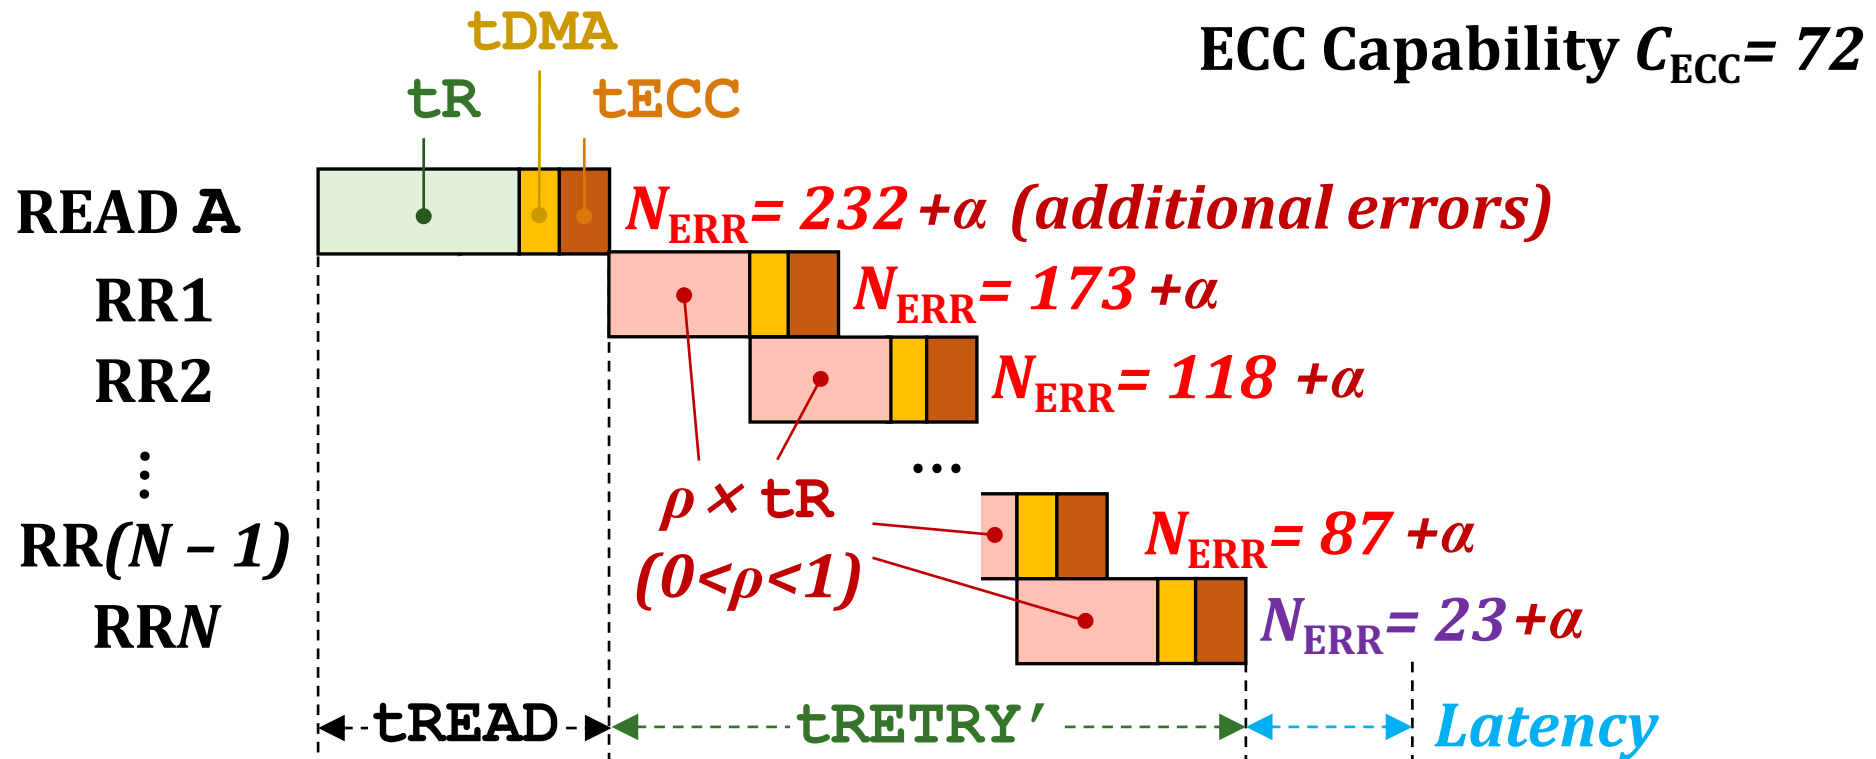
- Key idea: Reduce read-timing parameters for every retry step

ECC Capability  $C_{\text{ECC}} = 72$



# AR<sup>2</sup>: Adaptive Read-Retry

- Key idea: Reduce read-timing parameters for every retry step



Needs to ensure that  
# of additional errors  $<$  ECC margin

# Validation with Real 3D NAND Flash Chips

---

- **160 real** 48-layer Triple-Level Cell (TLC) NAND flash chips
- **Observation 1:** A **large ECC margin** in the final retry step even under **worst-case operating conditions**
  - Max. 40 errors per KiB under 1-year retention @ 2K P/E cycles
- **Observation 2:** A **large reliability margin** incorporated in **read-timing parameters**
  - 25%  $t_R$  reduction → Max. 23 additional errors

AR<sup>2</sup> can **easily work** in state-of-the-art NAND flash chips w/ **at least 25%  $t_R$  reduction**

# Talk Outline

---

- Read-Retry in Modern NAND Flash-Based SSDs
- PR<sup>2</sup>: Pipelined Read-Retry
- AR<sup>2</sup>: Adaptive Read-Retry
- Evaluation Results

# Evaluation Results

---

- Simulation using **MQSim** [Tavakkol, FAST18] and 12 real workloads
- Our proposal improves **SSD response time** by
  - Up to **51%** (**35%** on average) compared to a high-end SSD w/o read-retry mitigation
  - Up to **32%** (**17%** on average) compared to a **state-of-the-art read-retry mitigation technique**
- Many more detailed studies in the paper

# Reducing Solid-State Drive Read Latency by Optimizing Read-Retry

**Jisung Park**<sup>1</sup>, Myungsuk Kim<sup>2</sup>, Myoungjun Chun<sup>2</sup>,  
Lois Orosa<sup>1</sup>, Jihong Kim<sup>2</sup>, and Onur Mutlu<sup>1</sup>

<sup>1</sup> **SAFARI**  
**ETH** zürich



**ASPLOS 2021 (Session 17: Solid State Drives)**