A Case for Refresh Pausing in DRAM Memory Systems

Prashant Nair

Chia-Chen Chou

Moinuddin Qureshi

Introduction

- Dynamic Random Access Memory (DRAM) used as main memory
- DRAM stores data as charge on capacitor

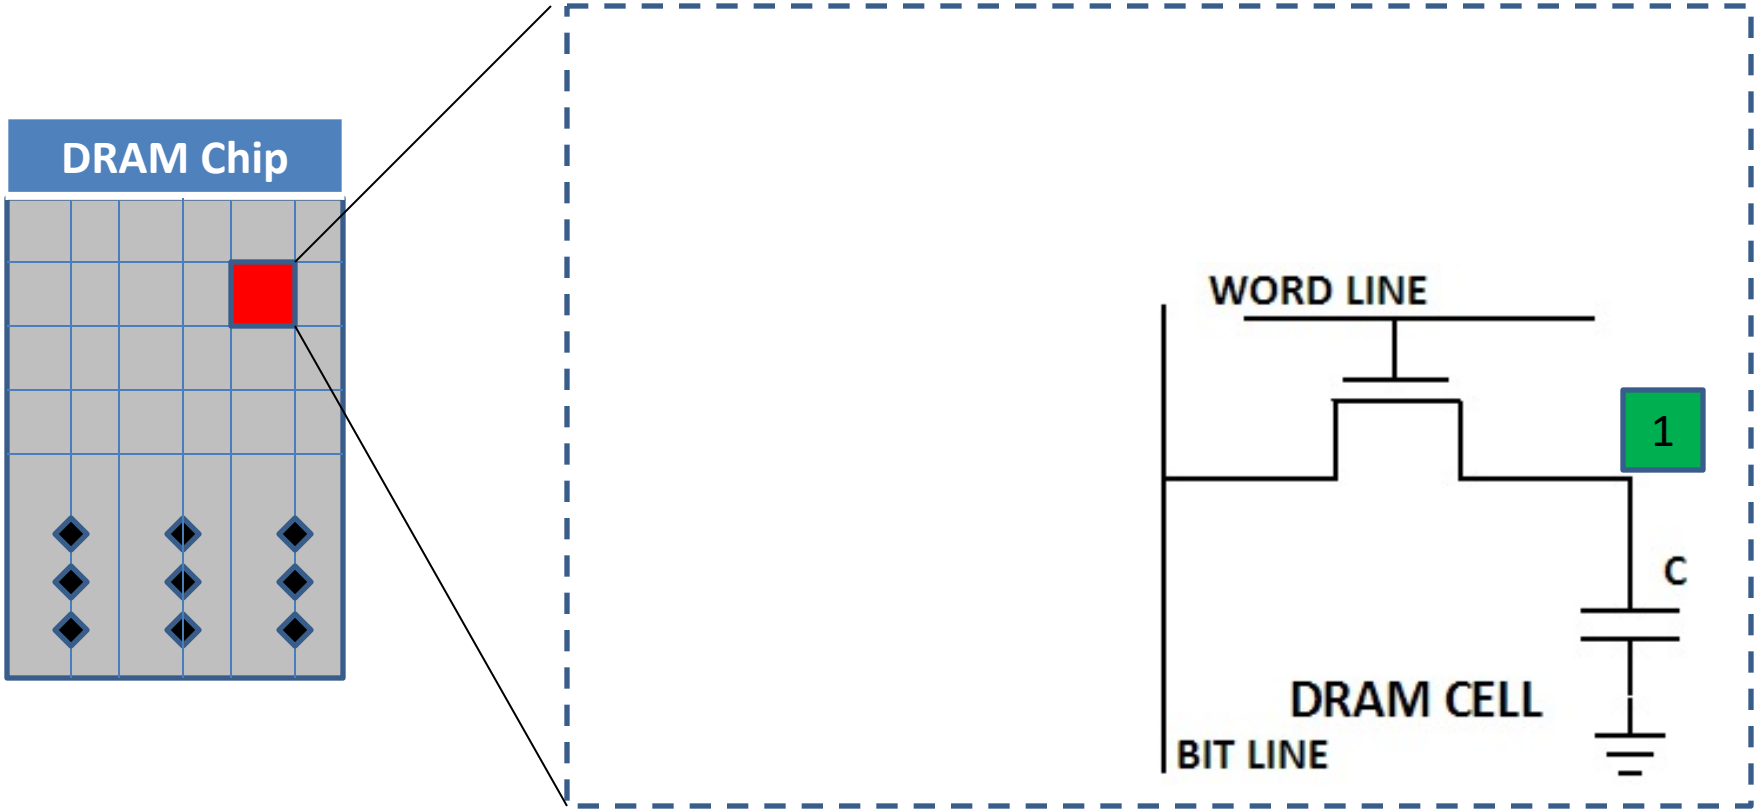

Introduction

- Dynamic Random Access Memory (DRAM) used as main memory
- DRAM stores data as charge on capacitor

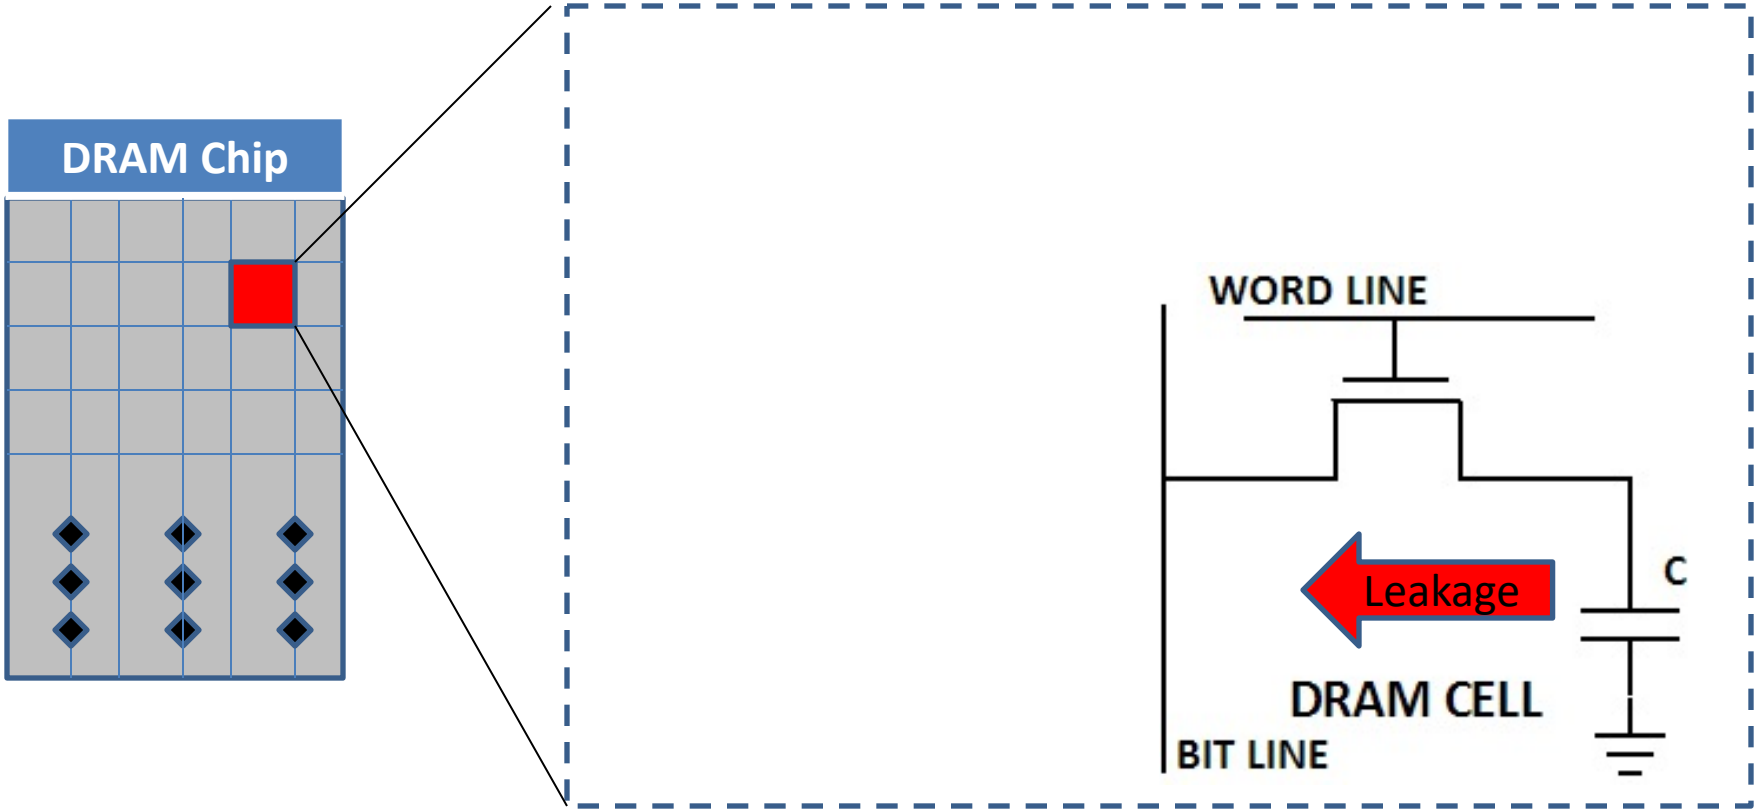

Introduction

- Dynamic Random Access Memory (DRAM) used as main memory
- DRAM stores data as charge on capacitor

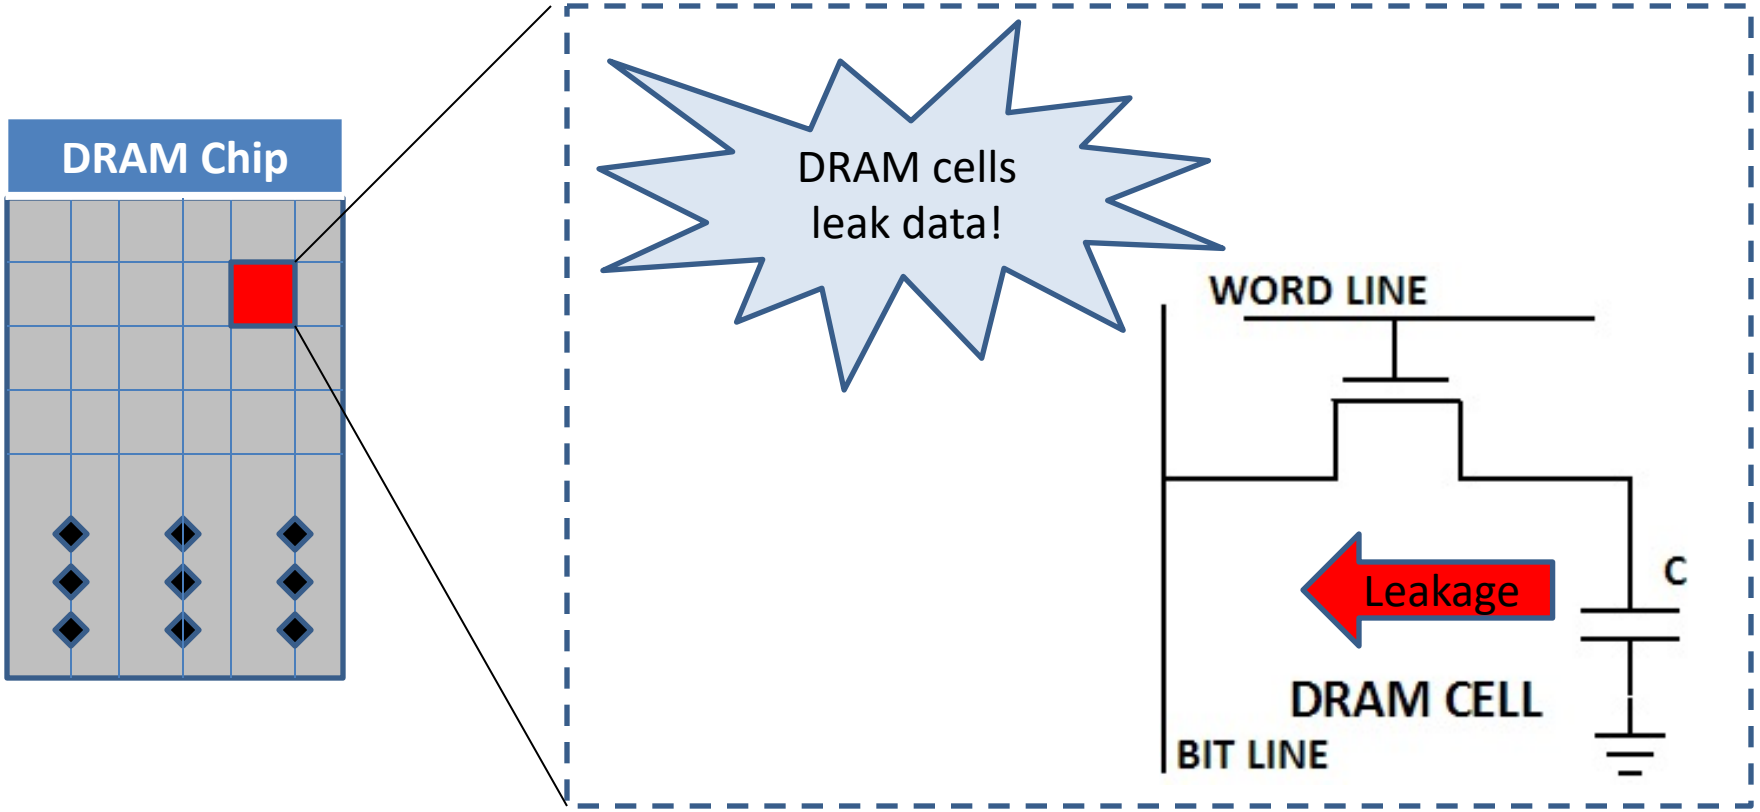

DRAM is a volatile memory ➔ Charge leaks quickly

Refresh: Restoring Data in DRAM

DRAM maintains data by Refresh operations

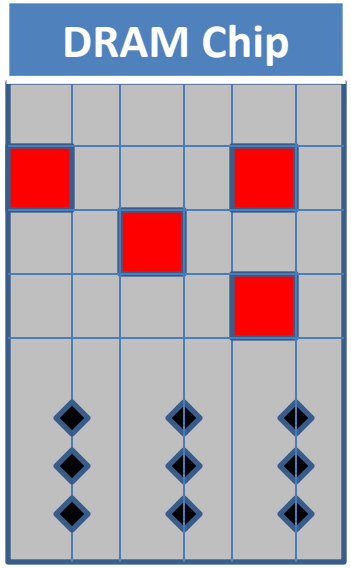

Refresh: Restoring Data in DRAM

DRAM maintains data by Refresh operations

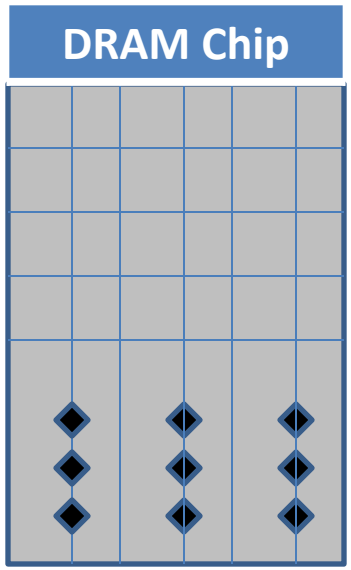

Charge on cells restored

Refresh: Restoring Data in DRAM

DRAM maintains data by Refresh operations

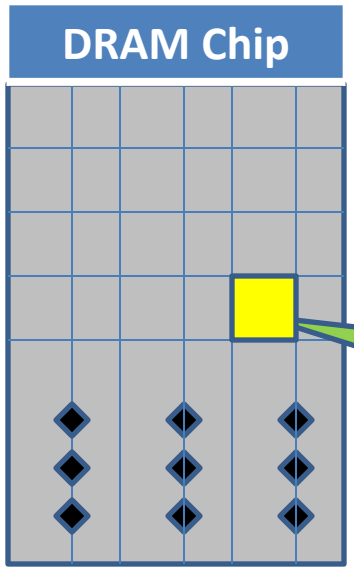

Charge on cells restored

JEDEC specified DRAM retention time:
64ms (< 85 C)
32ms (> 85 C)

Time between Refresh ≤ Retention Time

DRAM relies on Refresh for data integrity

Refresh: A Growing Problem

Time spent in Refresh proportional to number of Rows

Increasing memory capacity ➔ More time spent in Refresh

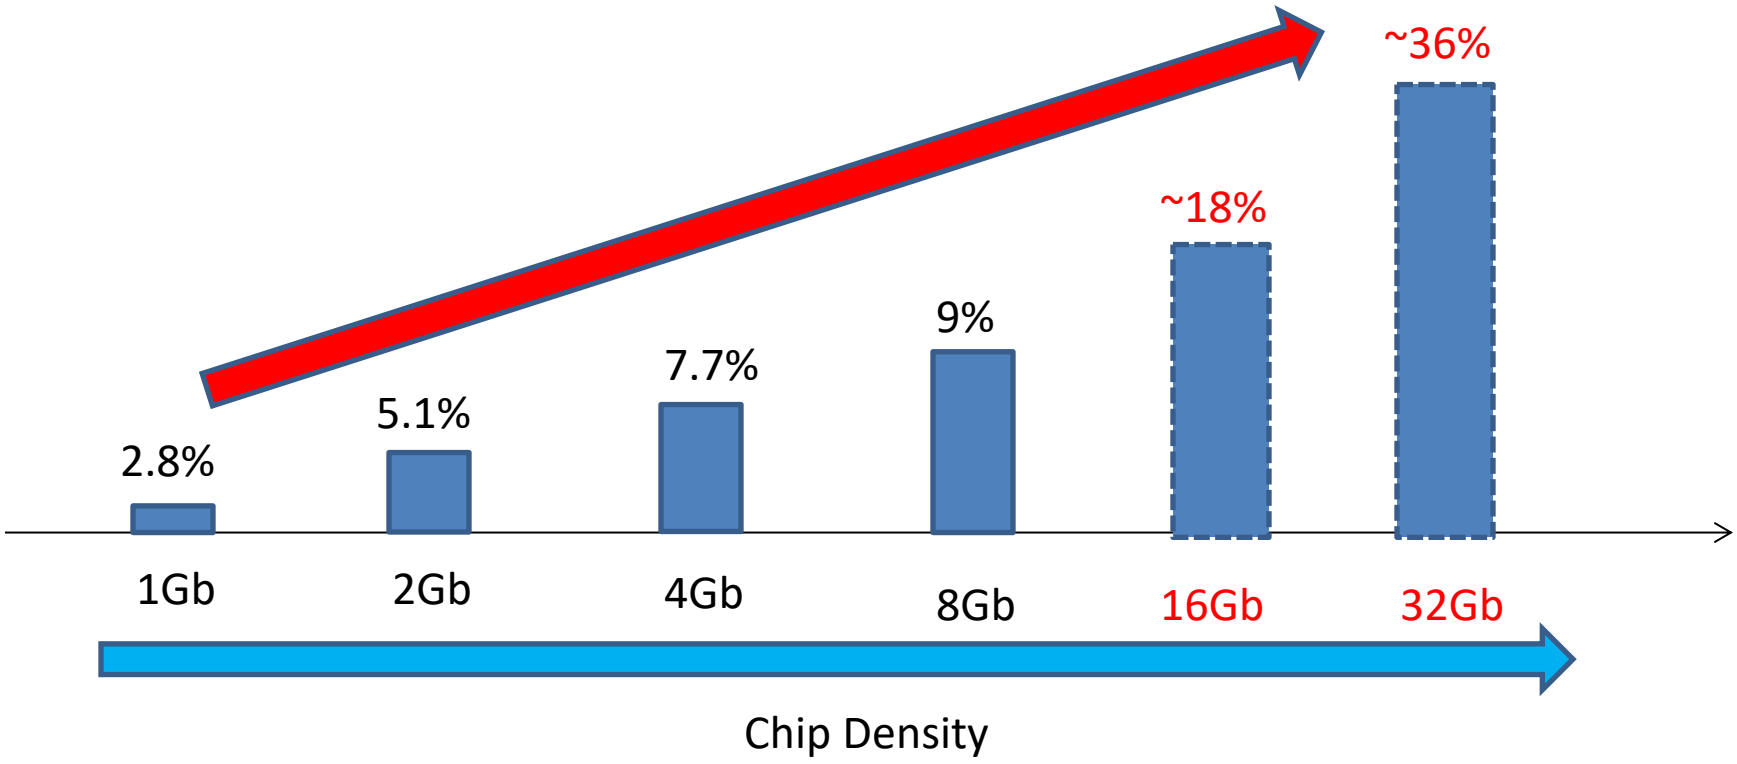

The time for doing Refresh is increasing with chip density

Refresh Blocks Reads

Memory unavailable for Read/Write during Refresh

Refresh Blocks Reads

Memory unavailable for Read/Write during Refresh

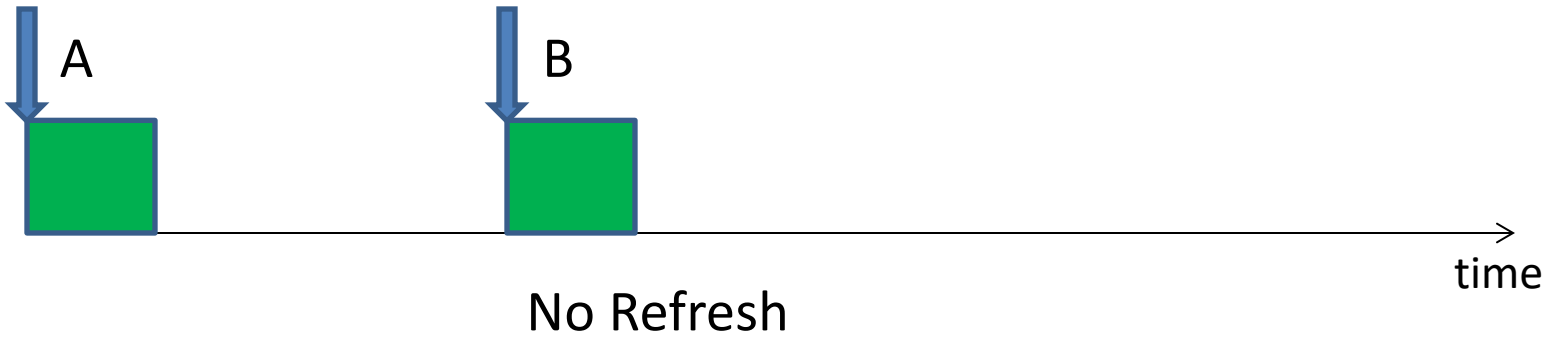

Refresh Blocks Reads

Memory unavailable for Read/Write during Refresh

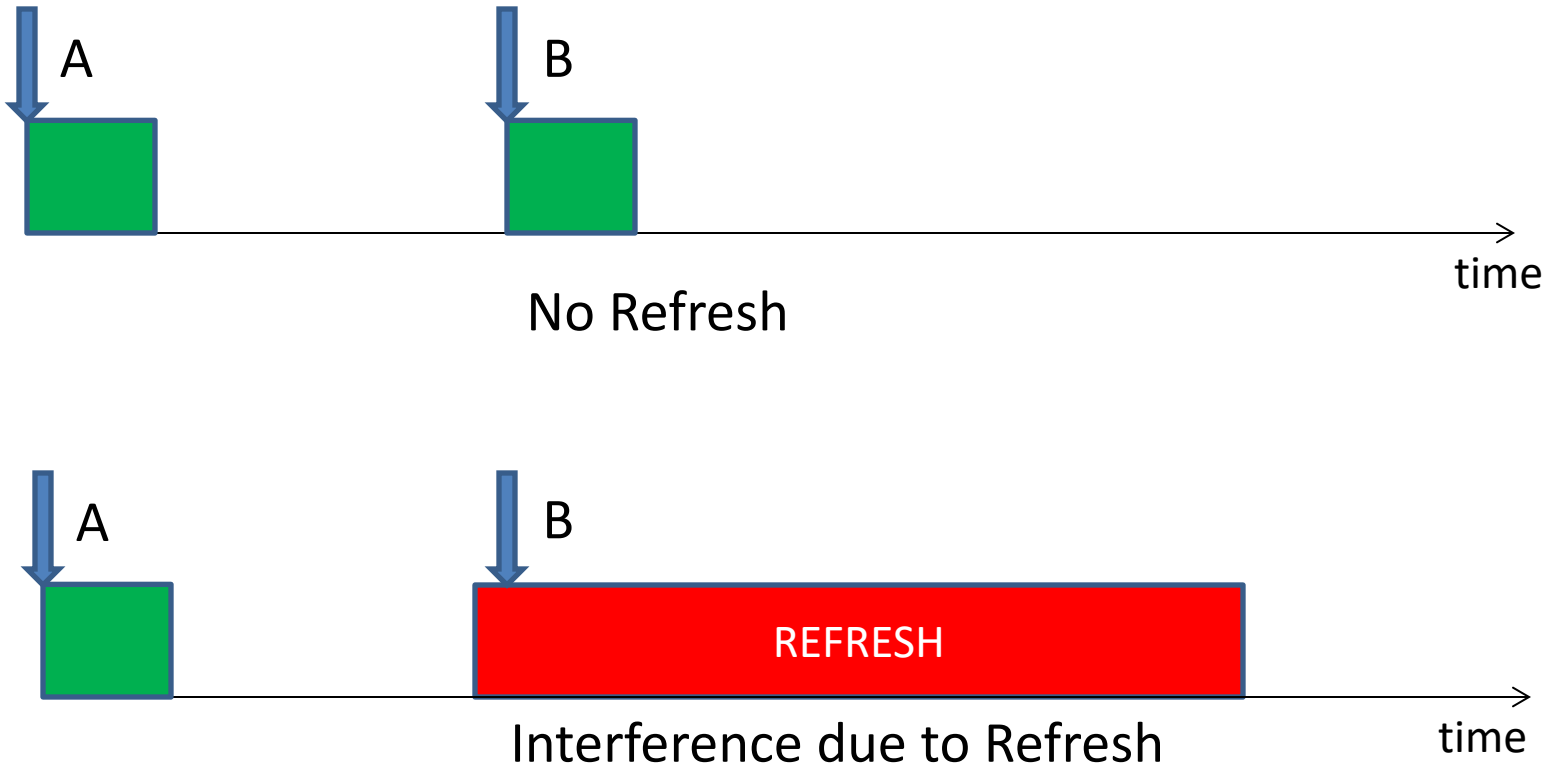

Refresh Blocks Reads

Memory unavailable for Read/Write during Refresh

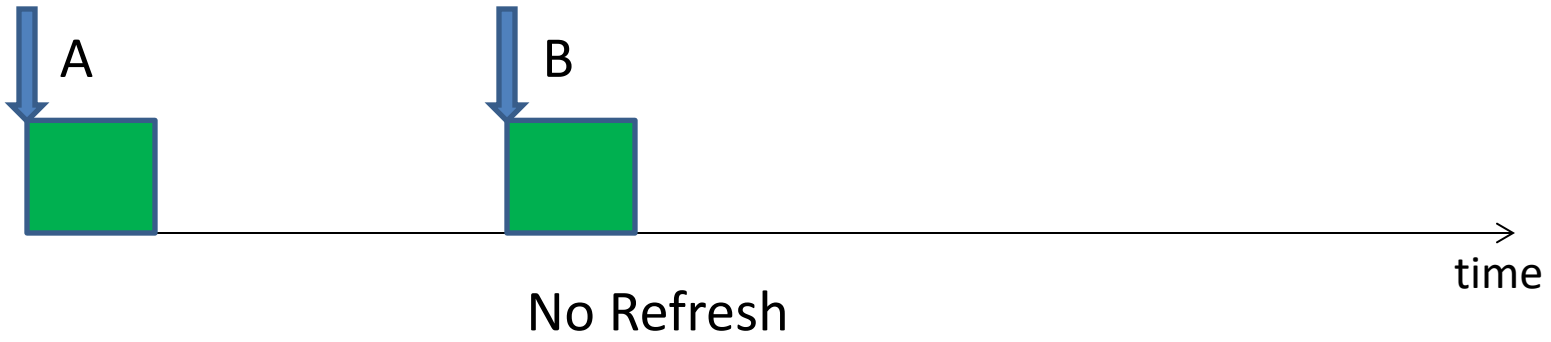

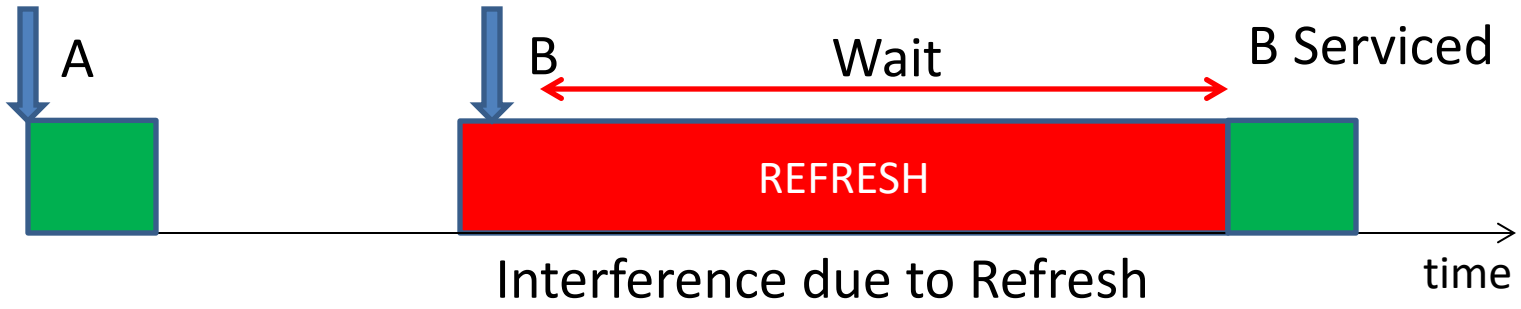

Refresh blocks reads ➔ Higher read latency

Impact of Refresh

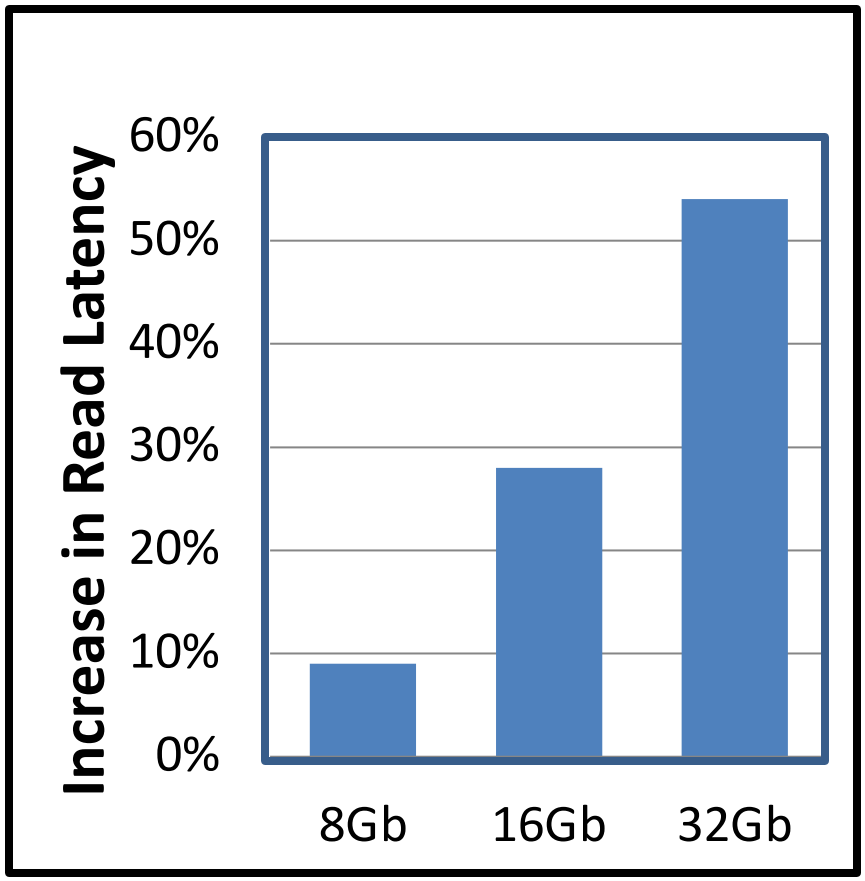

Impact of Refresh

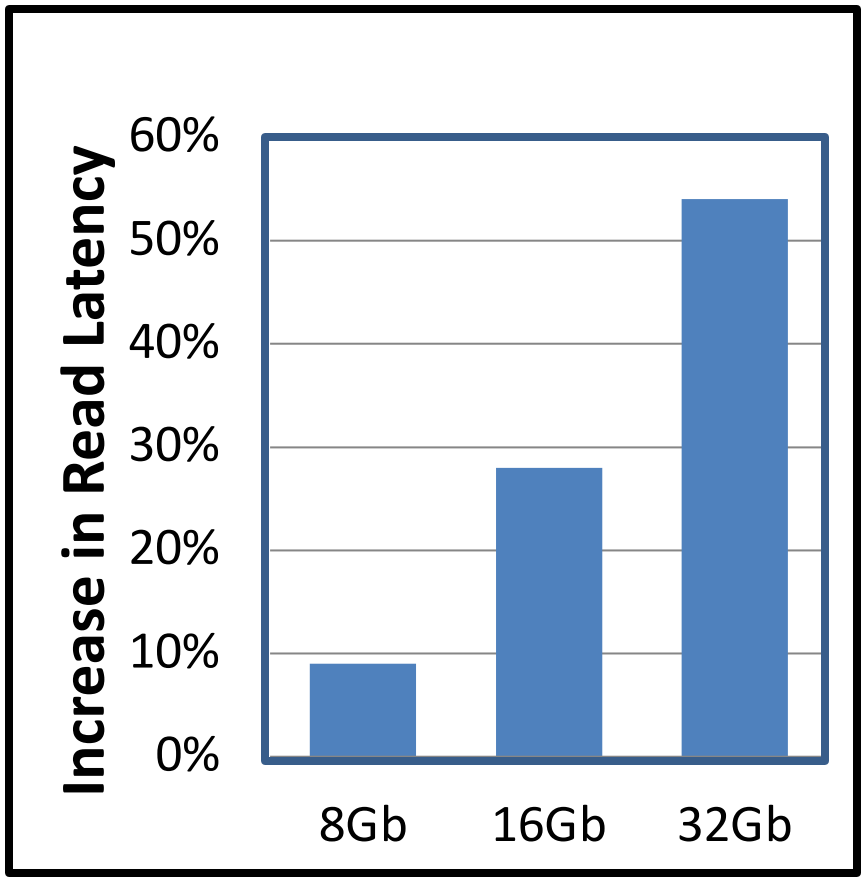

Impact of Refresh

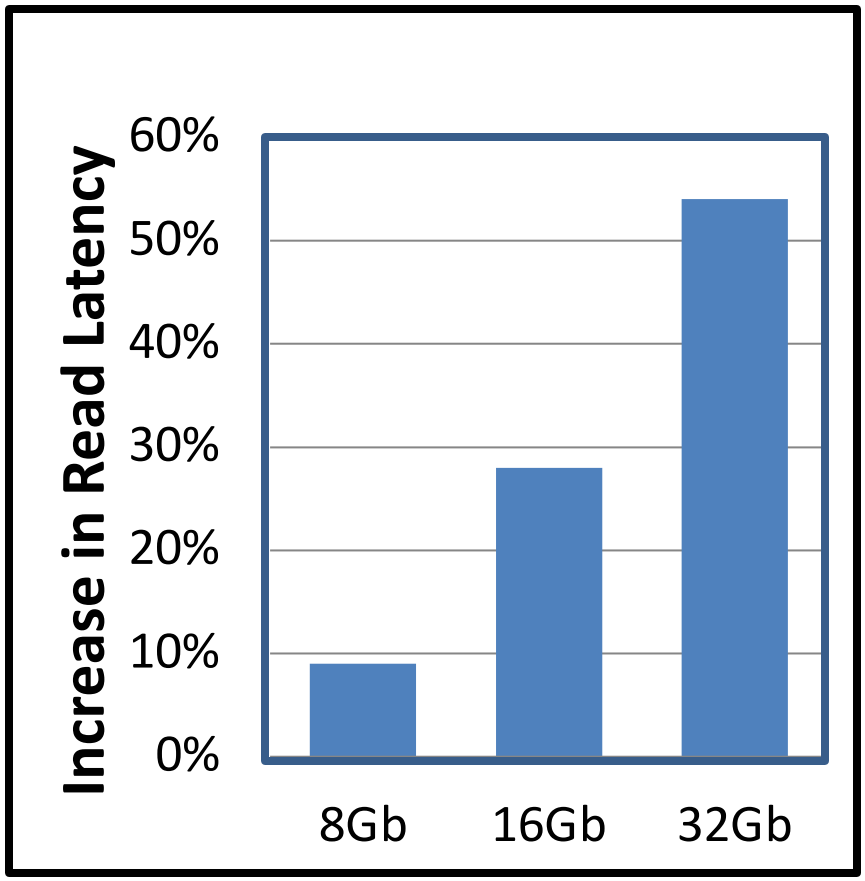

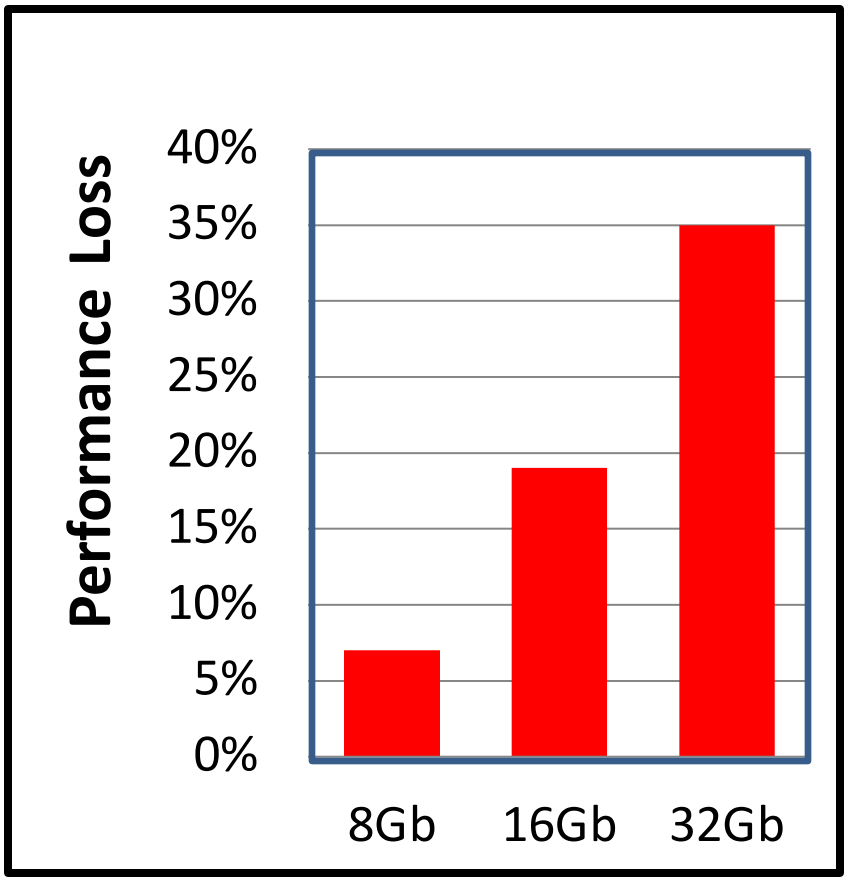

Impact of Refresh

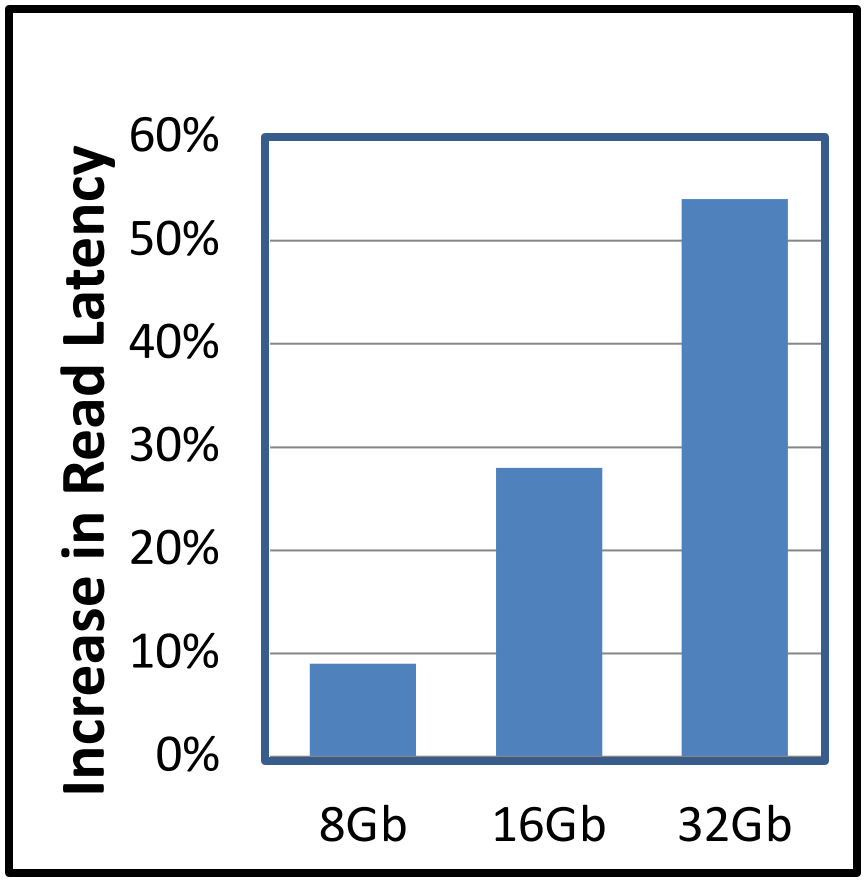

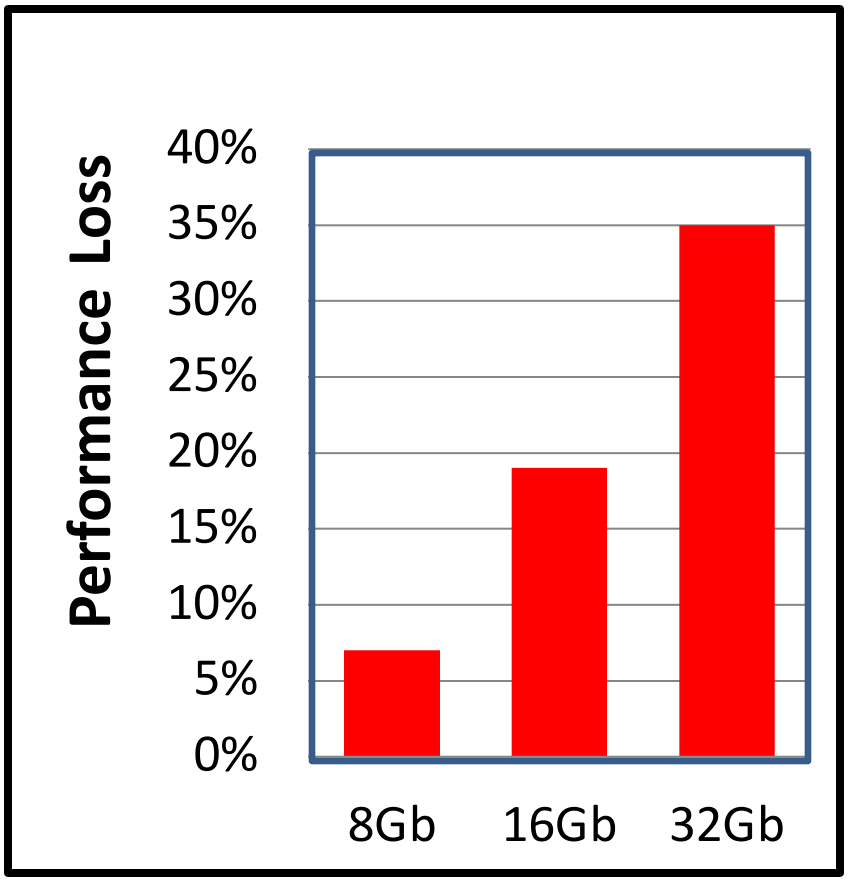

Impact of Refresh

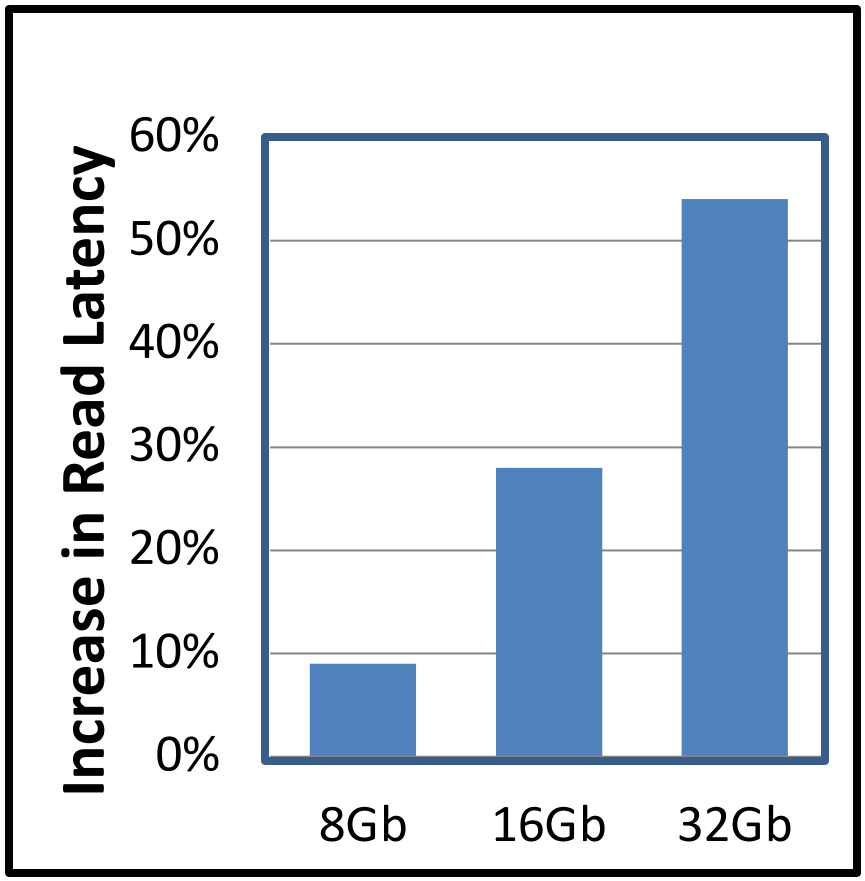

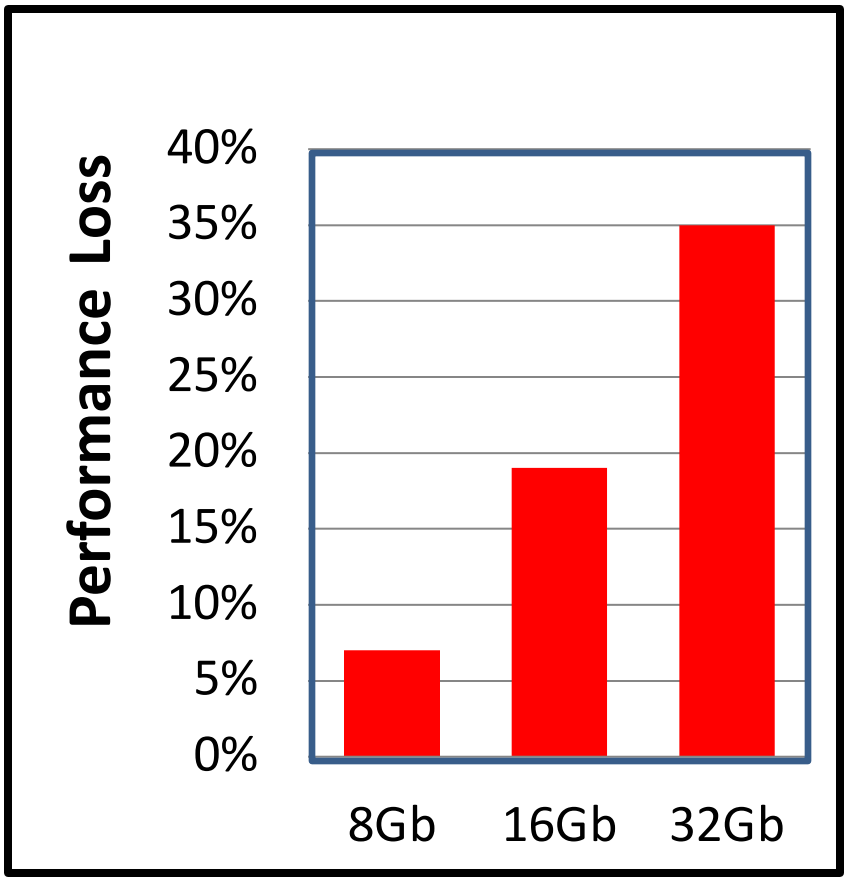

Impact of Refresh

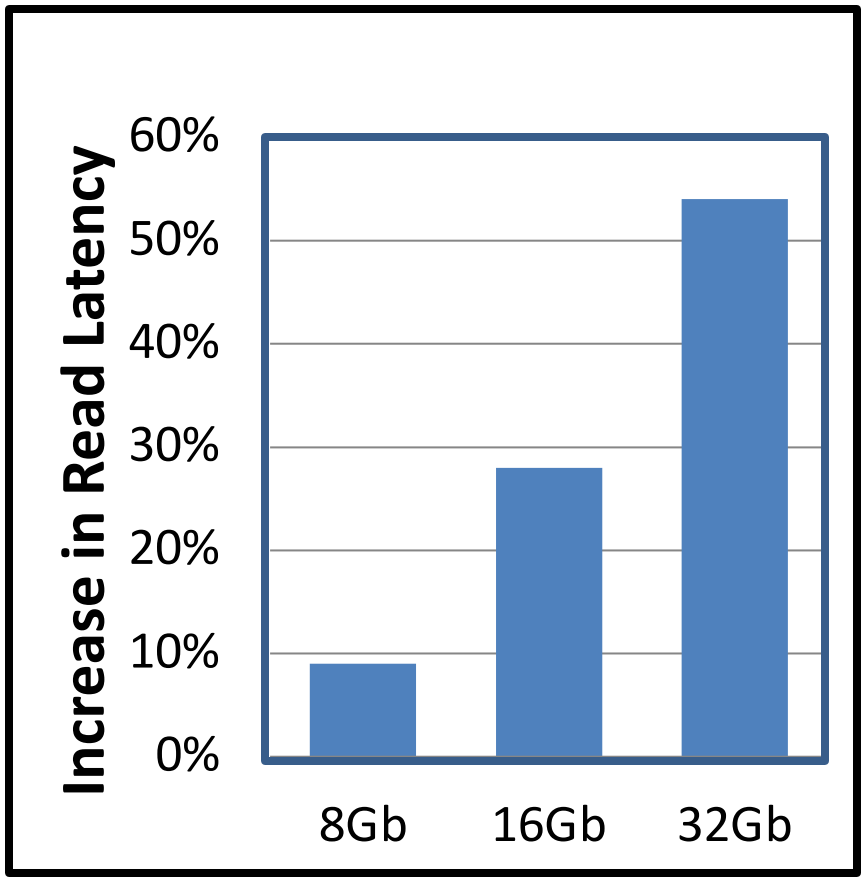

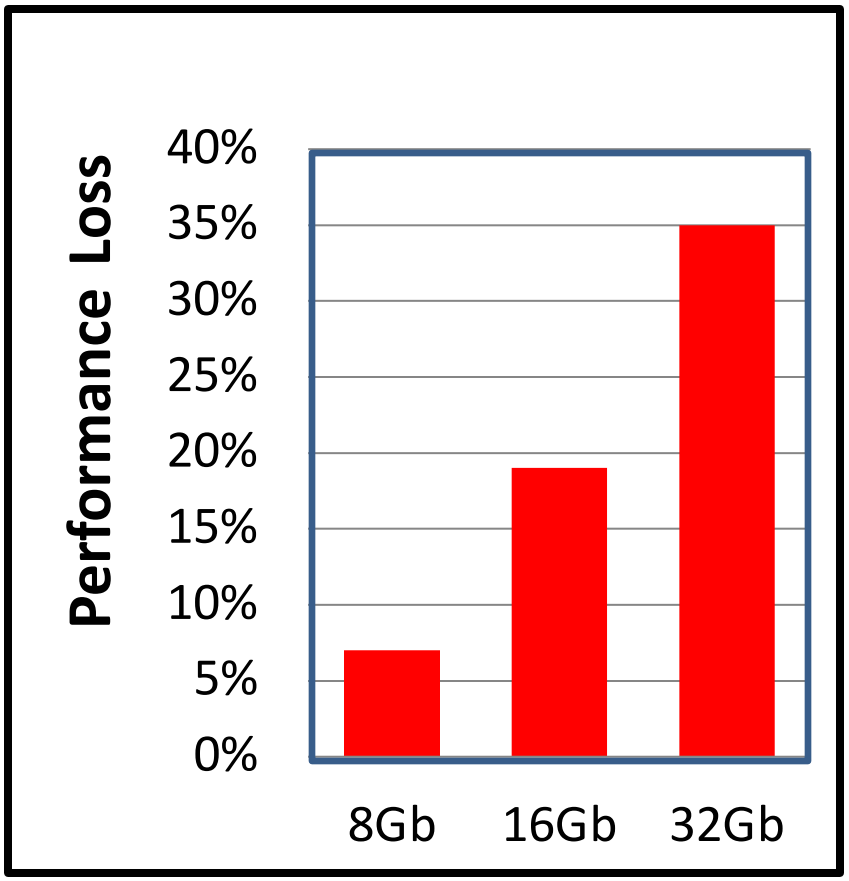

Impact of Refresh is significant, and increasing

Our Goal: Reduce the Read Latency impact of Refresh

Outline

- Introduction & Motivation
- Refresh Operation: Background
- Refresh Pausing
- Evaluation
- Alternative Proposals
- Summary

Refresh Operation

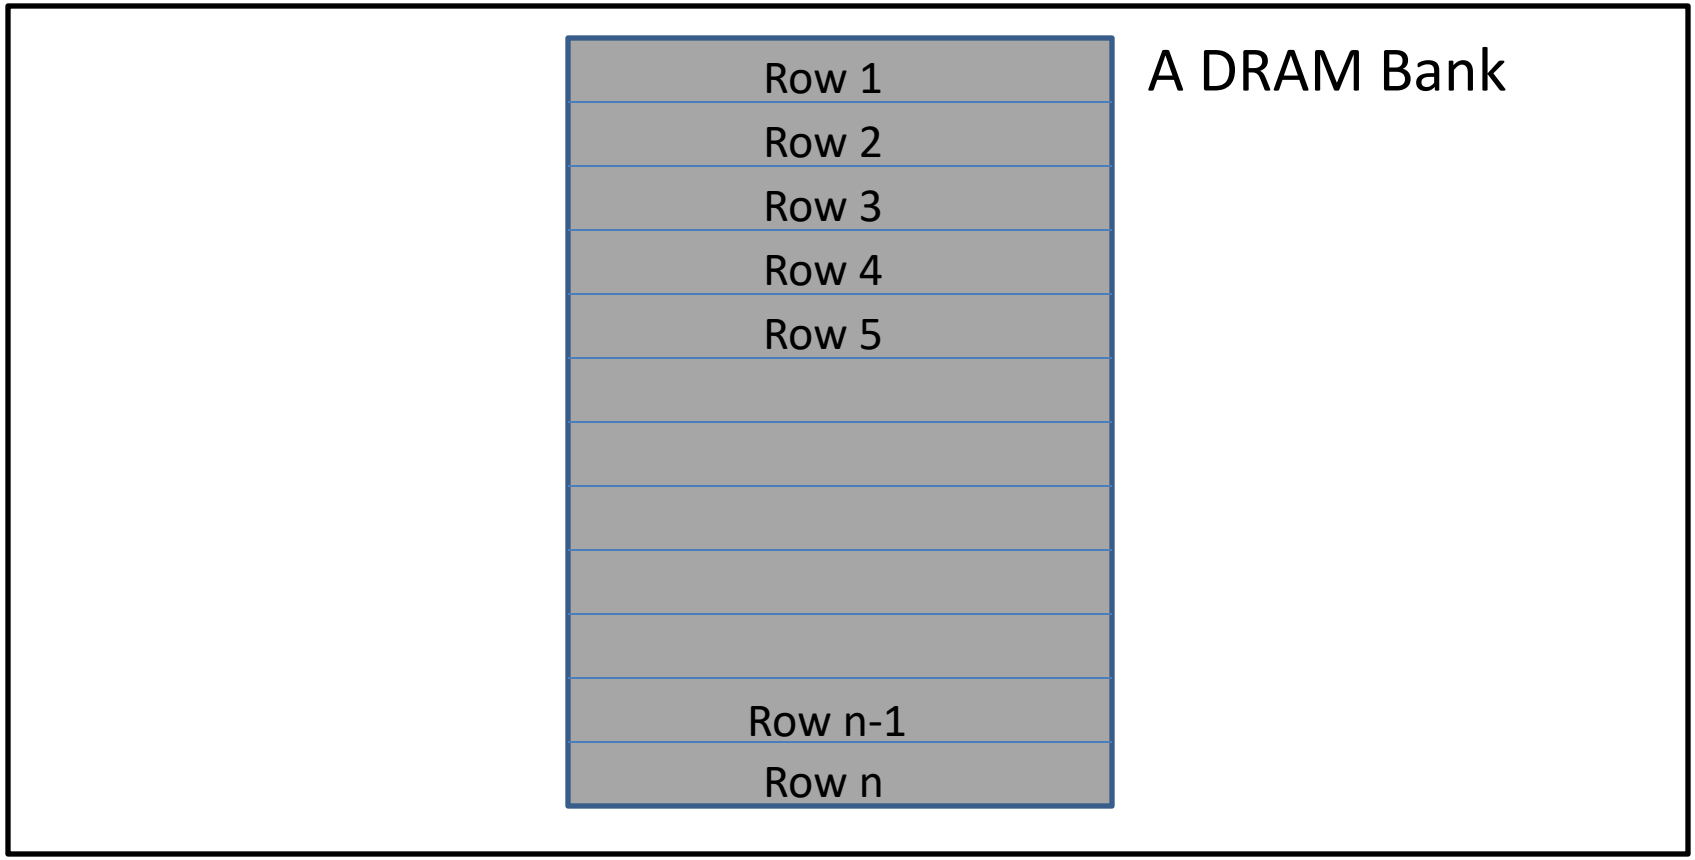

Refresh operates on a Row granularity

Refresh Operation

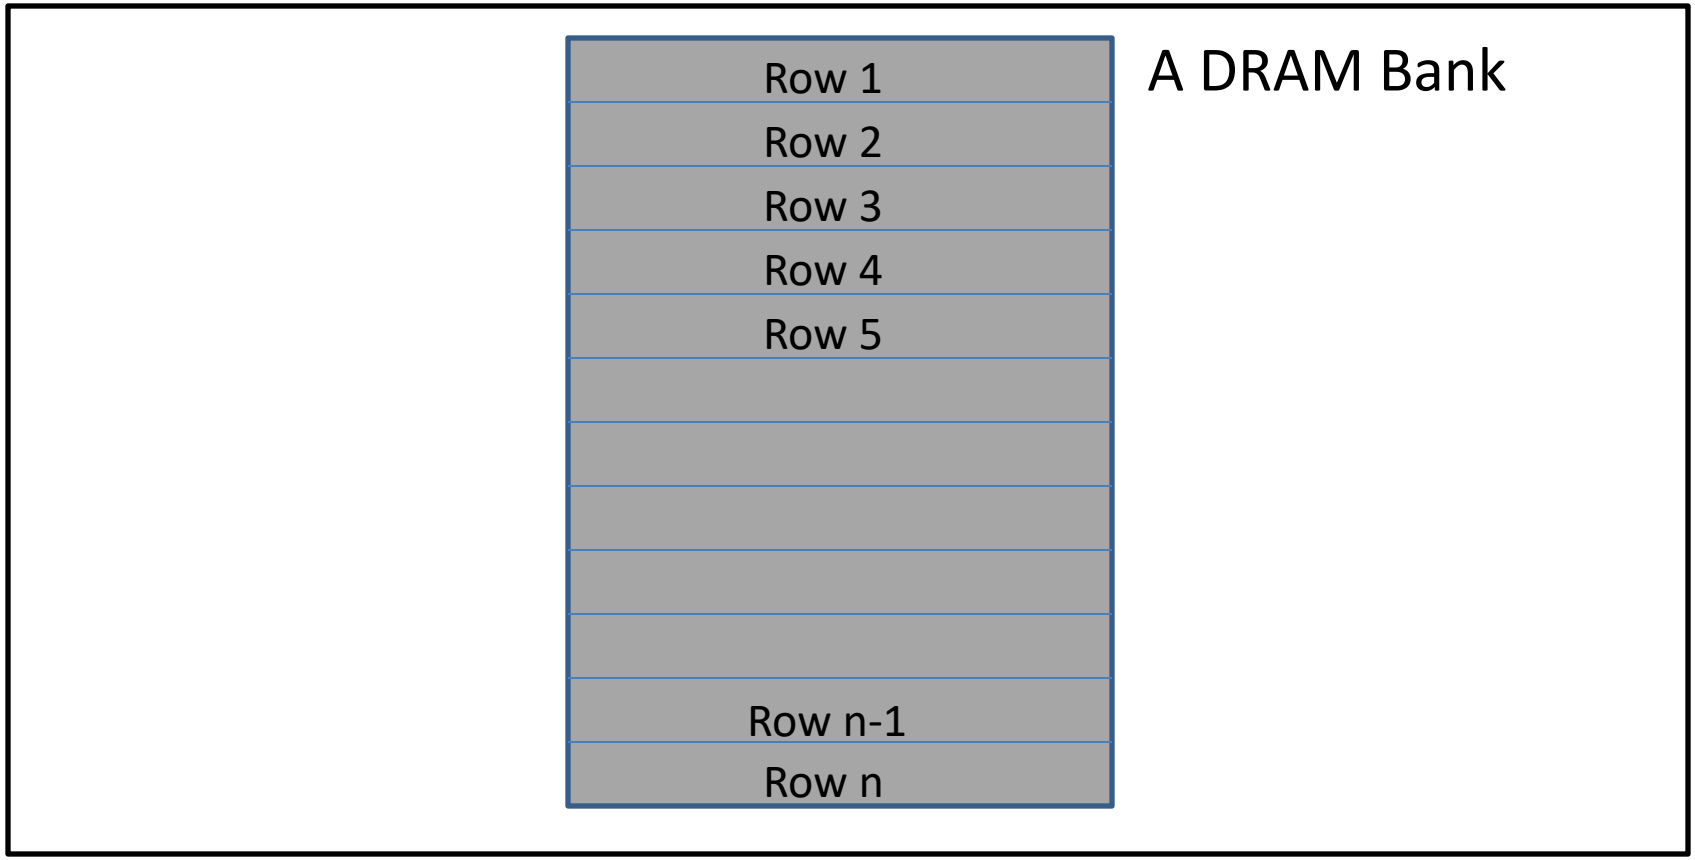

Refresh operates on a Row granularity

Refresh Modes

- Burst Mode:

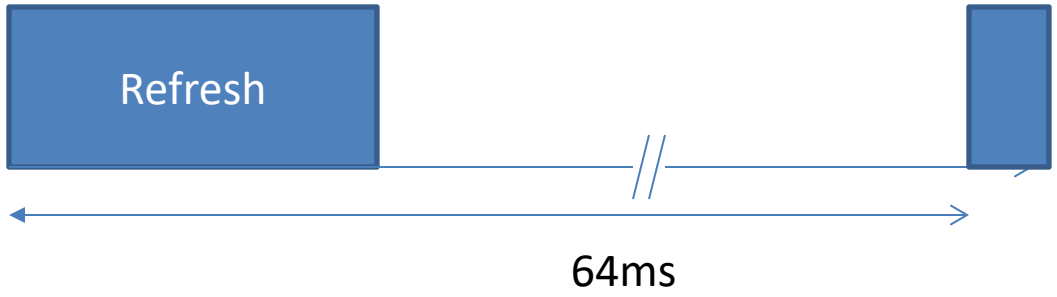

Memory unavailable until all rows finish refresh

- Distributed Mode:

Refresh Modes

- Burst Mode:

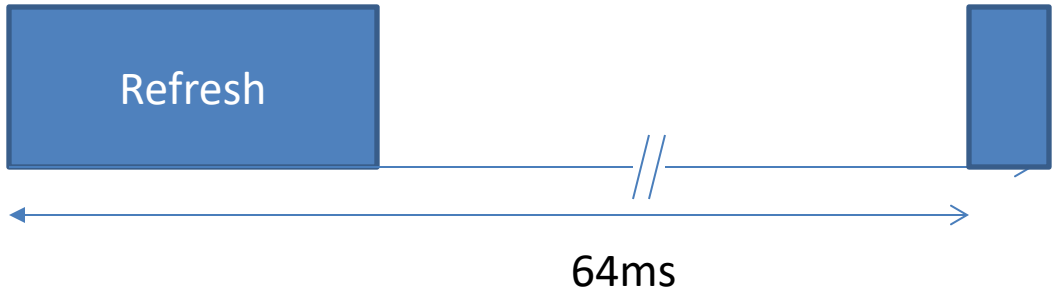

Memory unavailable until all rows finish refresh

- Distributed Mode:

8K refresh pulses in 64ms

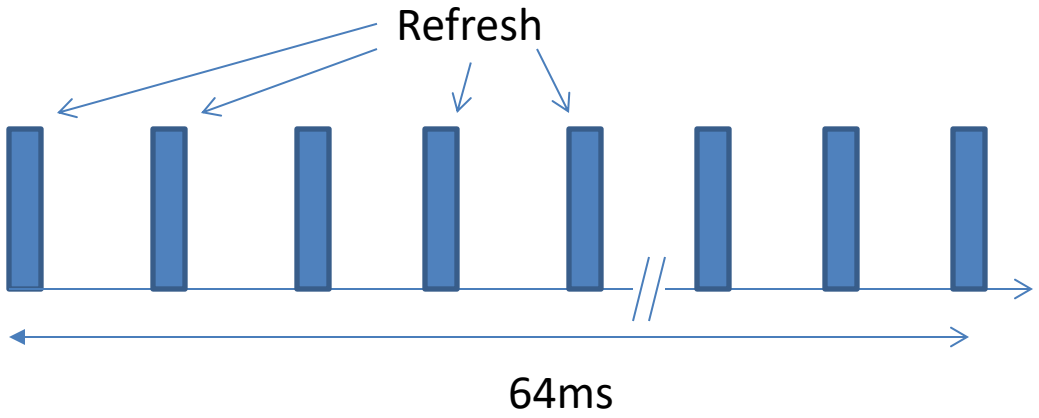

Refresh Modes

- Burst Mode:

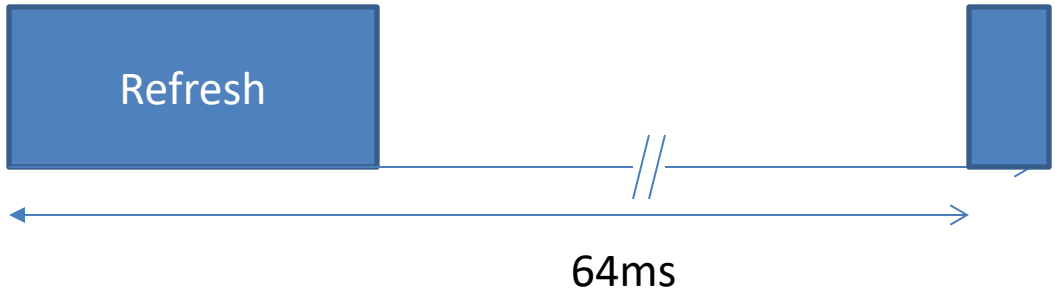

Memory unavailable until all rows finish refresh

- Distributed Mode:

8K refresh pulses in 64ms

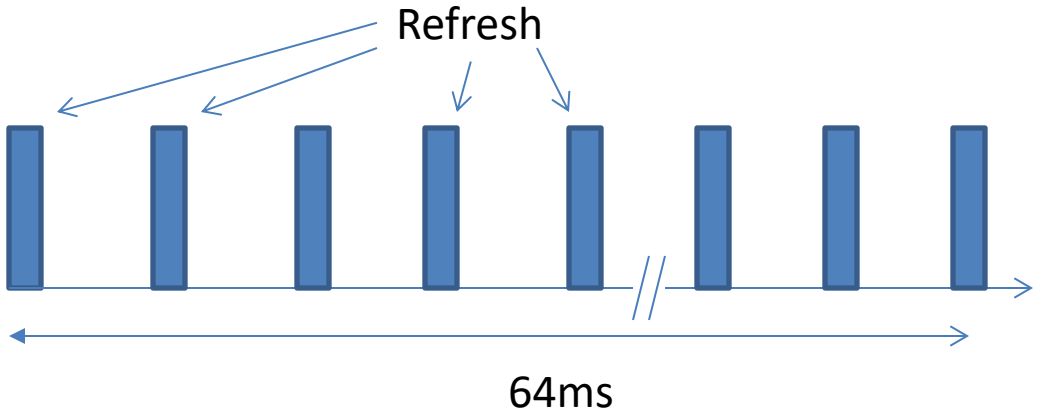

Distributed mode reduces contention from Refresh

Refresh Bundle

Every pulse refreshes a ‘Bundle of rows’

Chip Size	Rows in a Refresh bundle (per bank)
512 Mb	1
1Gb	2
2Gb	4
4Gb or 8Gb (Twin 4Gb die)	8

Refresh Bundle

Every pulse refreshes a ‘Bundle of rows’

Chip Size	Rows in a Refresh bundle (per bank)
512 Mb	1
1Gb	2
2Gb	4
4Gb or 8Gb (Twin 4Gb die)	8

Refresh Bundle currently have upto 8 rows, and increasing

The Latency Wall of Refresh

T_{RFC} is the time to do refresh for every refresh pulse

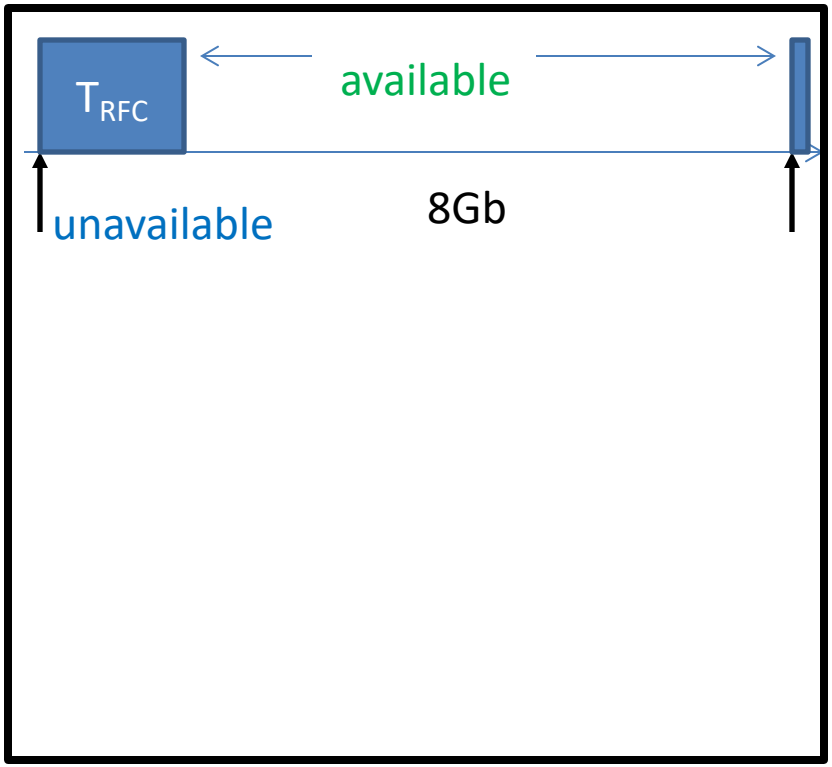

The Latency Wall of Refresh

T_{RFC} is the time to do refresh for every refresh pulse

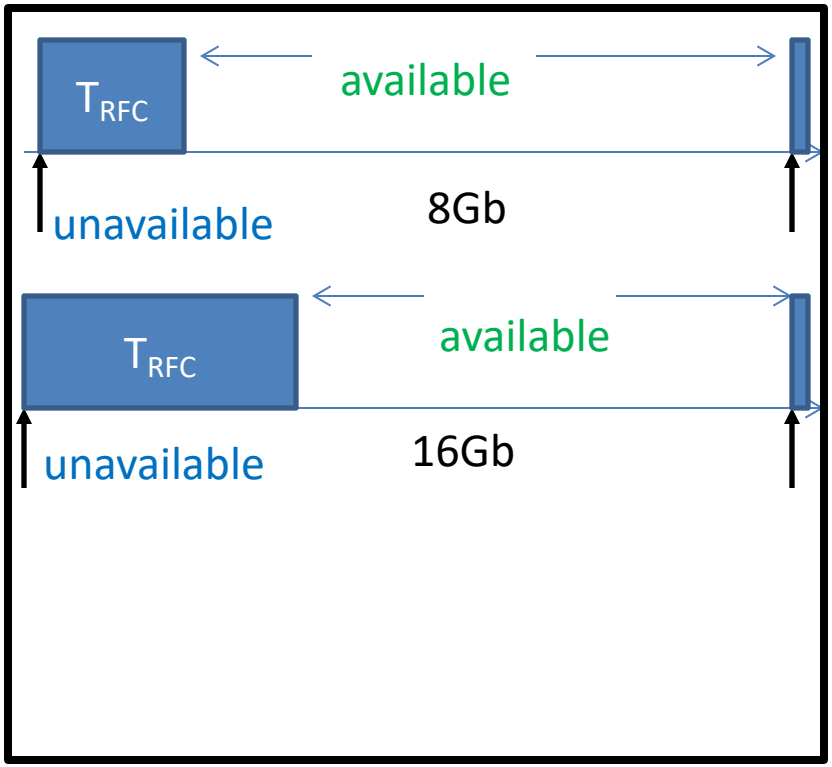

The Latency Wall of Refresh

T_{RFC} is the time to do refresh for every refresh pulse

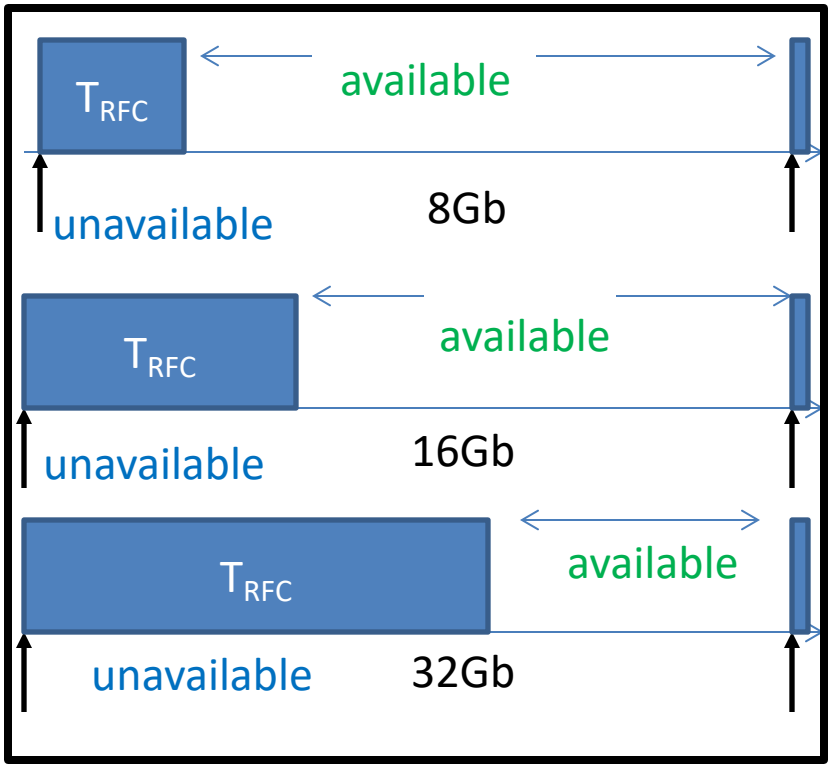

The Latency Wall of Refresh

T_{RFC} is the time to do refresh for every refresh pulse

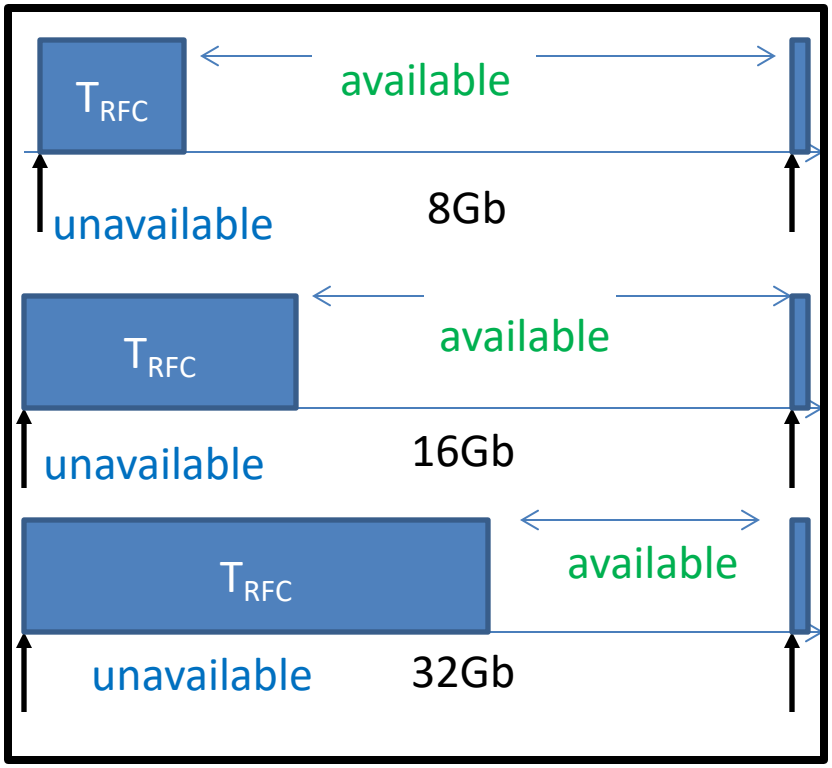

Current 8Gb chips have T_{RFC} of 350ns >> read latency

High T_{RFC} ➔ Read waits for refresh for long time

Outline

- Introduction & Motivation
- Refresh Operation: Background
- **Refresh Pausing**
- Evaluation
- Alternative Proposals
- Summary

Refresh Pausing

Insight: Make Refresh Operations Interruptible

Refresh Pausing

Insight: Make Refresh Operations Interruptible

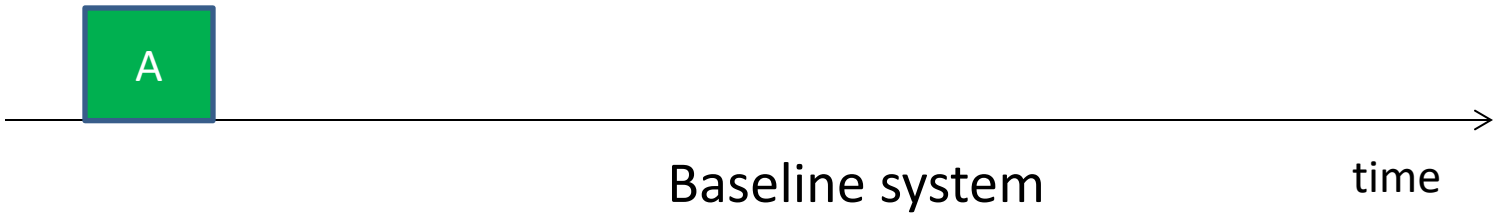

Refresh Pausing

Insight: Make Refresh Operations Interruptible

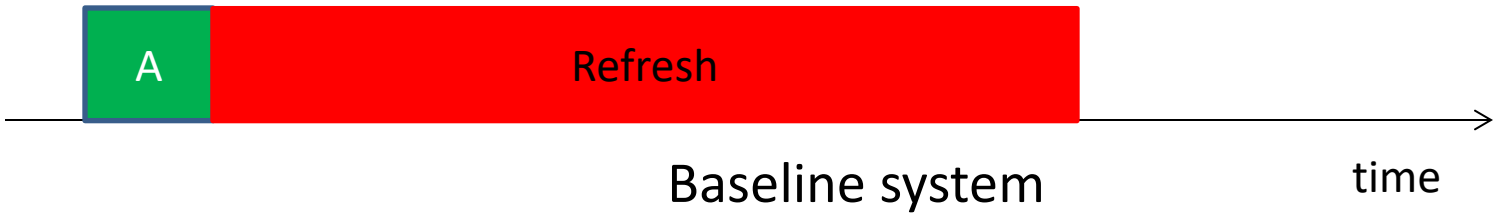

Refresh Pausing

Insight: Make Refresh Operations Interruptible

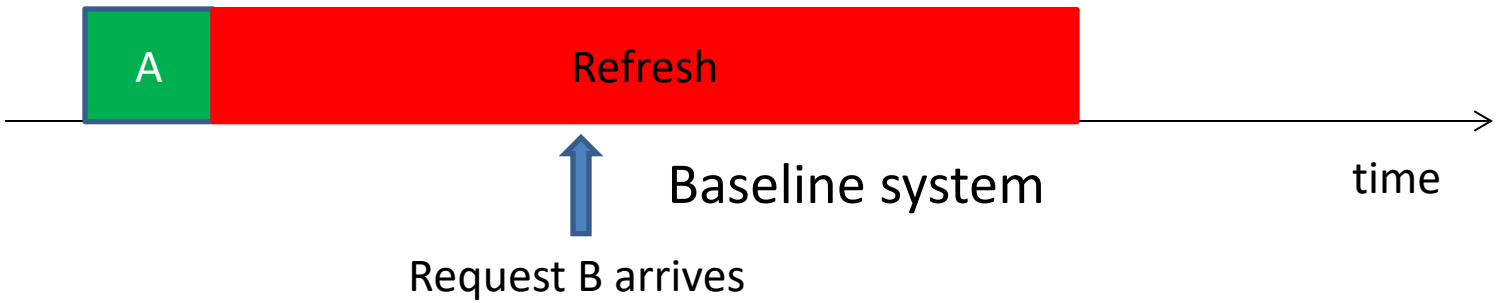

Refresh Pausing

Insight: Make Refresh Operations Interruptible

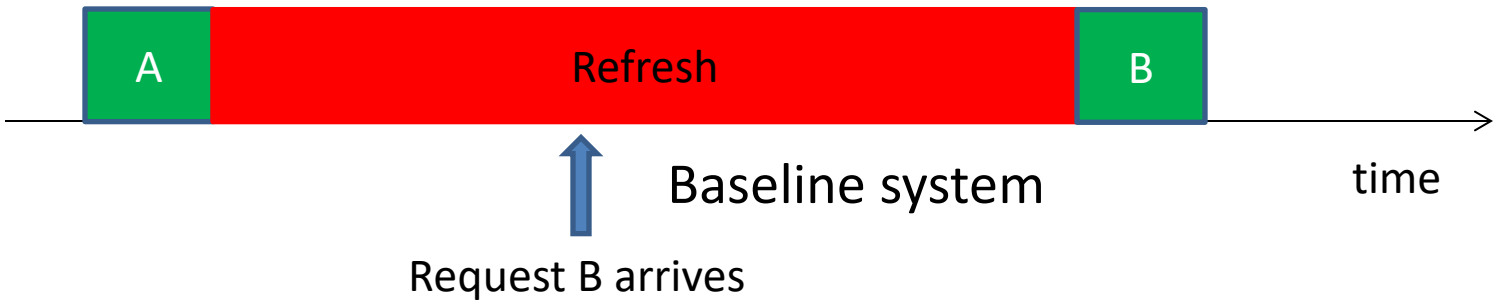

Refresh Pausing

Insight: Make Refresh Operations Interruptible

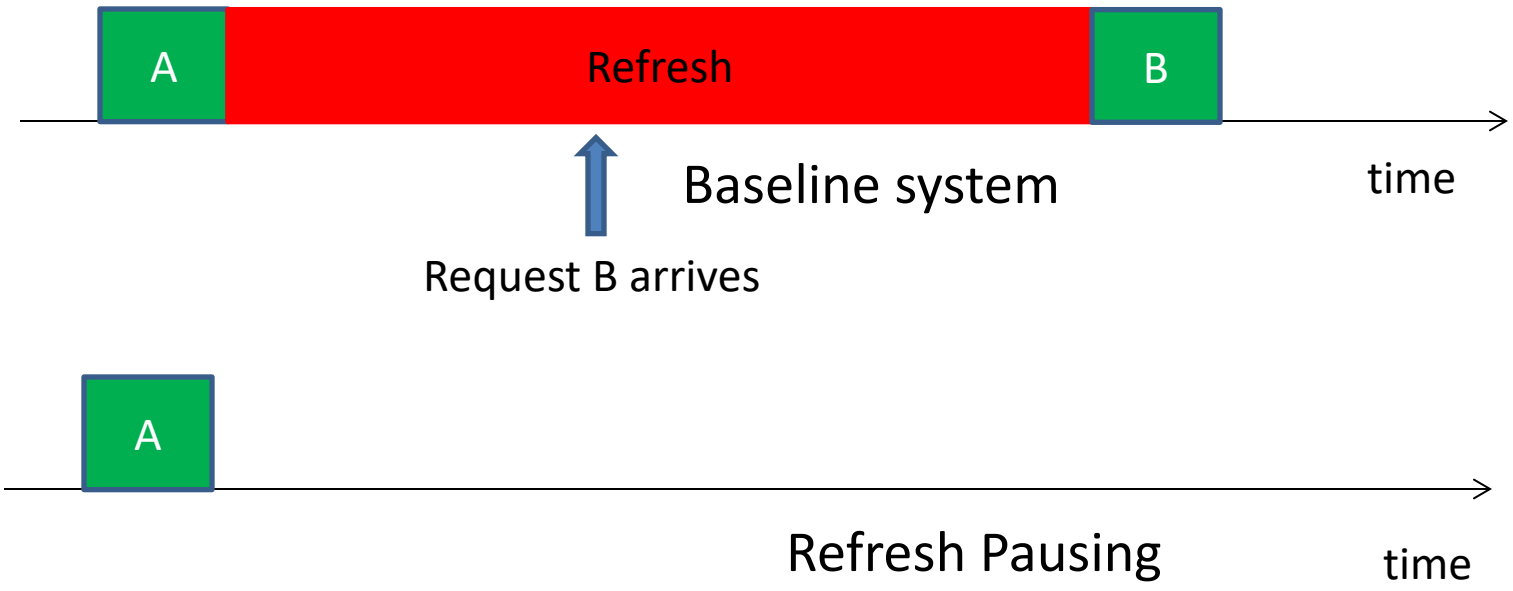

Refresh Pausing

Insight: Make Refresh Operations Interruptible

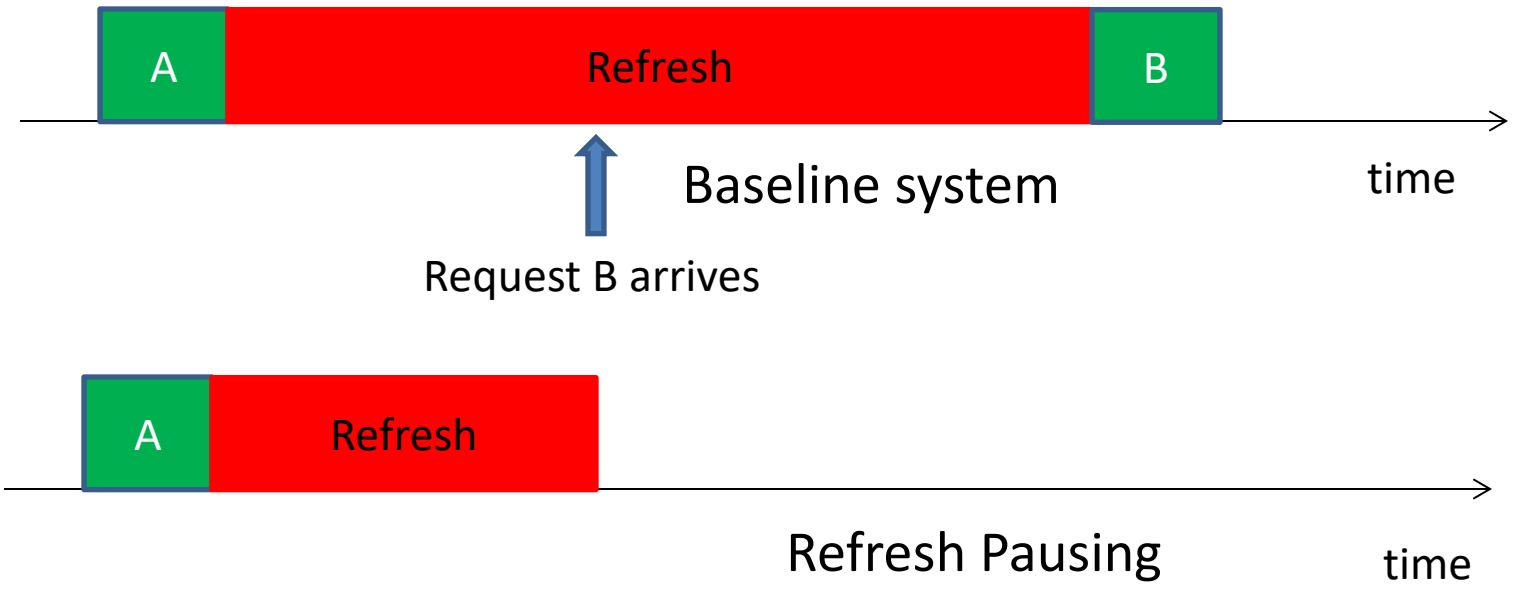

Refresh Pausing

Insight: Make Refresh Operations Interruptible

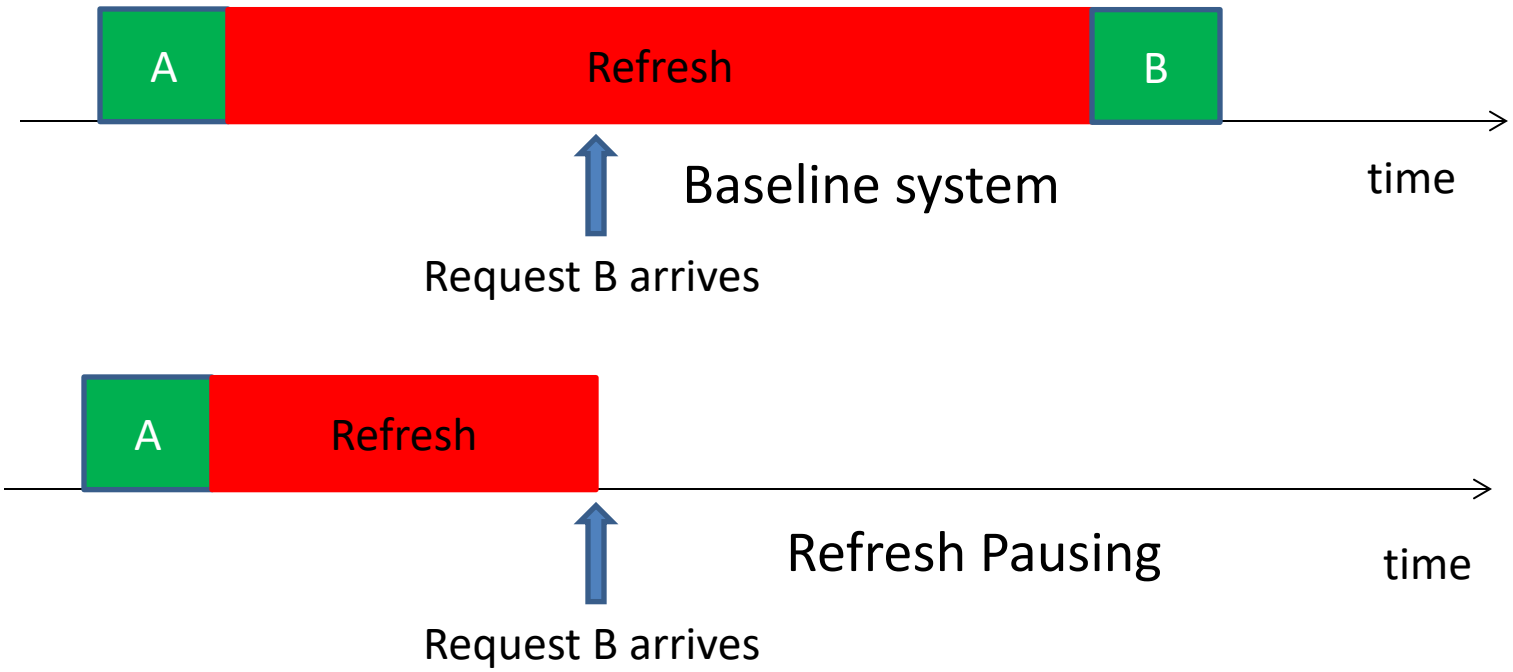

Refresh Pausing

Insight: Make Refresh Operations Interruptible

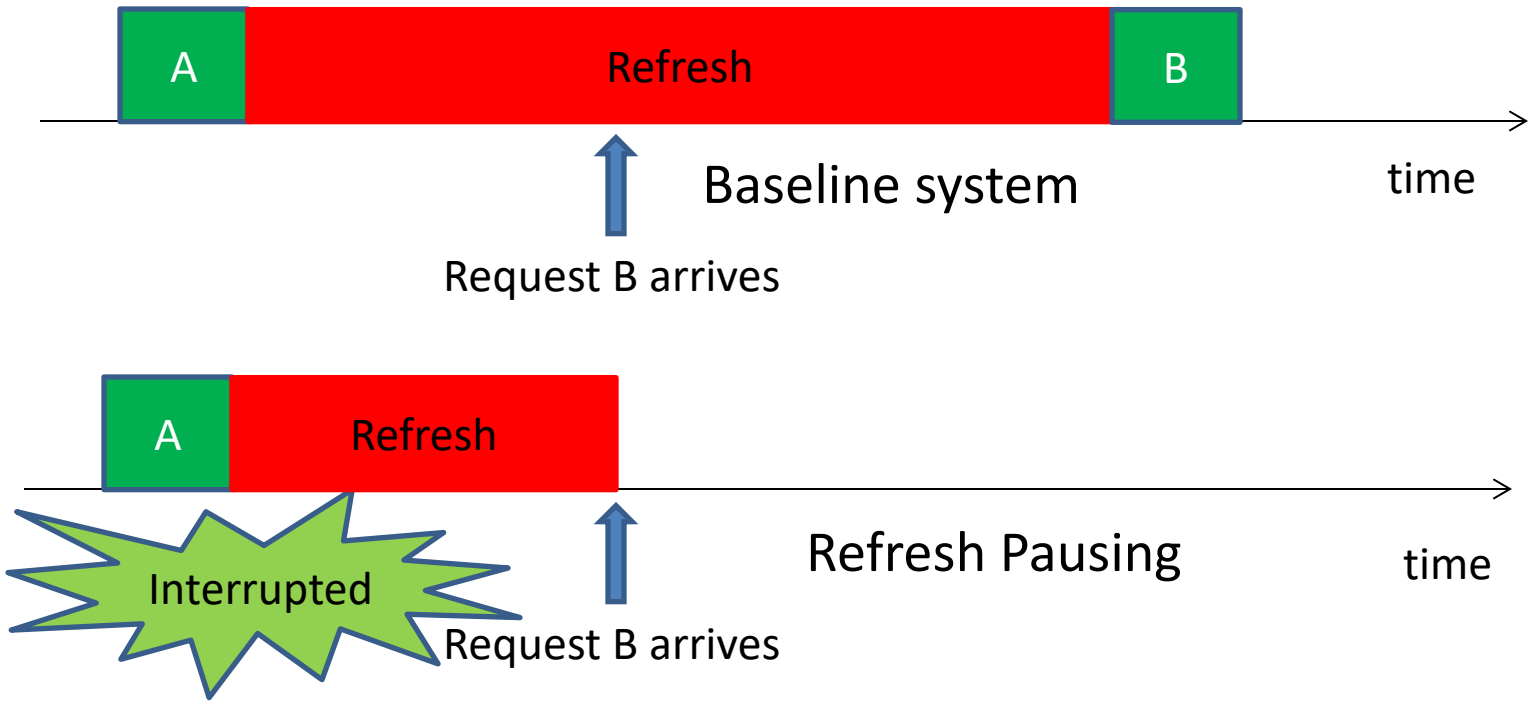

Refresh Pausing

Insight: Make Refresh Operations Interruptible

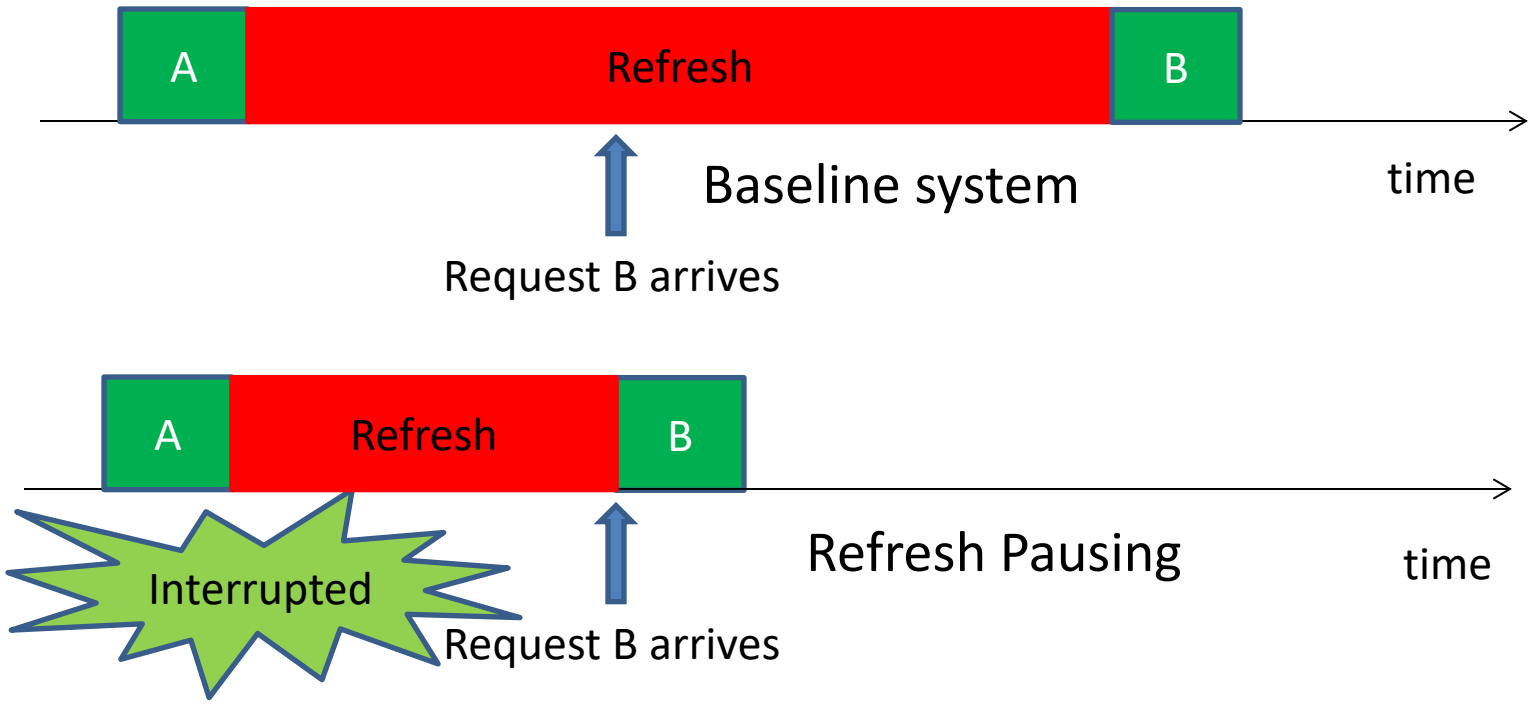

Refresh Pausing

Insight: Make Refresh Operations Interruptible

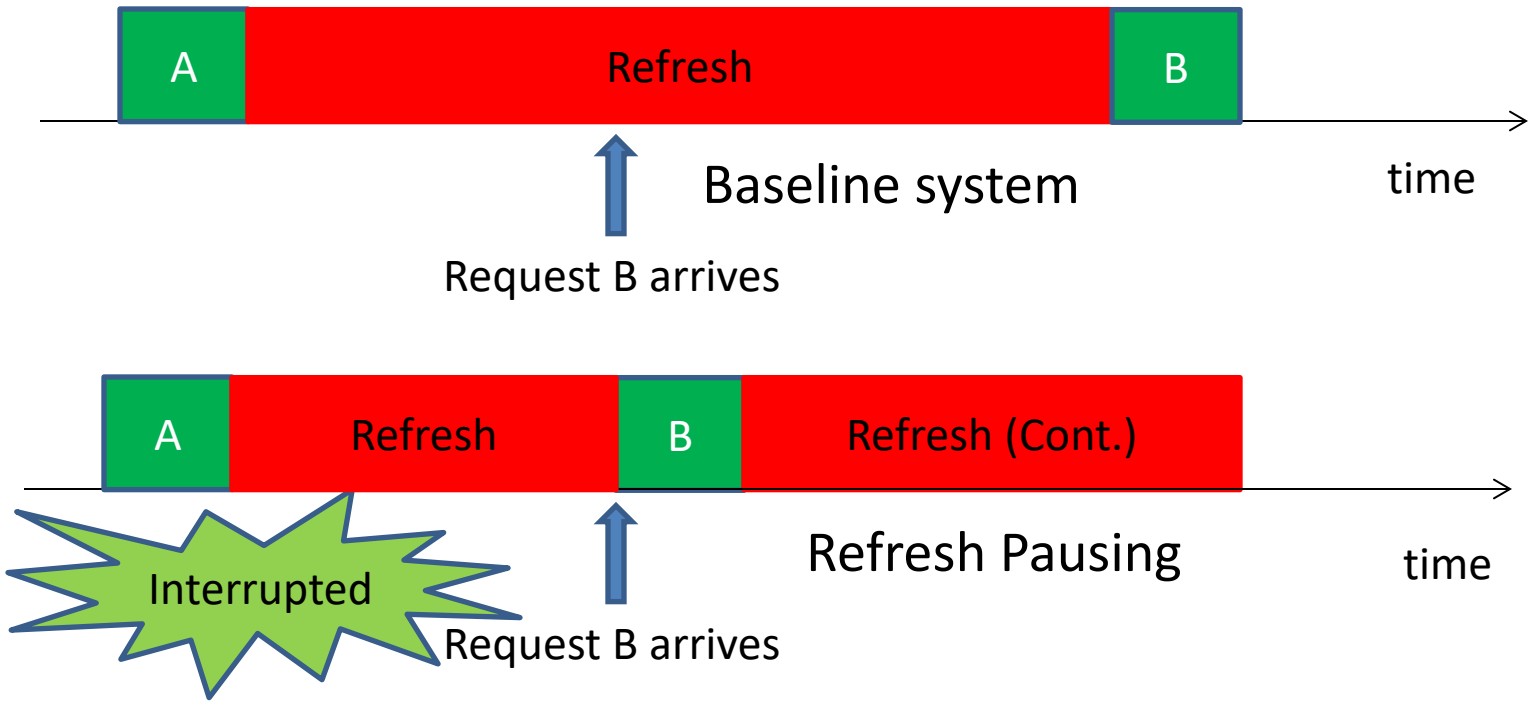

Refresh Pausing

Insight: Make Refresh Operations Interruptible

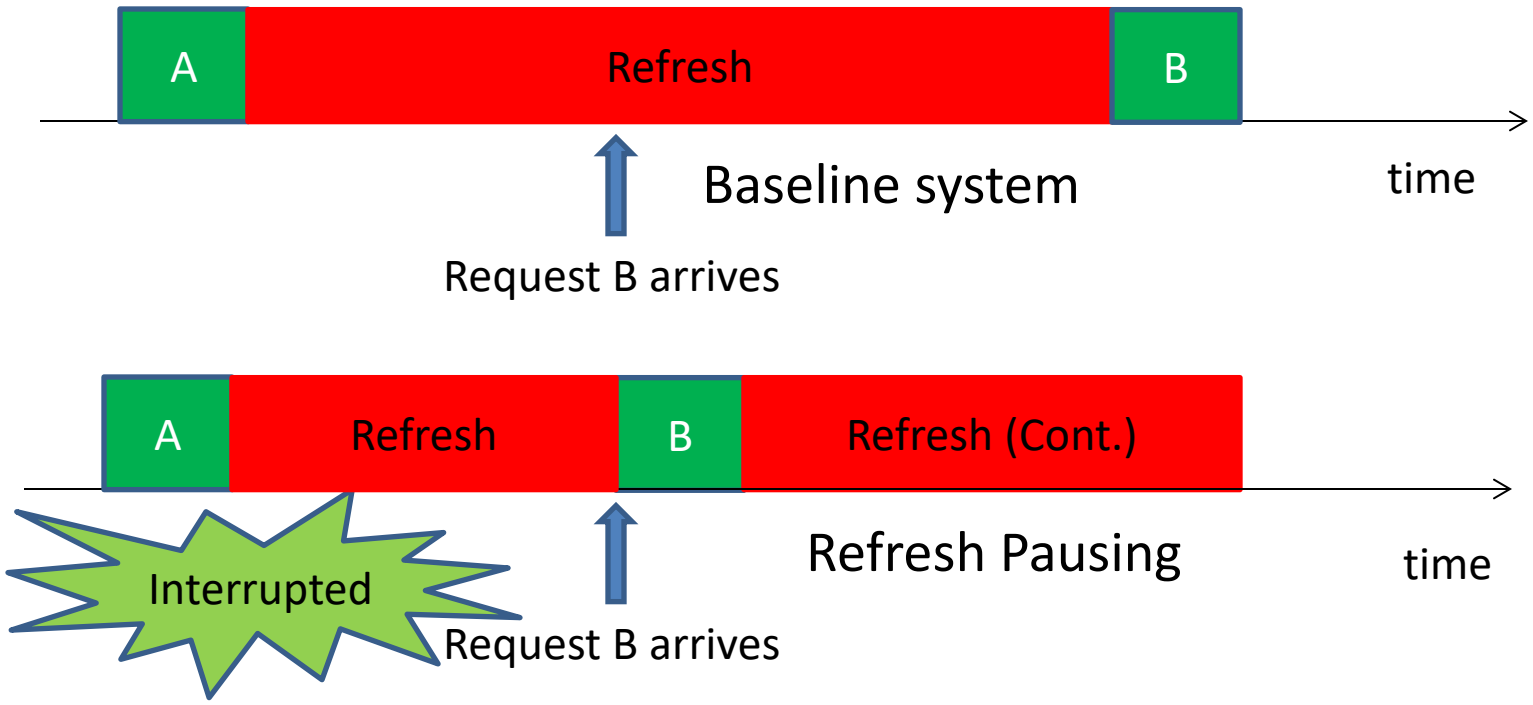

Pausing Refresh reduces wait time for Reads

Refresh Pausing

Insight: Make Refresh Operations Interruptible

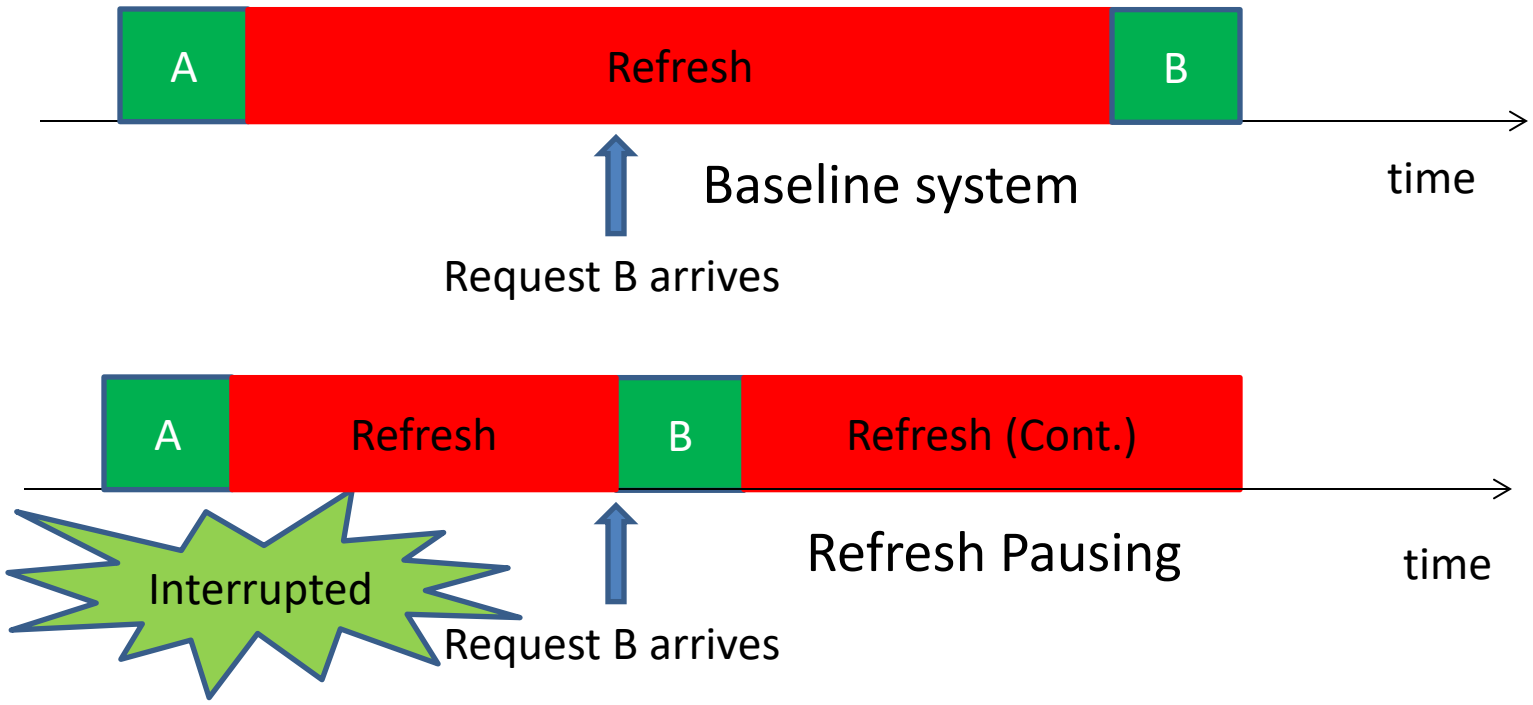

Pausing at arbitrary point can cause data loss

Pausing Refresh reduces wait time for Reads

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

Without Refresh Pausing

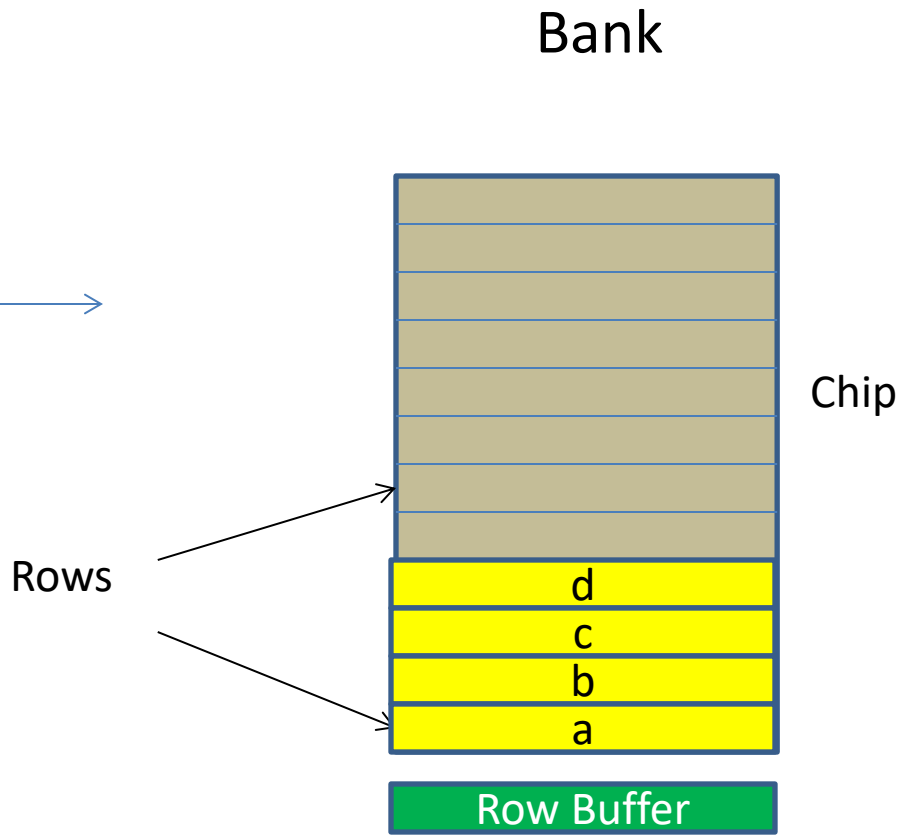

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

Without Refresh Pausing

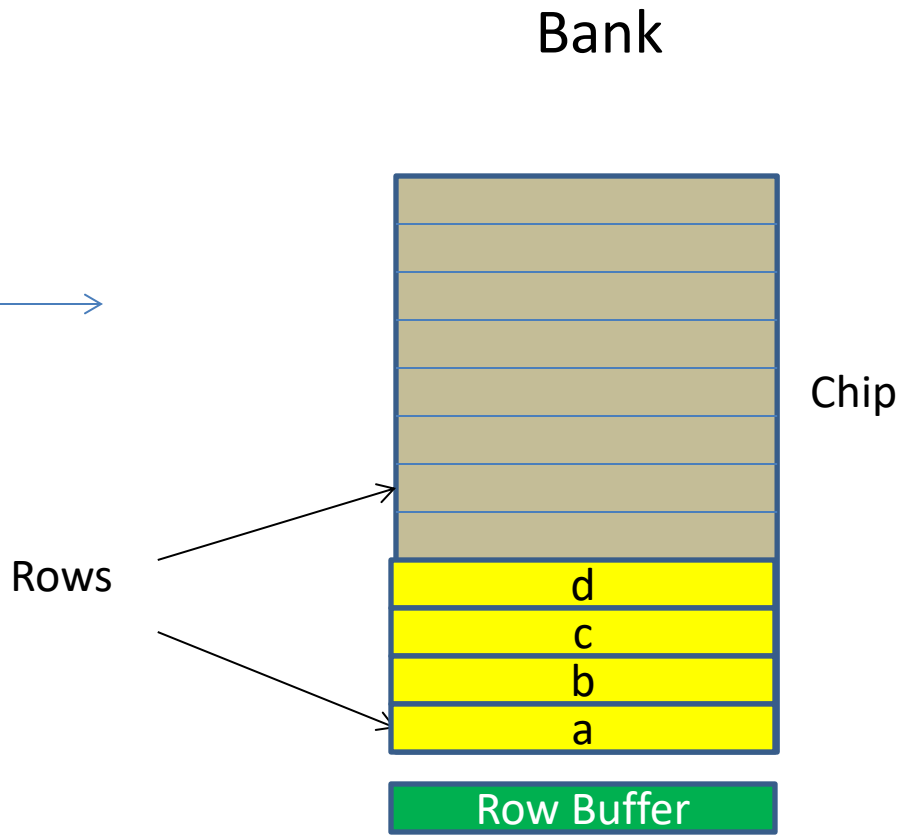

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

Without Refresh Pausing

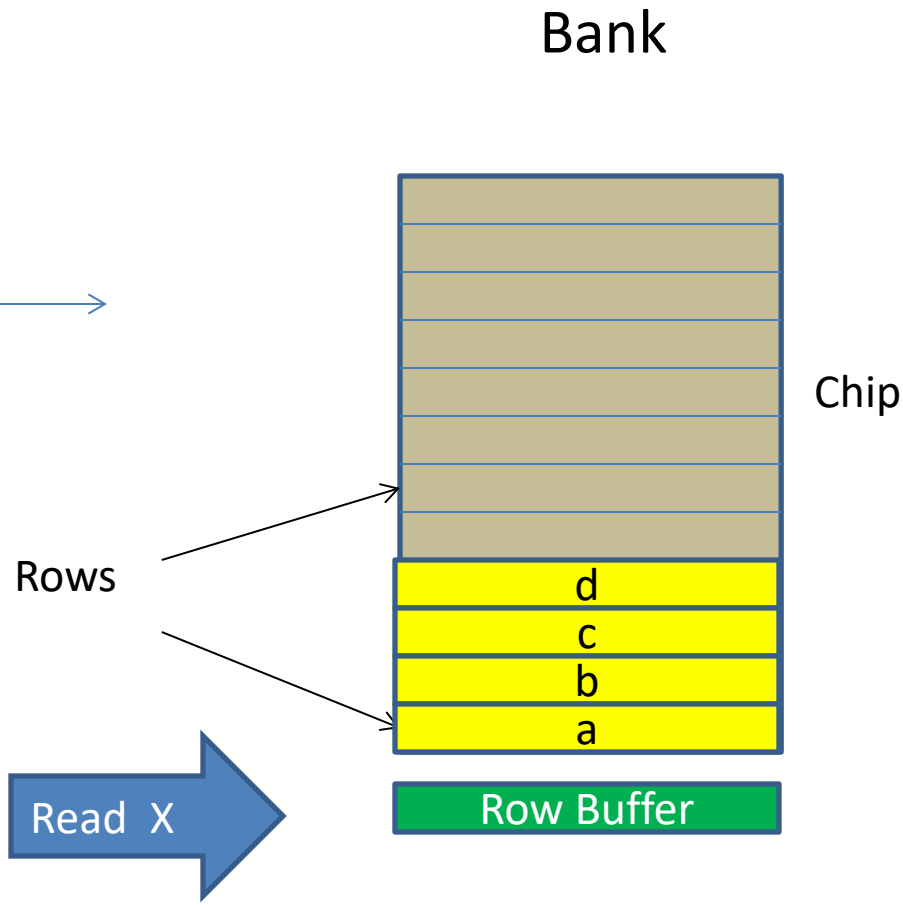

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

Without Refresh Pausing

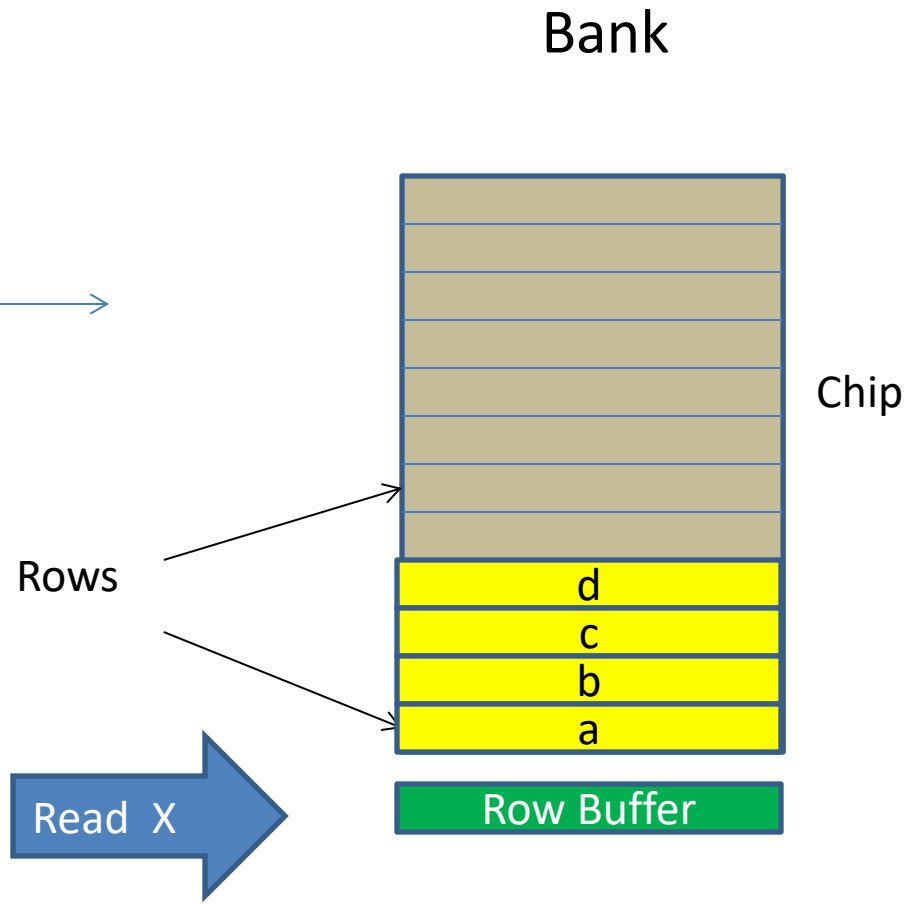

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

Without Refresh Pausing

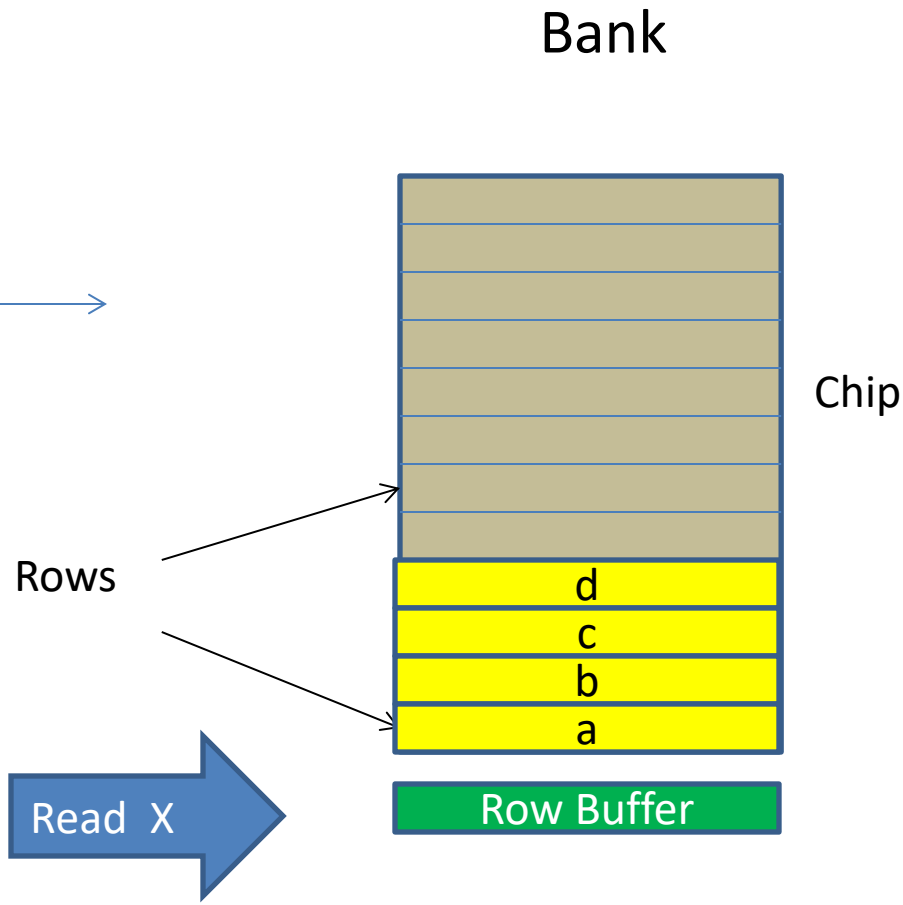

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

Without Refresh Pausing

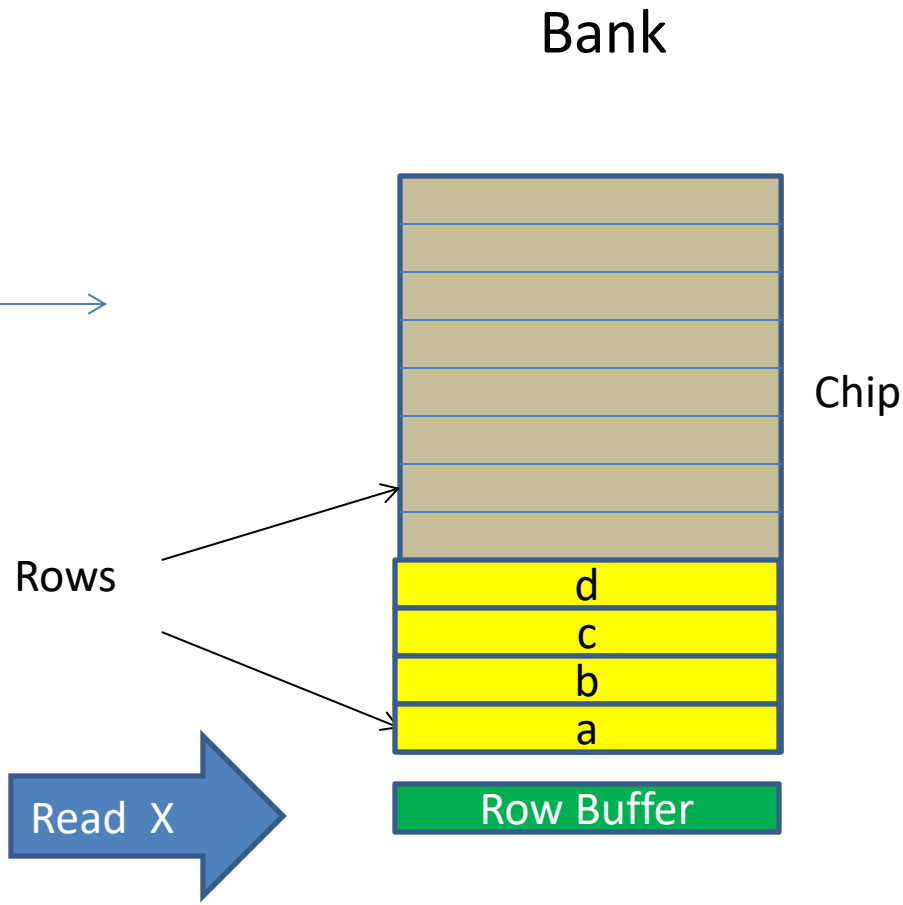

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

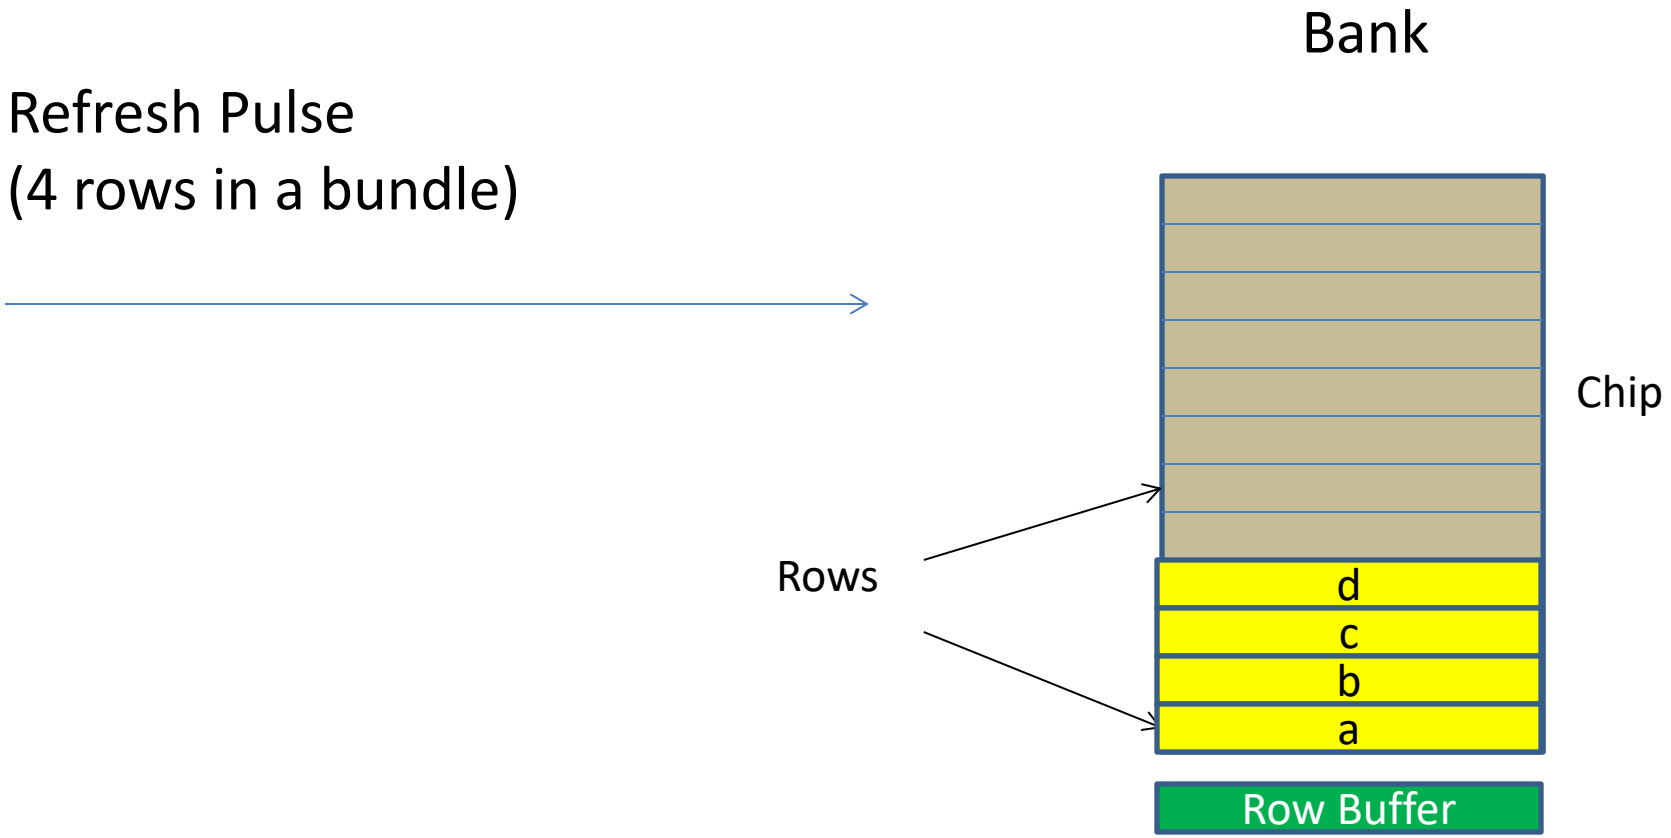

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

With Refresh Pausing

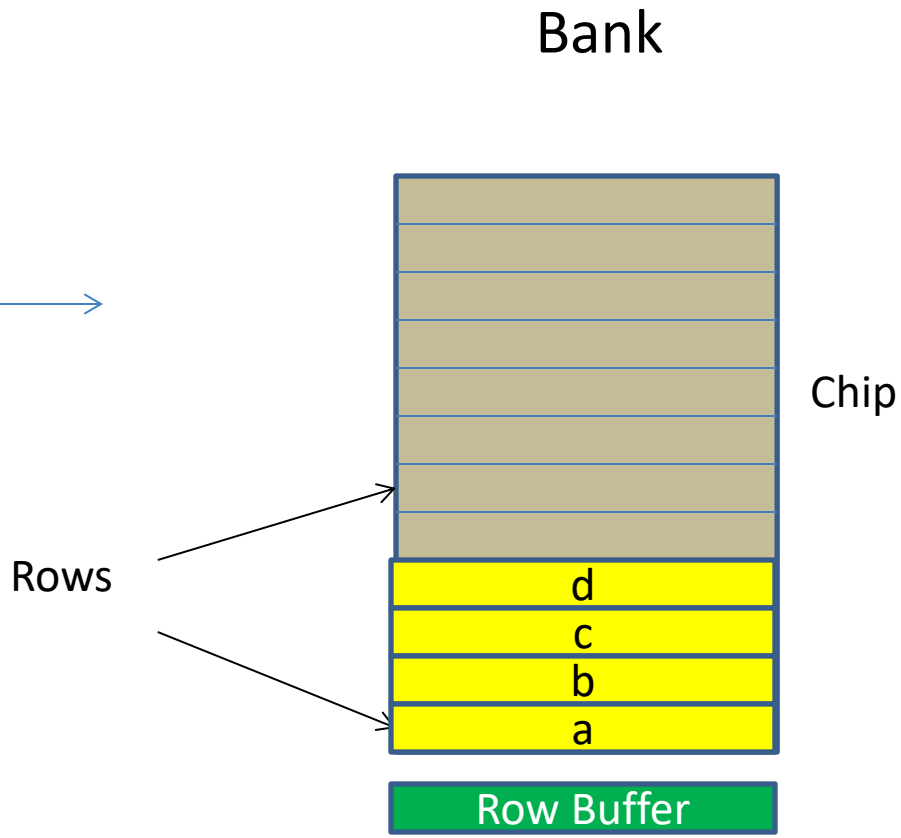

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

With Refresh Pausing

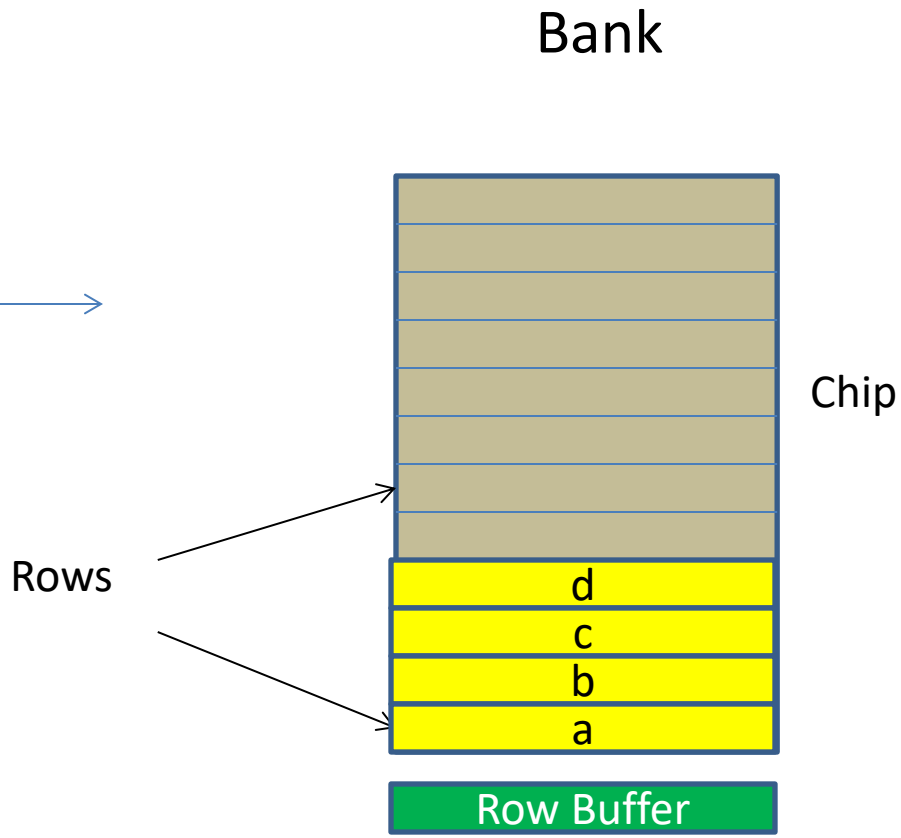

Refresh Pausing: When to Pause?

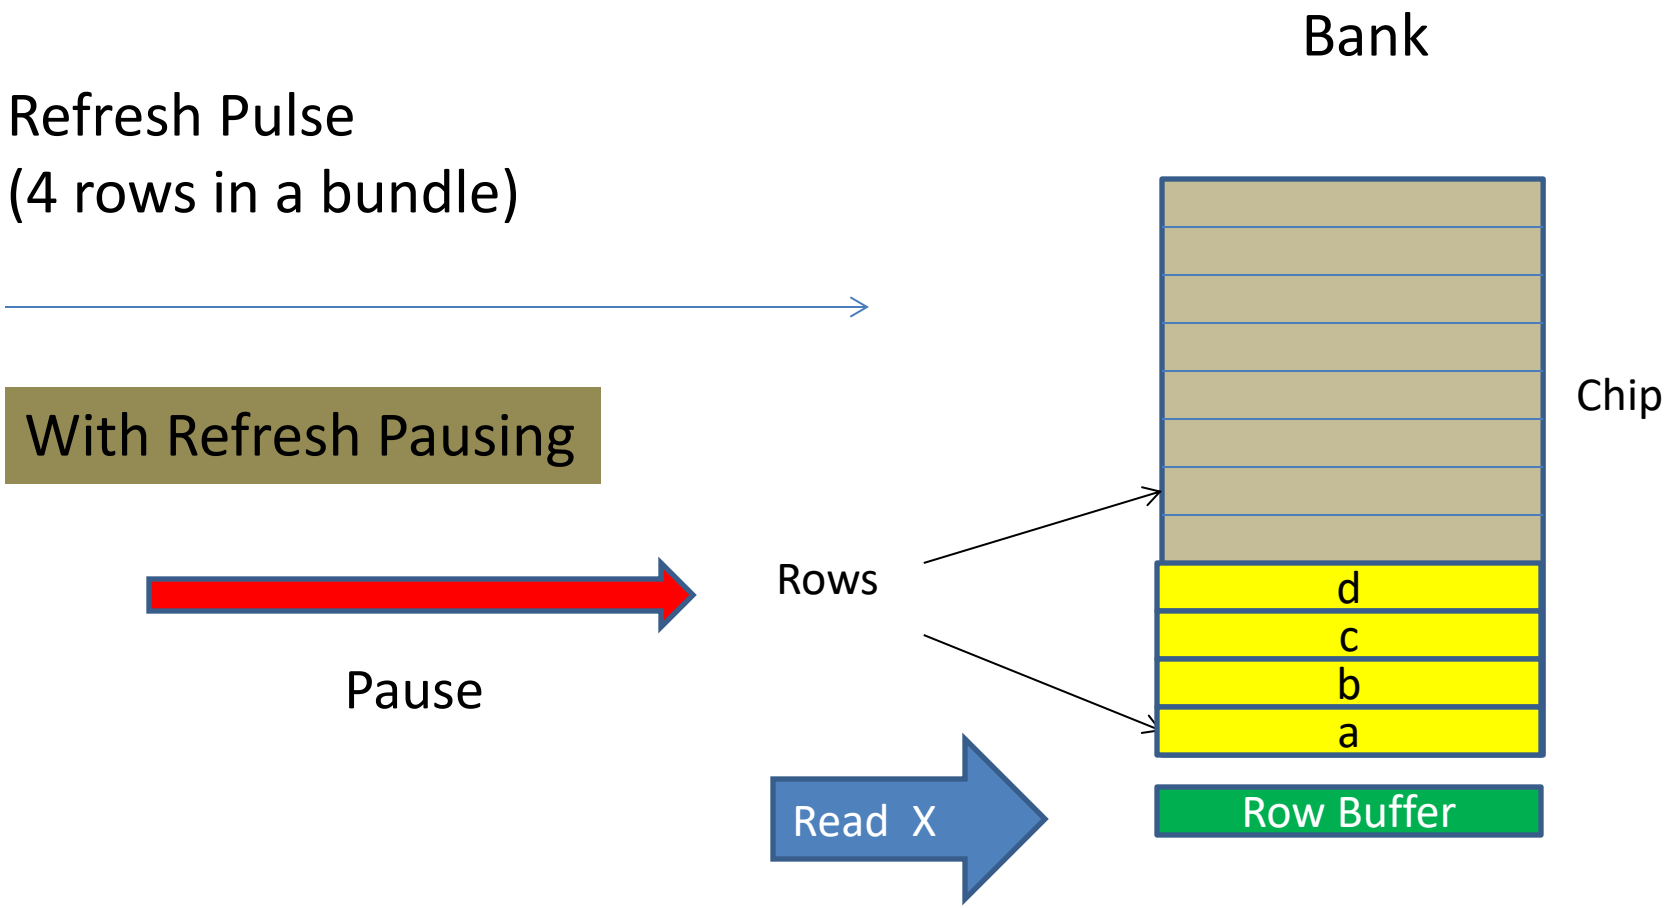

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

With Refresh Pausing

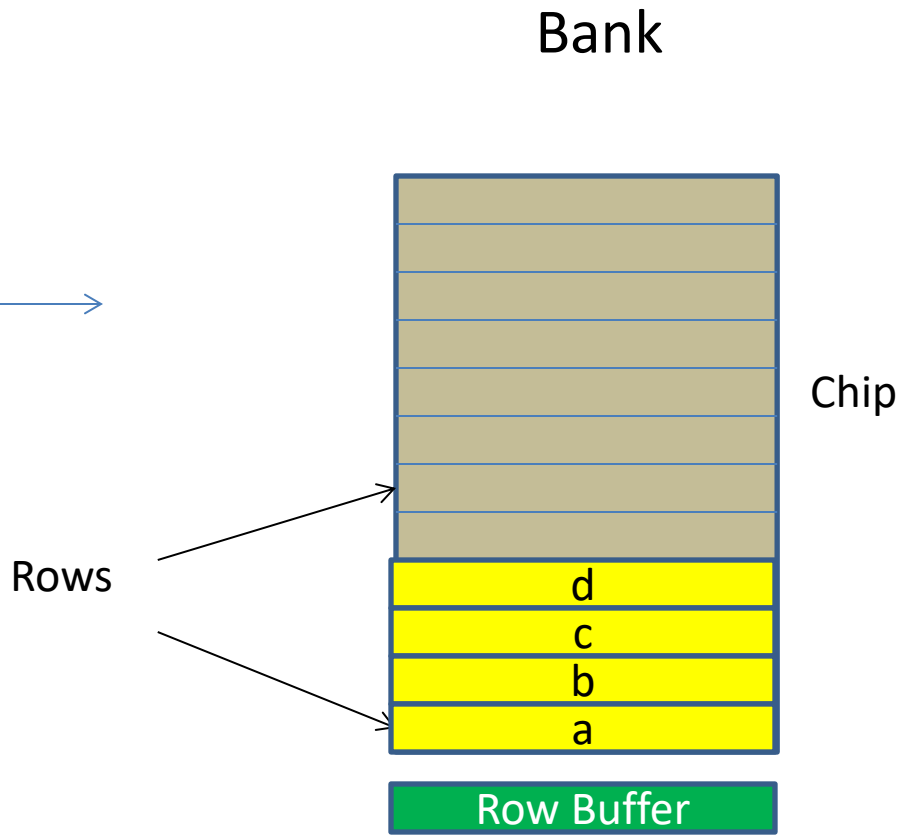

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

With Refresh Pausing

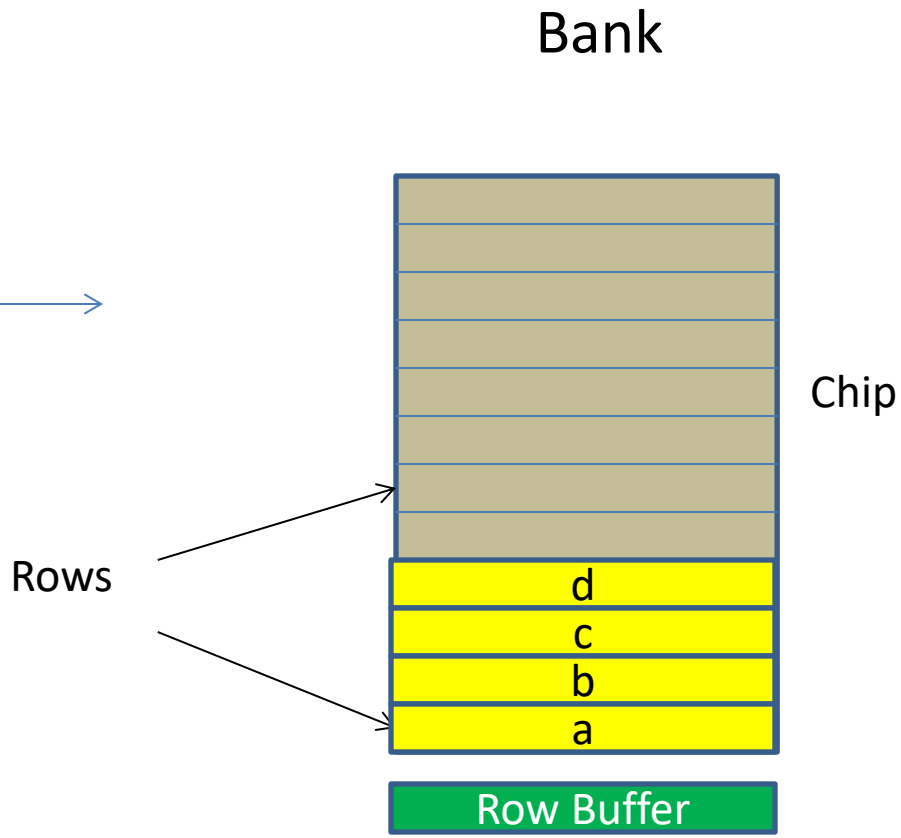

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

With Refresh Pausing

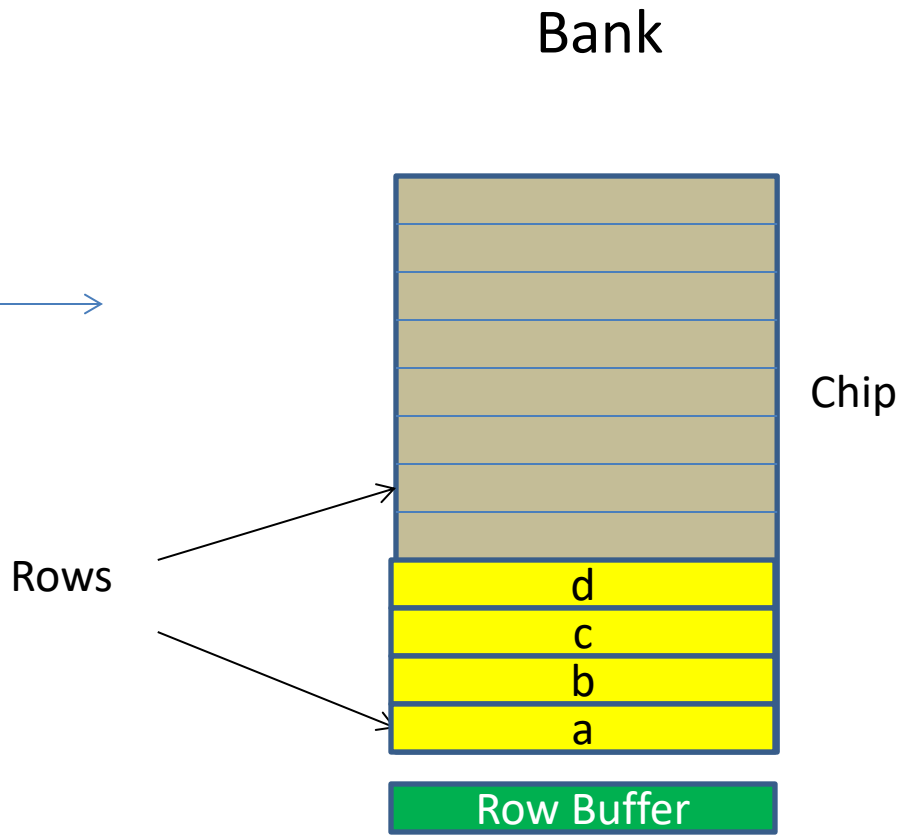

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

With Refresh Pausing

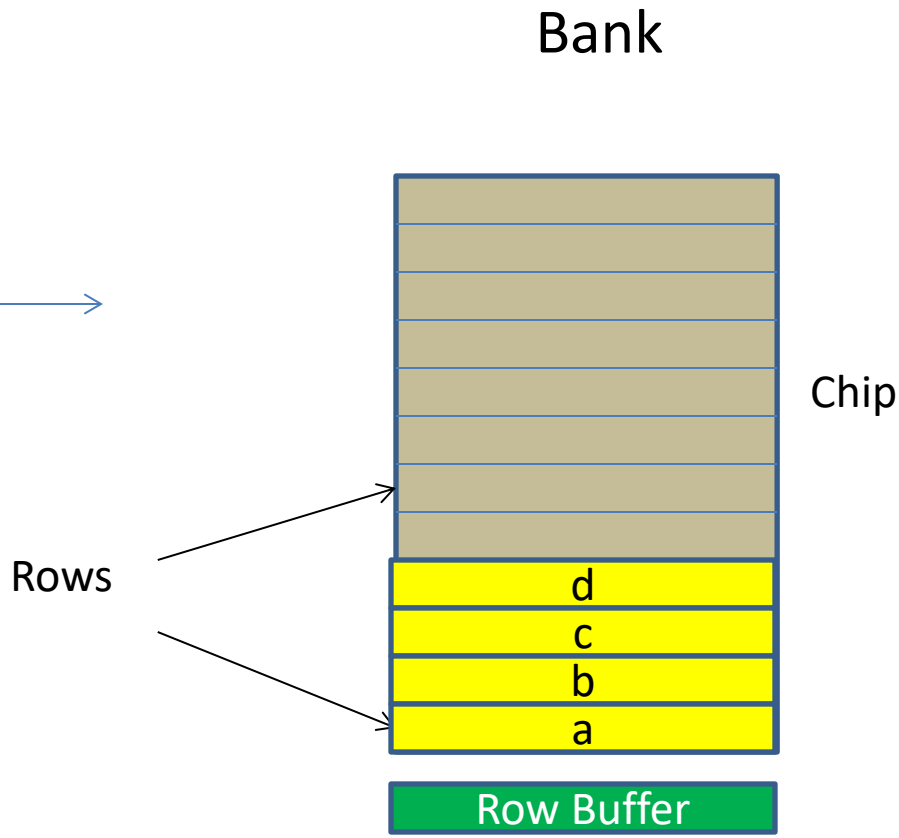

Refresh Pausing: When to Pause?

Refresh Pulse
(4 rows in a bundle)

With Refresh Pausing

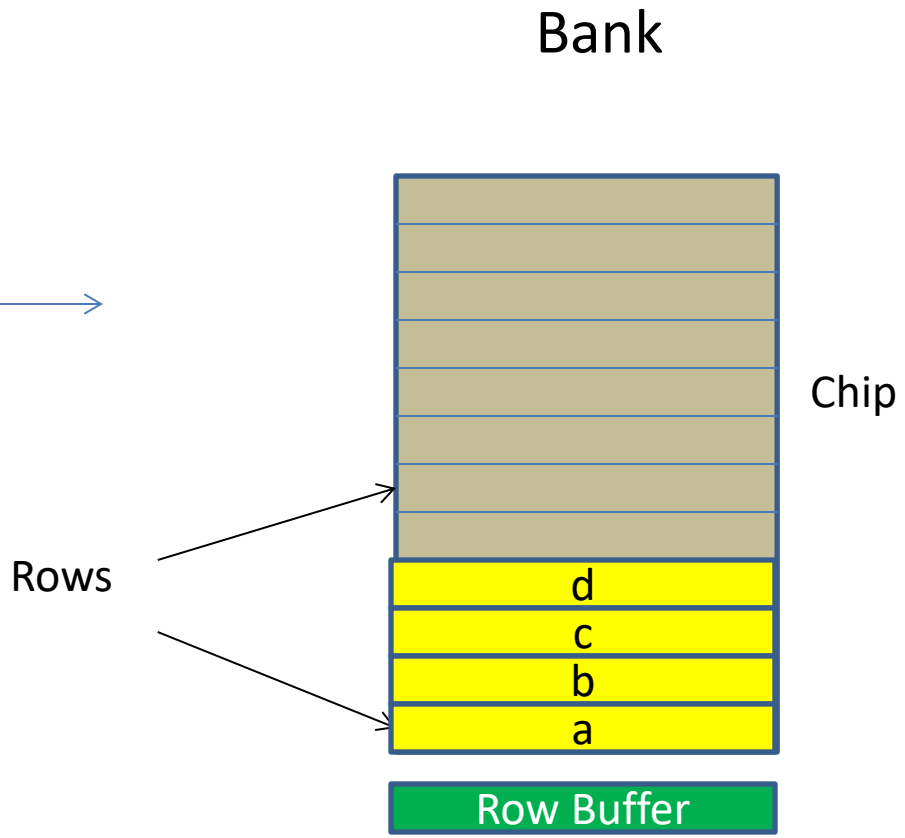

Refresh Pausing at Row boundary to service read

Refresh Pausing: Interface Details

- Memory Controller generates a Refresh Enable (RE) signal
- Pausing requires '*active low*' detection of RE
- One way communication only

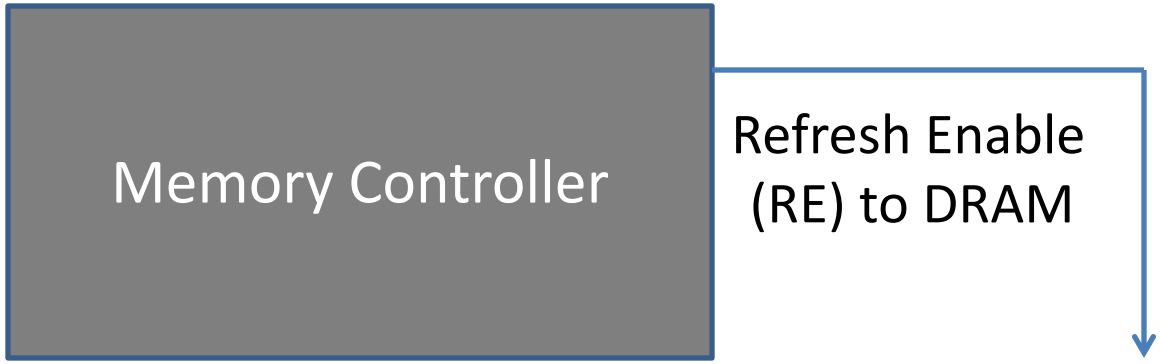

Refresh Pausing: Interface Details

- Memory Controller generates a Refresh Enable (RE) signal
- Pausing requires '*active low*' detection of RE
- One way communication only

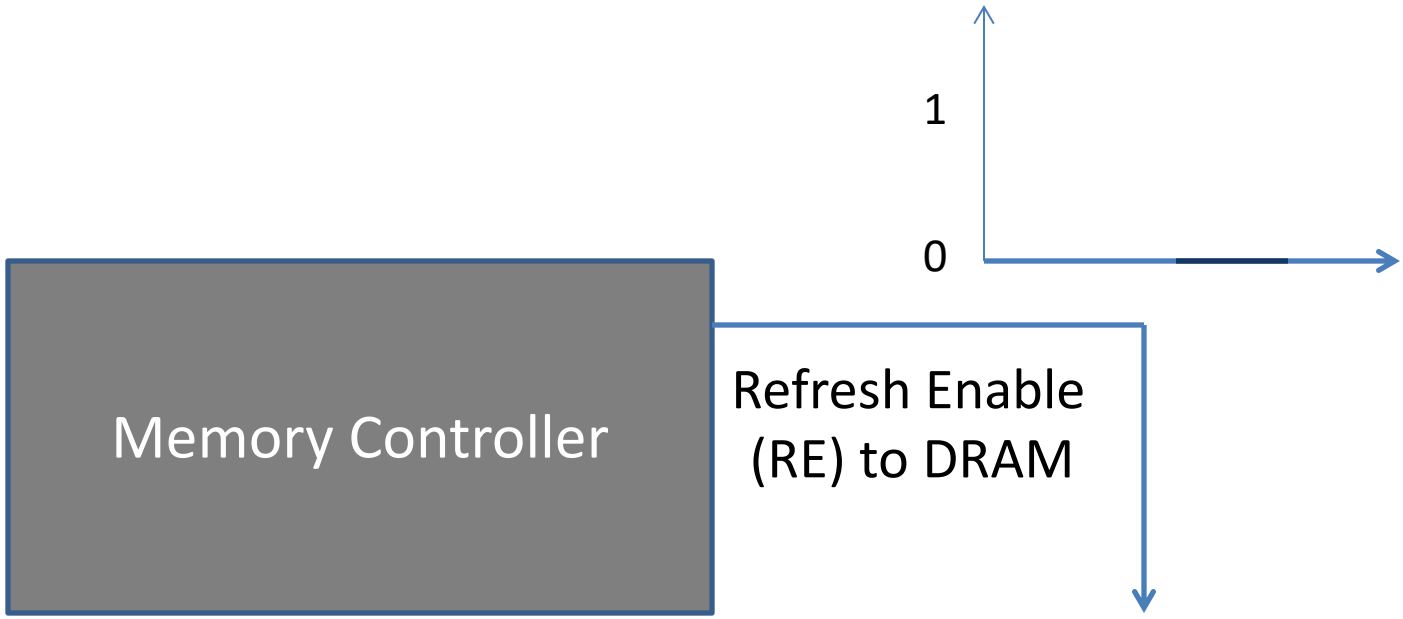

Refresh Pausing: Interface Details

- Memory Controller generates a Refresh Enable (RE) signal
- Pausing requires '*active low*' detection of RE
- One way communication only

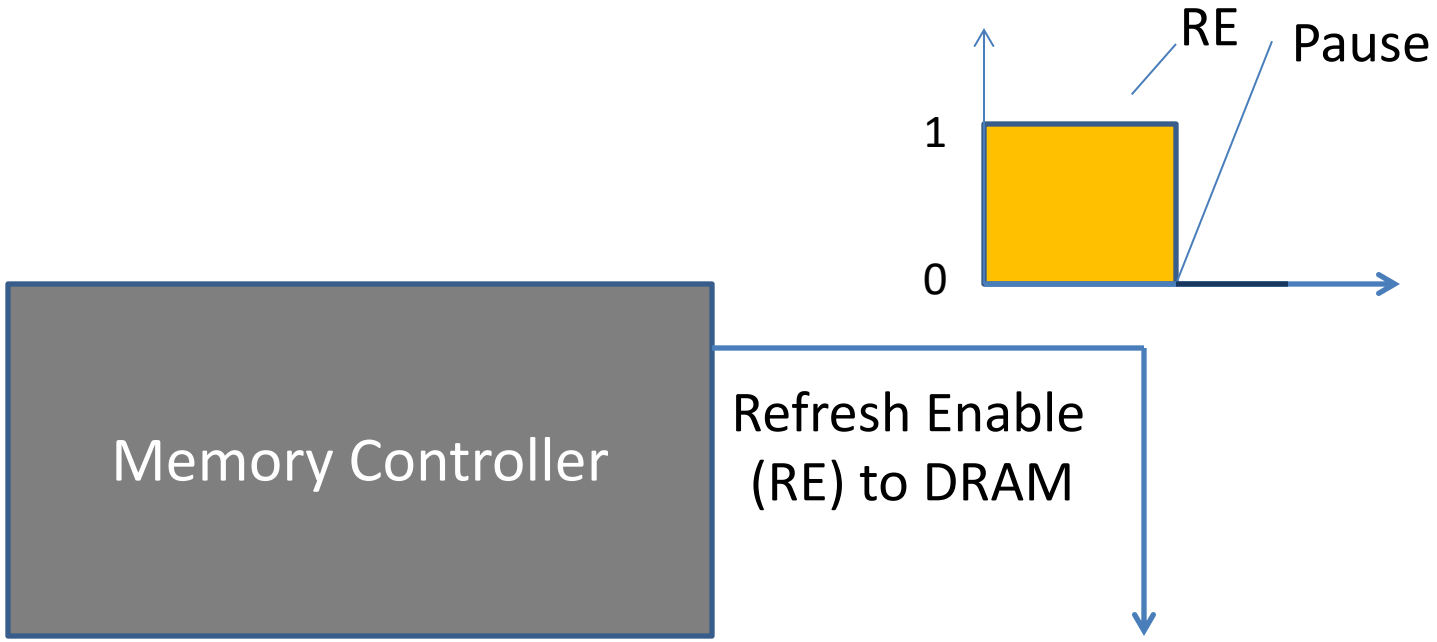

Refresh Pausing: Interface Details

- Memory Controller generates a Refresh Enable (RE) signal
- Pausing requires '*active low*' detection of RE
- One way communication only

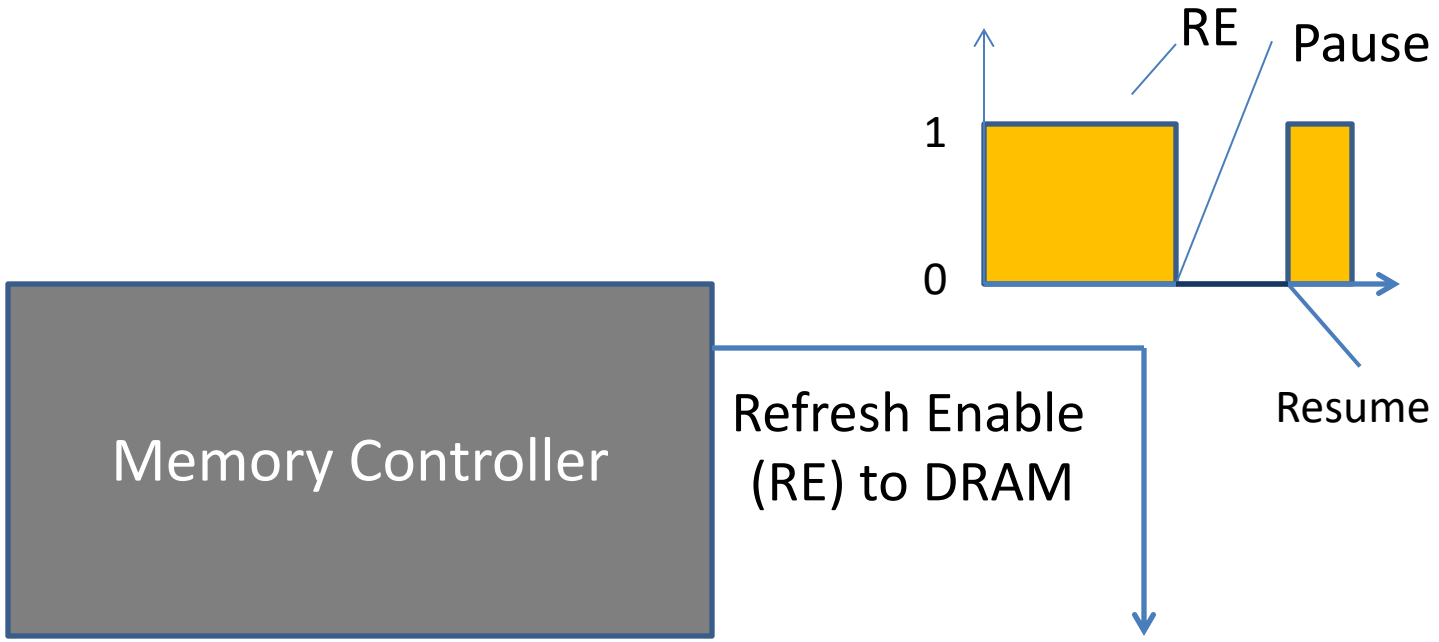

Refresh Pausing: Track a Paused Row

- Row Address Counter increments the addresses

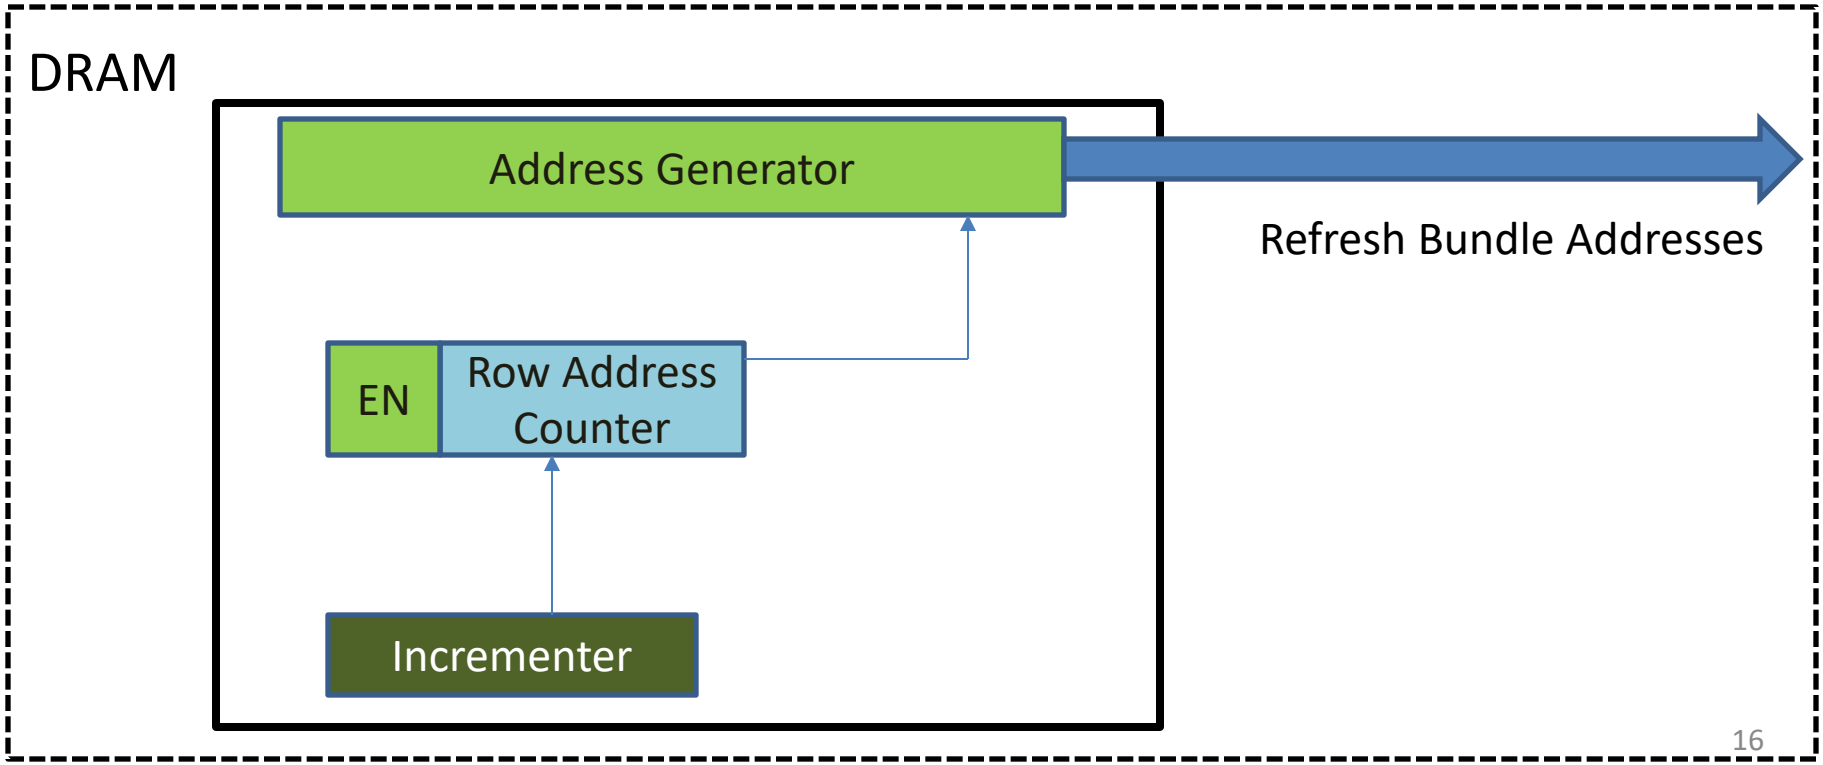

Refresh Pausing: Track a Paused Row

- Row Address Counter increments the addresses
- Stop the increment using a simple AND gate

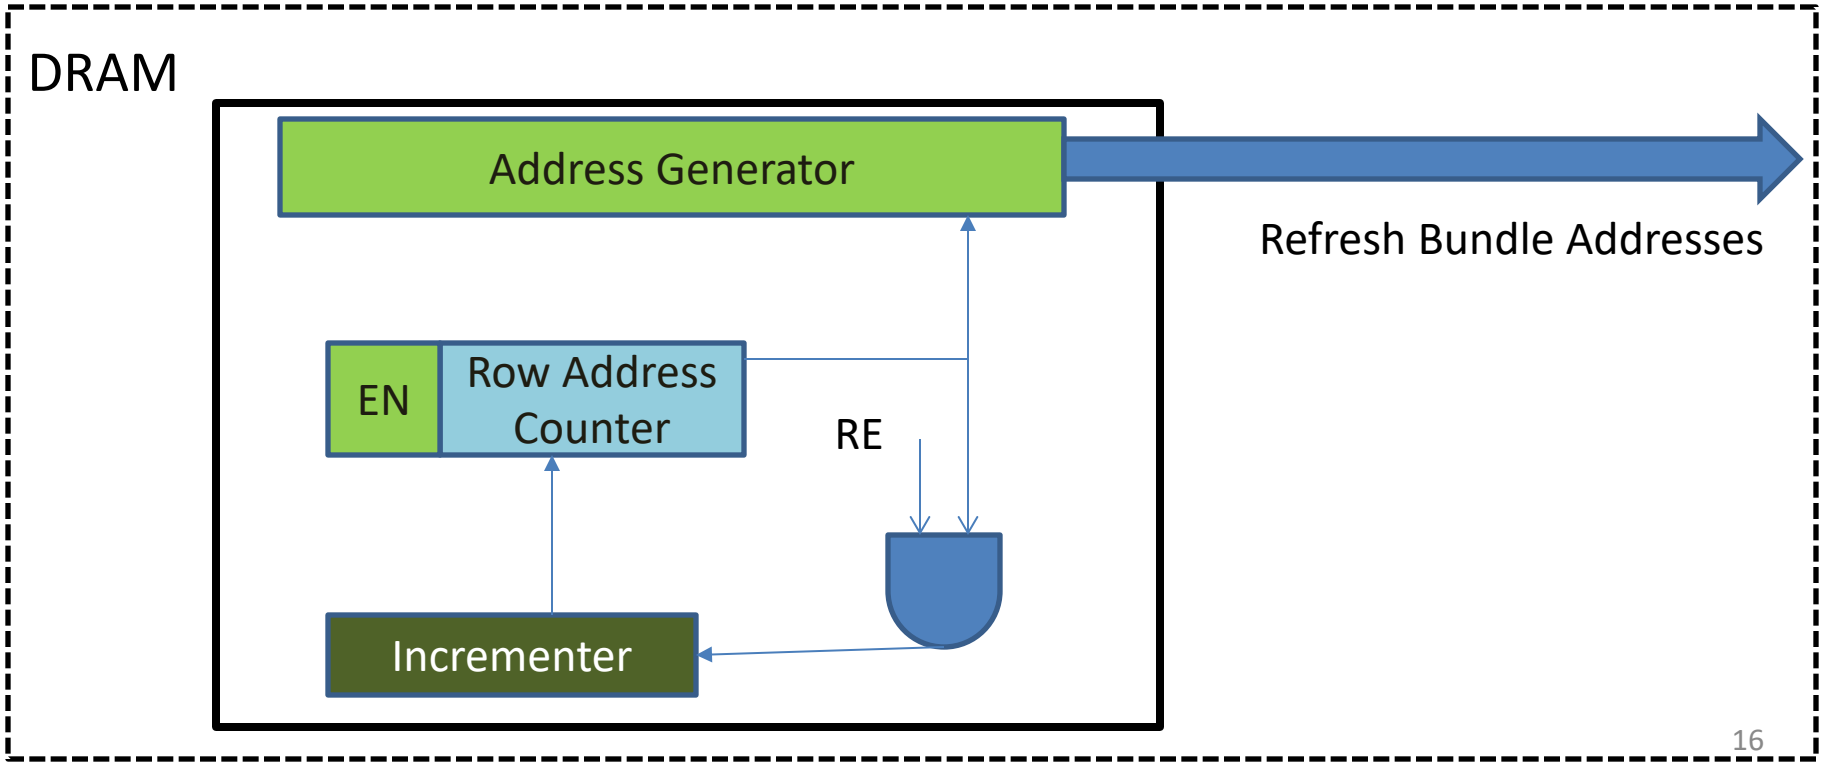

Refresh Pausing: Track a Paused Row

- Row Address Counter increments the addresses
- Stop the increment using a simple AND gate
- Active Low Refresh Enable as ‘Refresh Pause’

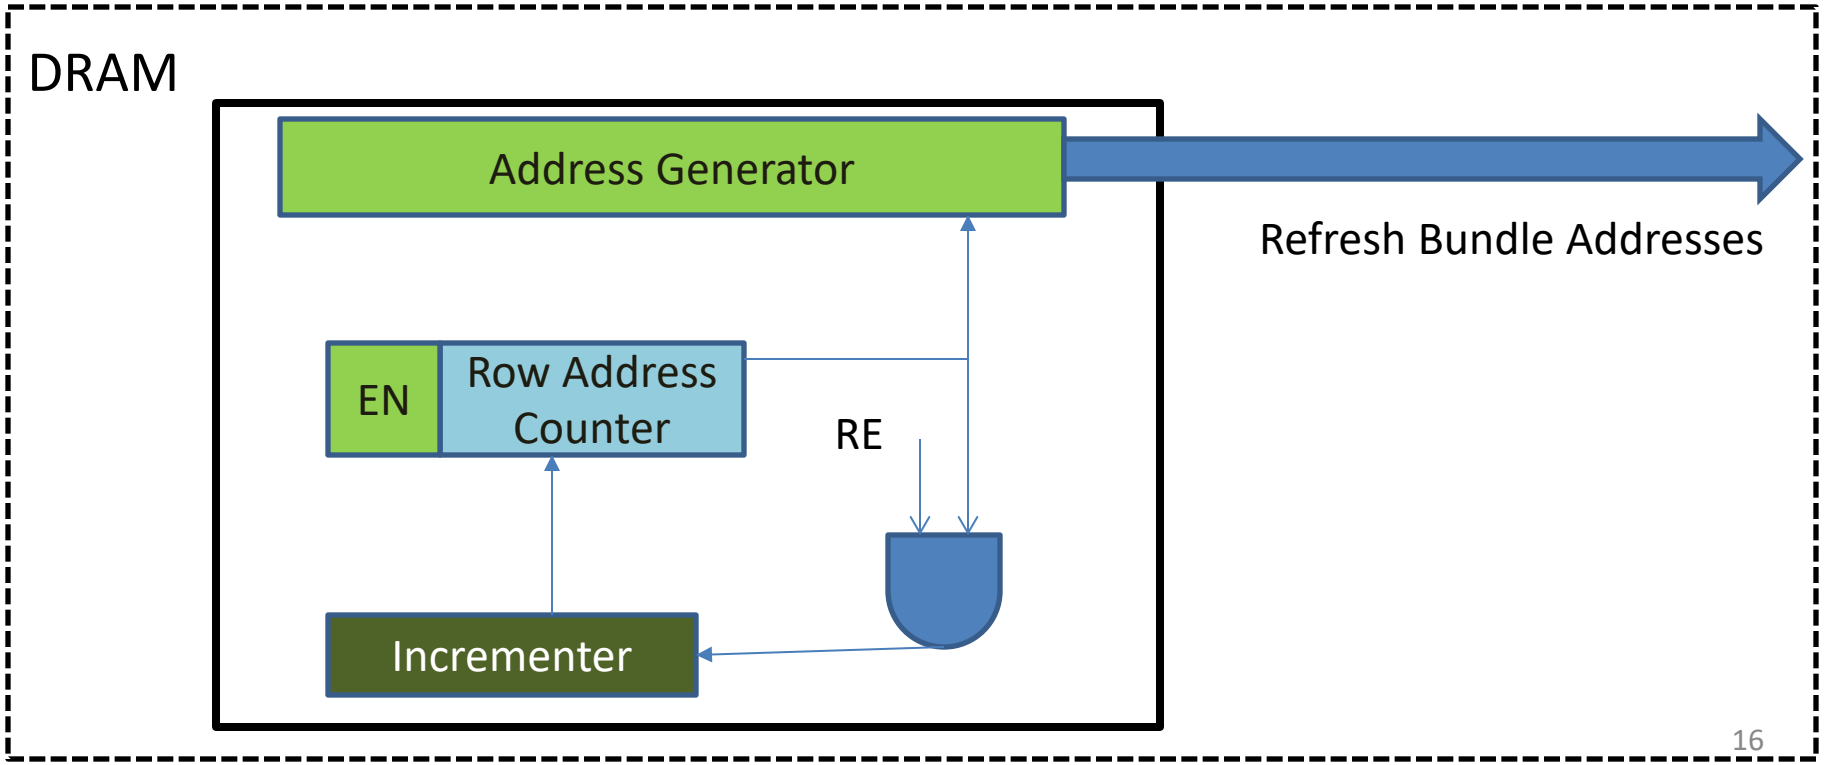

Refresh Pausing: Memory Scheduler

- Scheduler schedules: Read, Write, and Refresh
- Responsible for Pausing Refresh for Read
- Keeps track of refresh time done before Pause

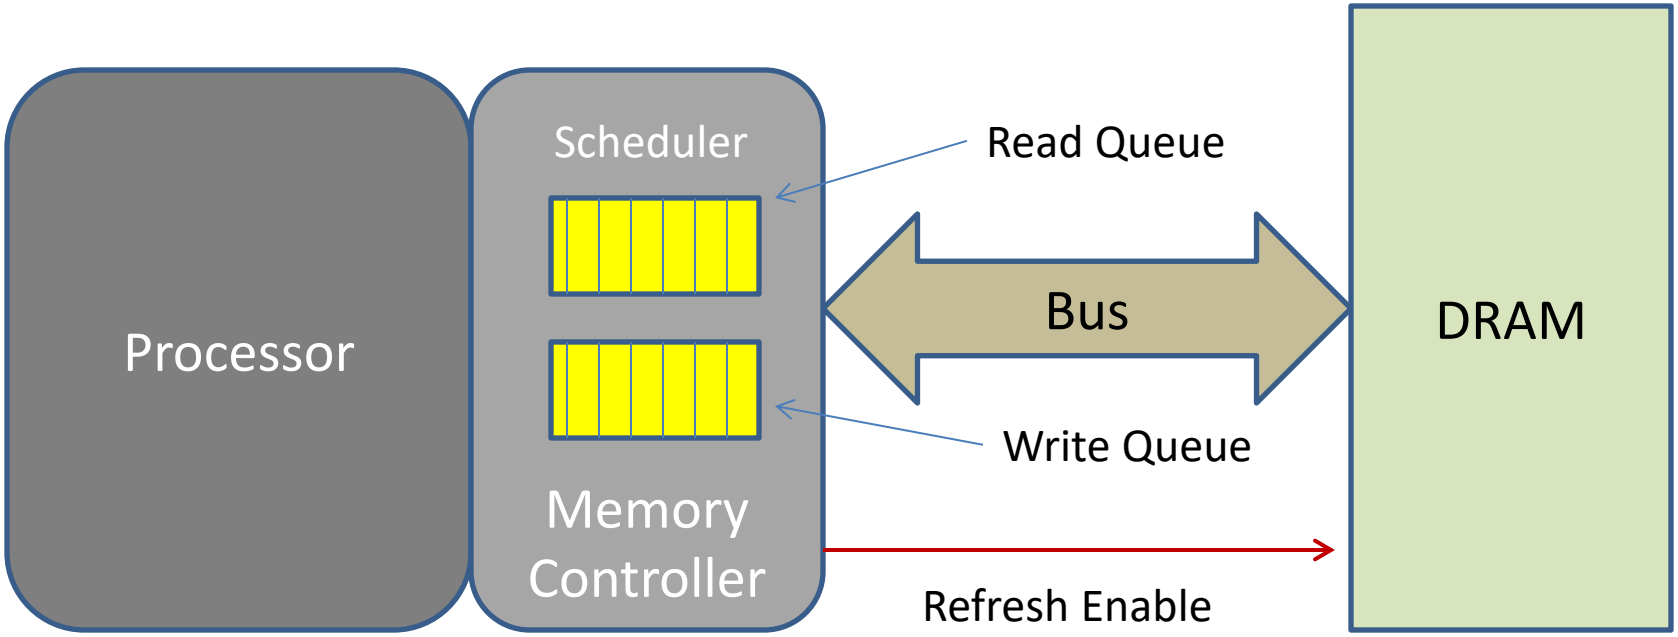

Forced Refresh

- Pausing can delay Refresh

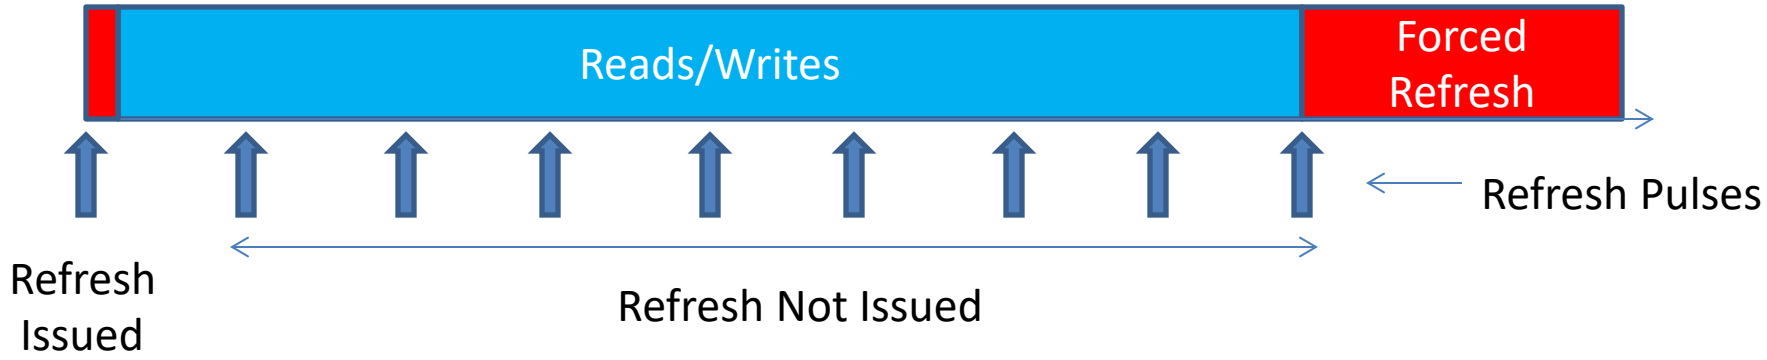

- JEDEC allows delay of up-to 8 pending refresh

Forced Refresh

- Pausing can delay Refresh

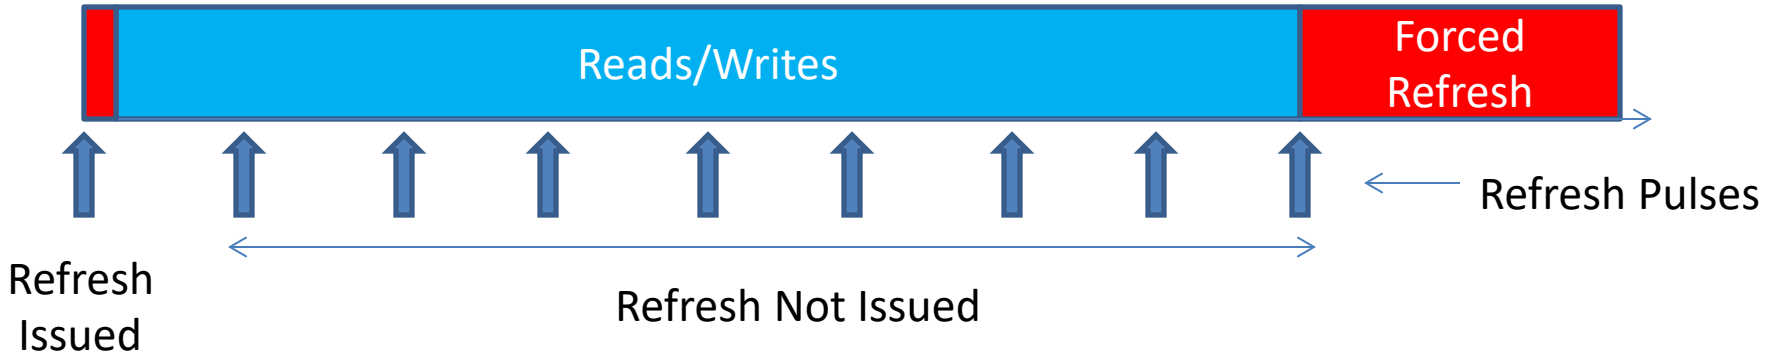

- JEDEC allows delay of up-to 8 pending refresh
- If 8 pending refresh, then issue ‘Forced Refresh’

Forced Refresh

- Pausing can delay Refresh

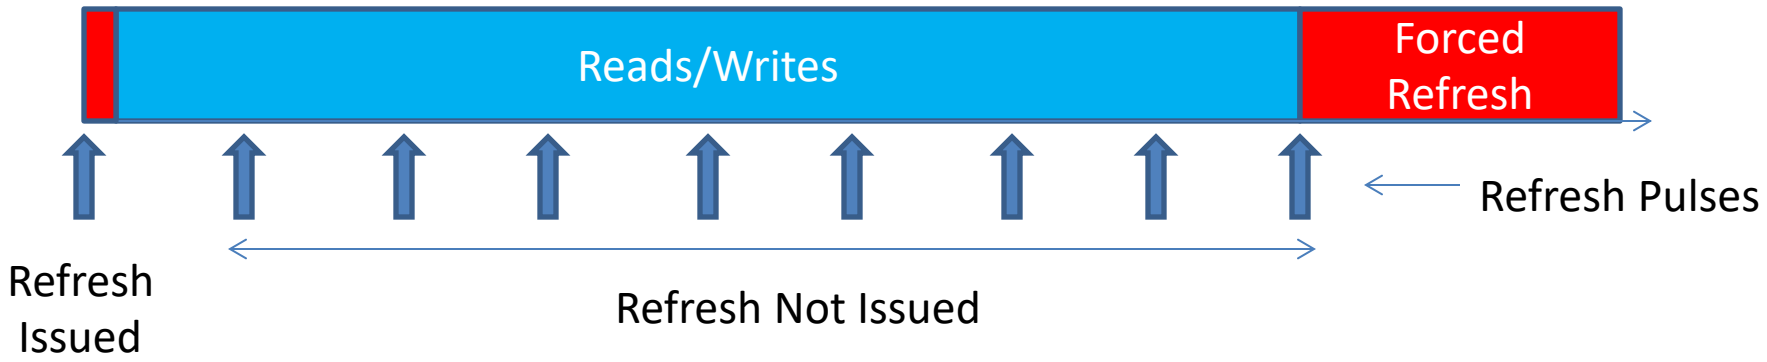

- JEDEC allows delay of up-to 8 pending refresh
- If 8 pending refresh, then issue ‘Forced Refresh’
- Forced Refresh cannot be Paused

Forced Refresh for data integrity

Outline

- Introduction & Motivation
- Refresh Operation: Background
- Refresh Pausing
- **Evaluation**
- Alternative Proposals
- Summary

Experimental Setup

- **Simulator:** uSIMM from Memory Scheduling Championship (MSC)

- **Workloads**: MSC Suite
  COMMERCIAL(5), PARSEC(9), BIOBENCH(2) and SPEC(2)

- **Configuration:**

Number of Cores	4
Last Level Cache	1MB
DRAM (DDR3)	8 Chips/Rank, 8Gb/Chip
Channels, Ranks, Banks	4,2,8
Refresh (Baseline)	Distributed (JEDEC)

- Results presented for temperature > 85C (paper also has <85C)

Results: Read Latency

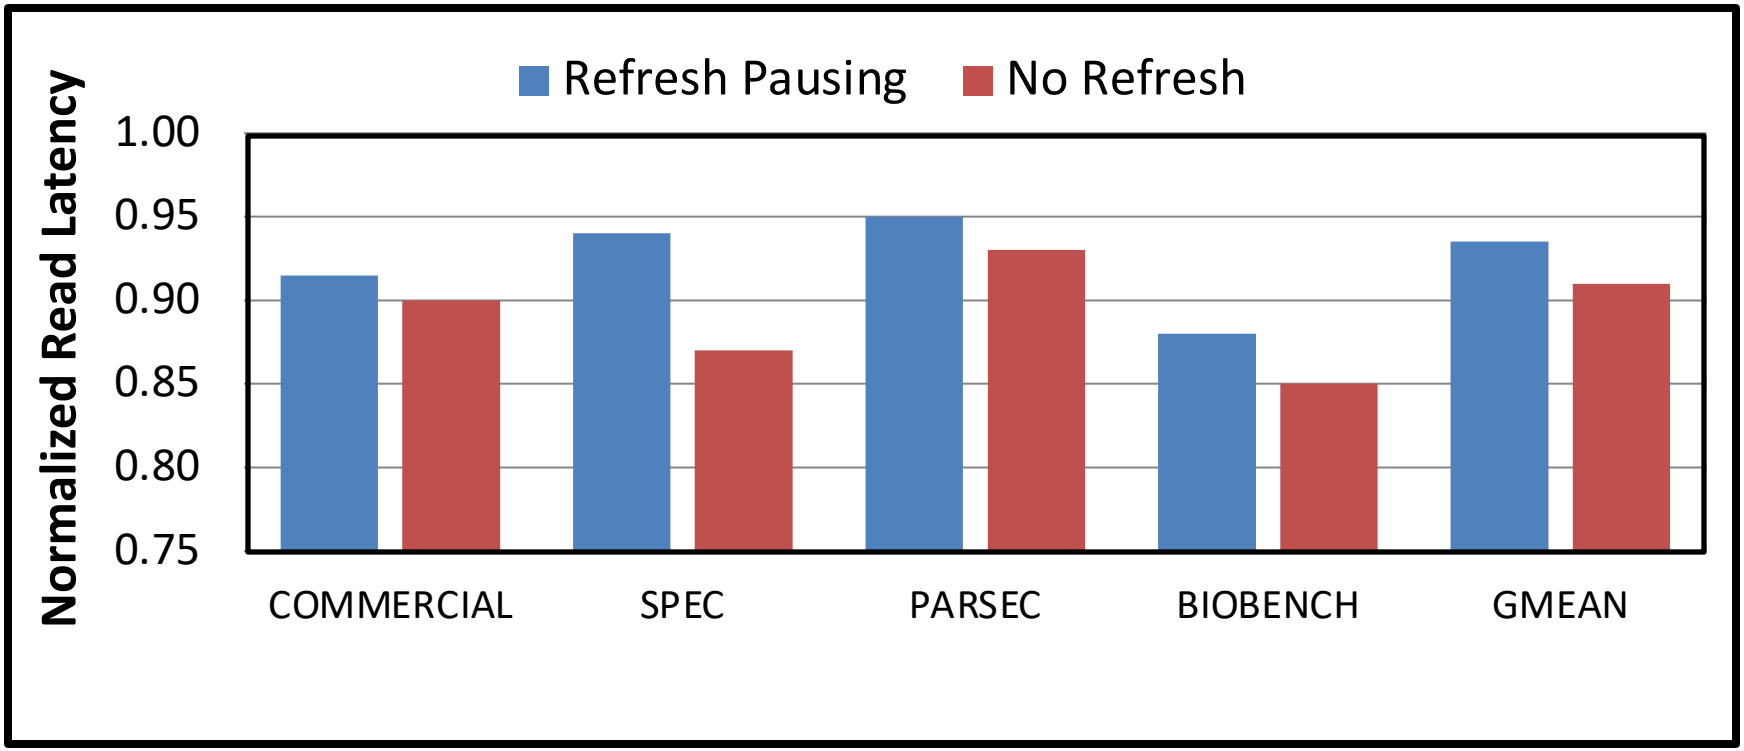

Results: Read Latency

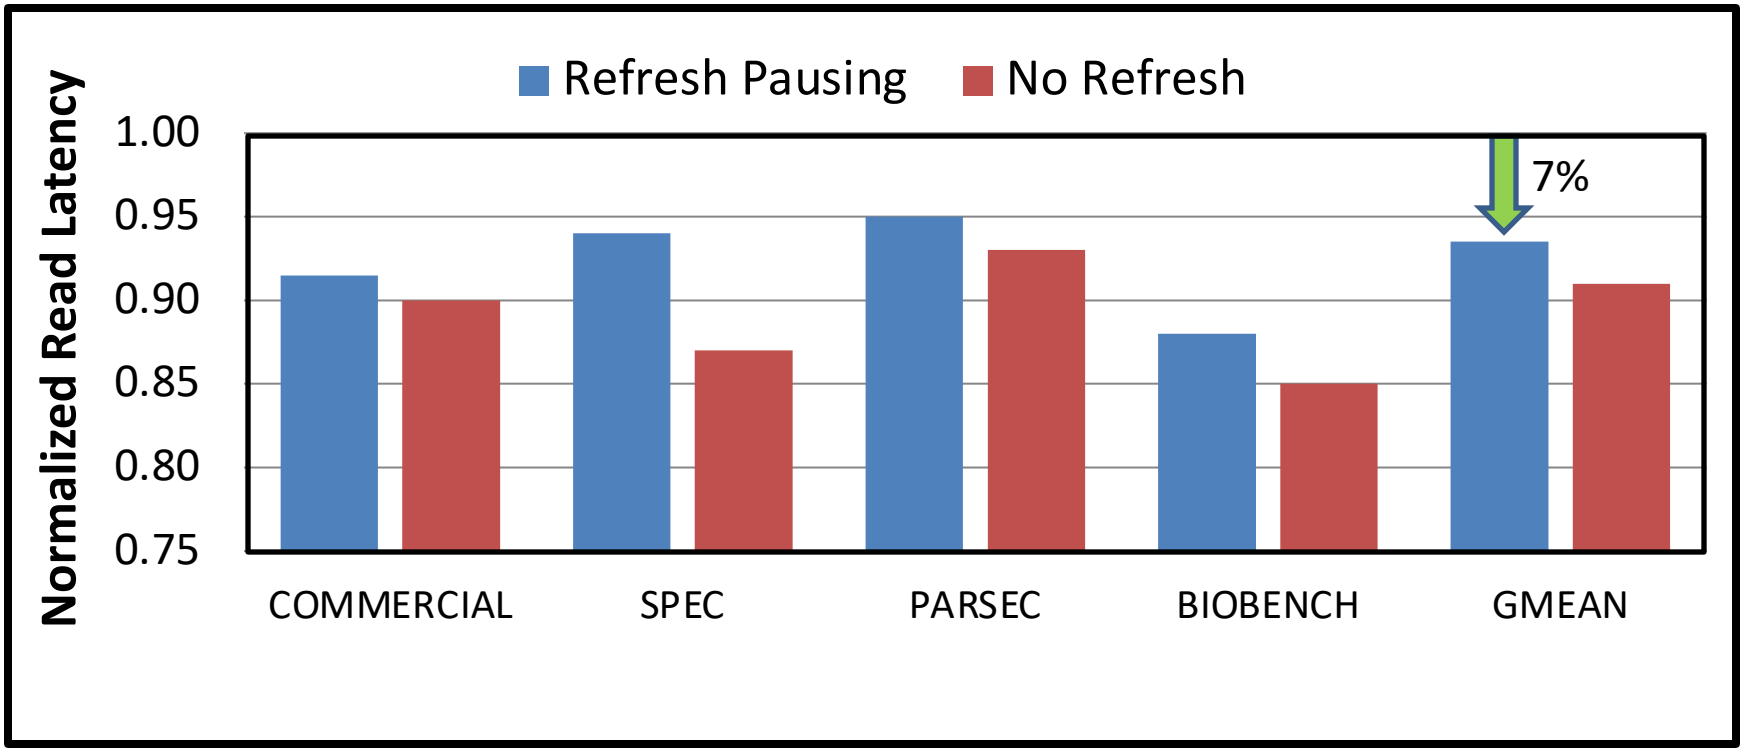

- Refresh Pausing gives ~7% read latency reduction for an 8Gb chip

Results: Performance

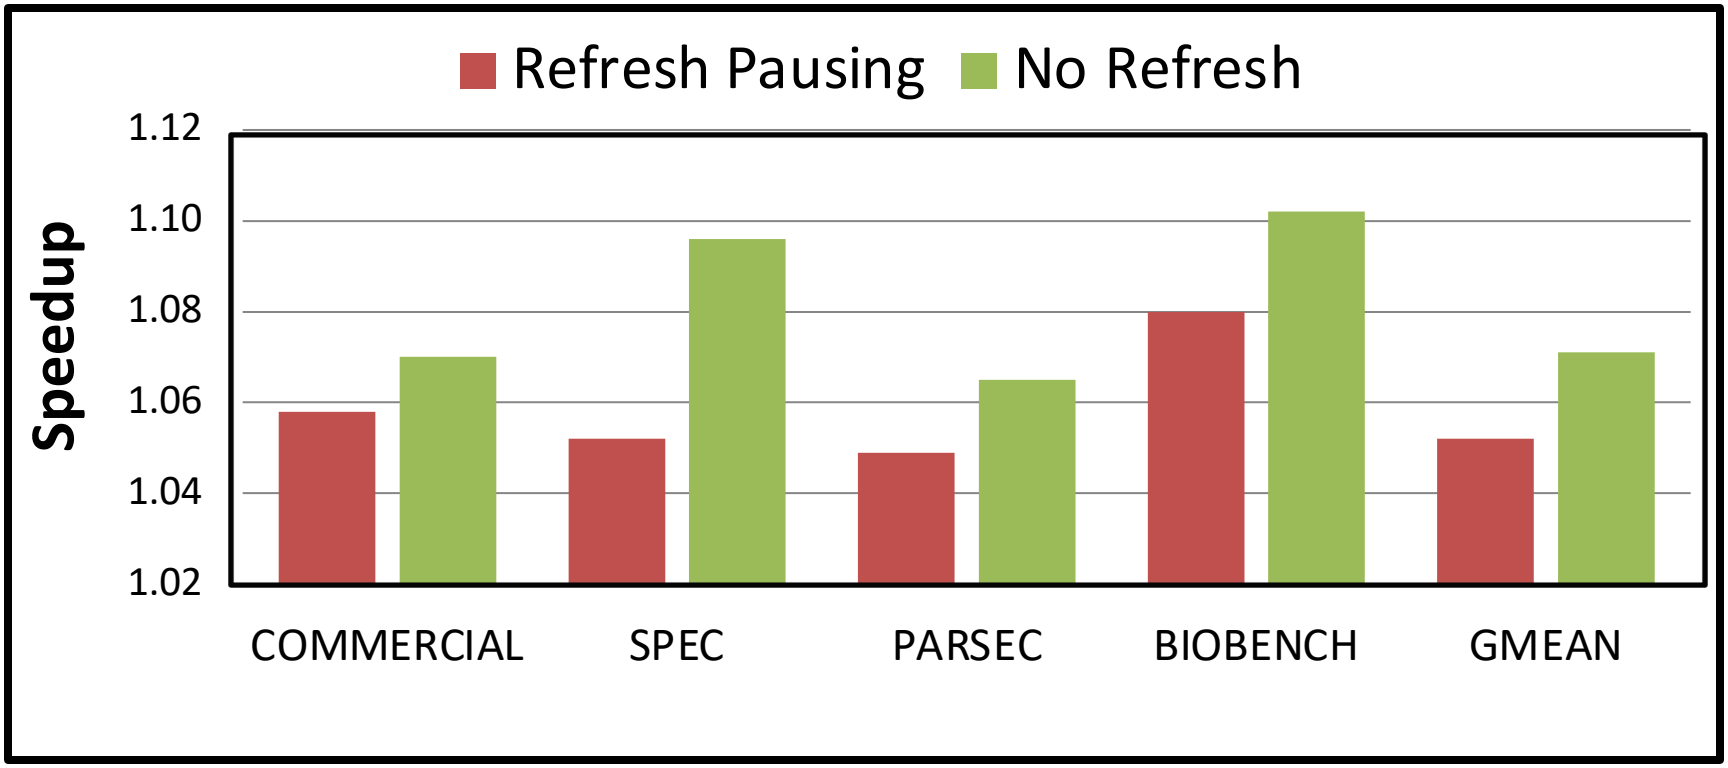

Results: Performance

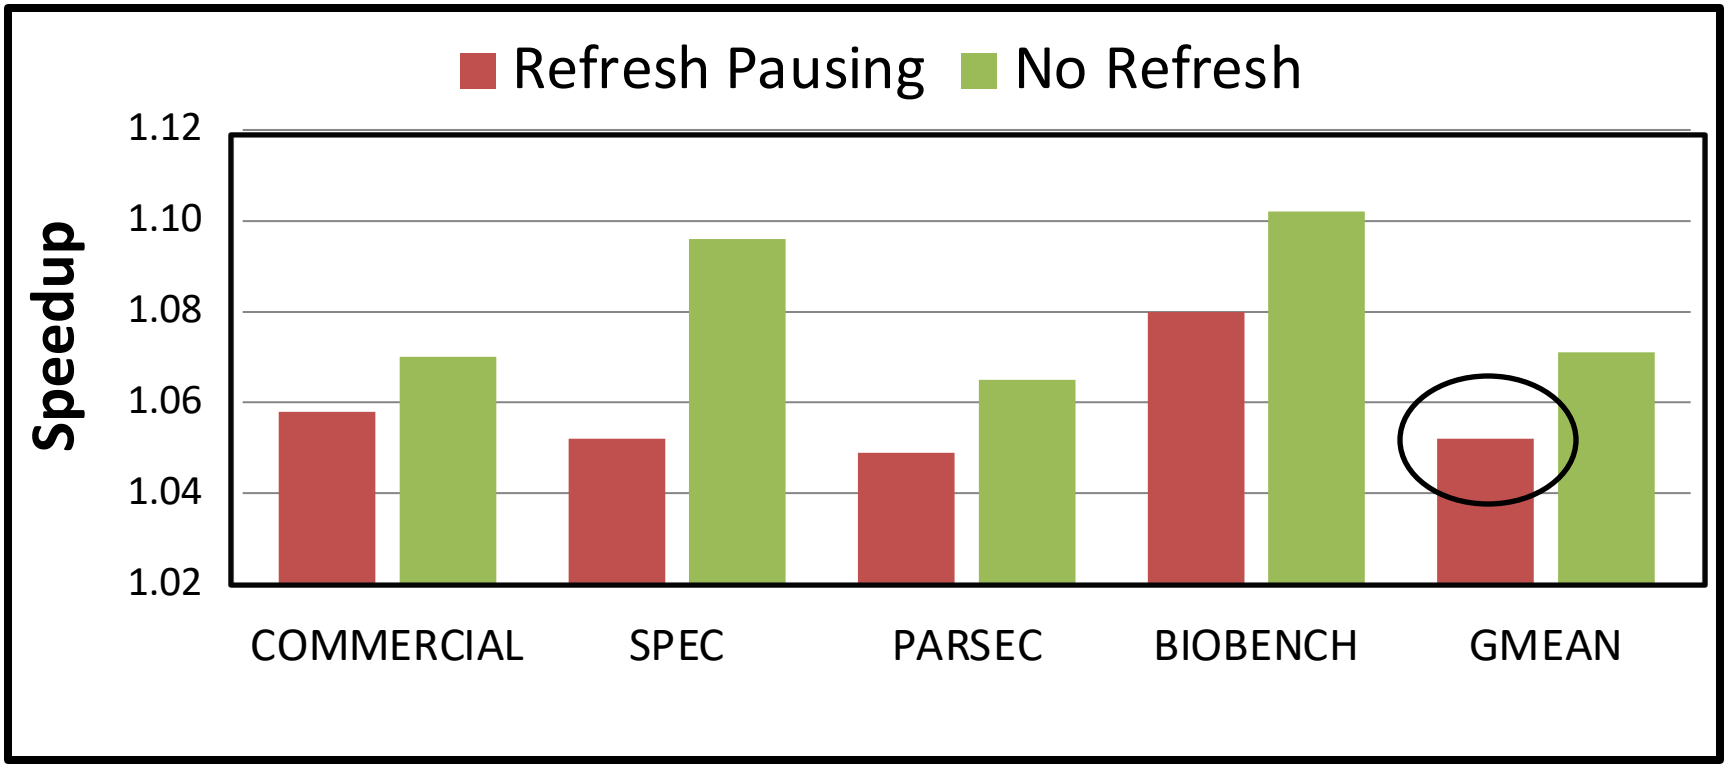

- Refresh Pausing gives ~5% performance improvement for an 8Gb chip

Results: Impact of Chip Density

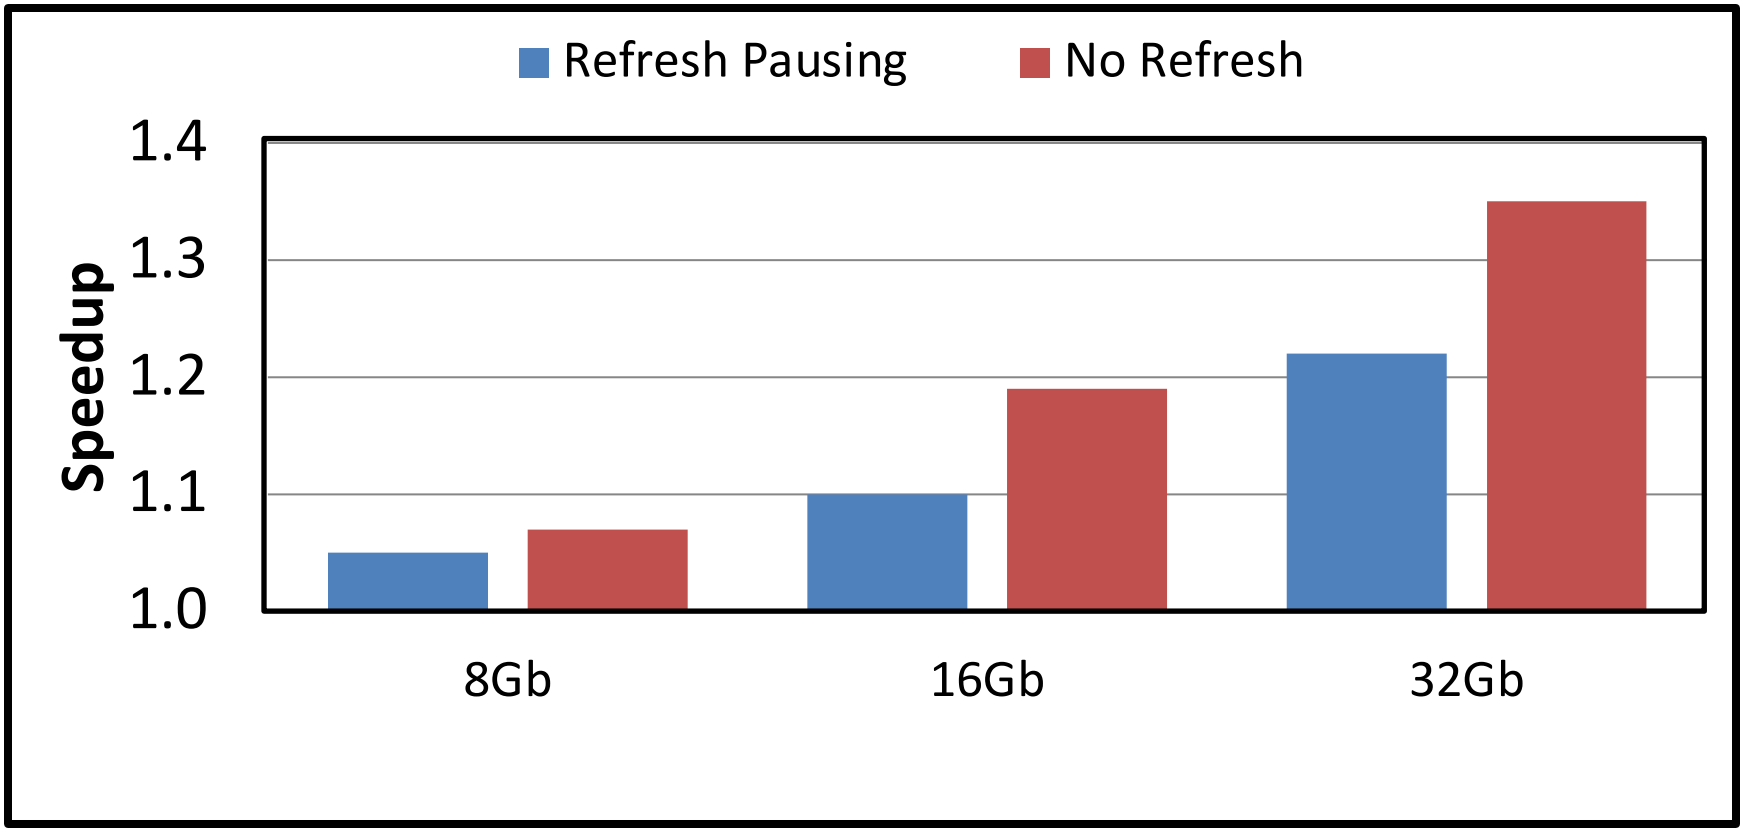

Results: Impact of Chip Density

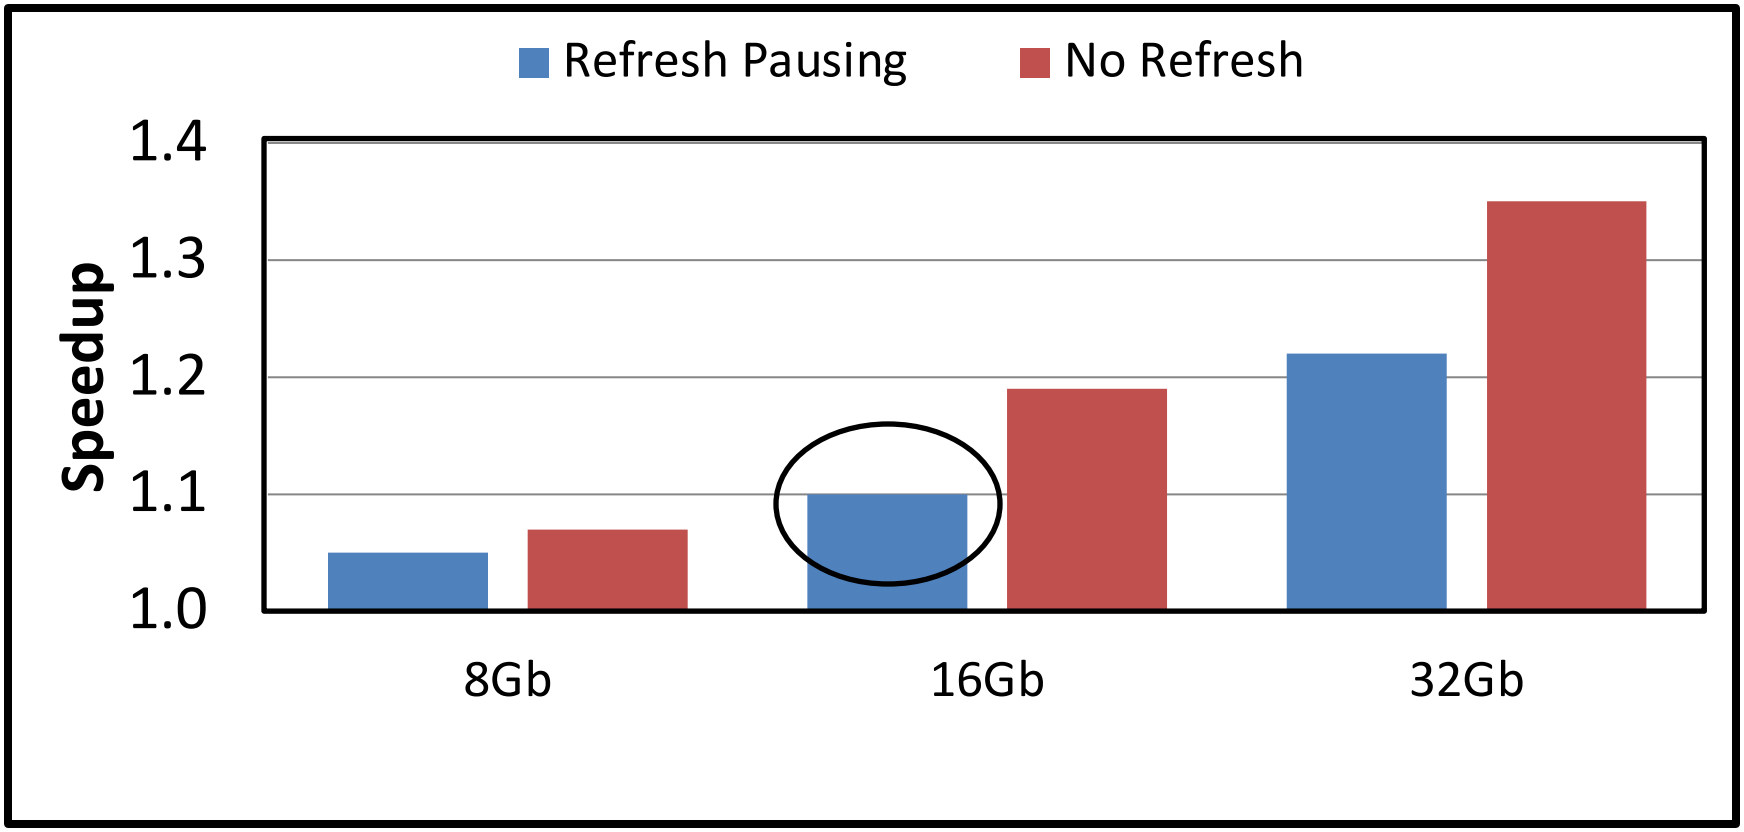

Results: Impact of Chip Density

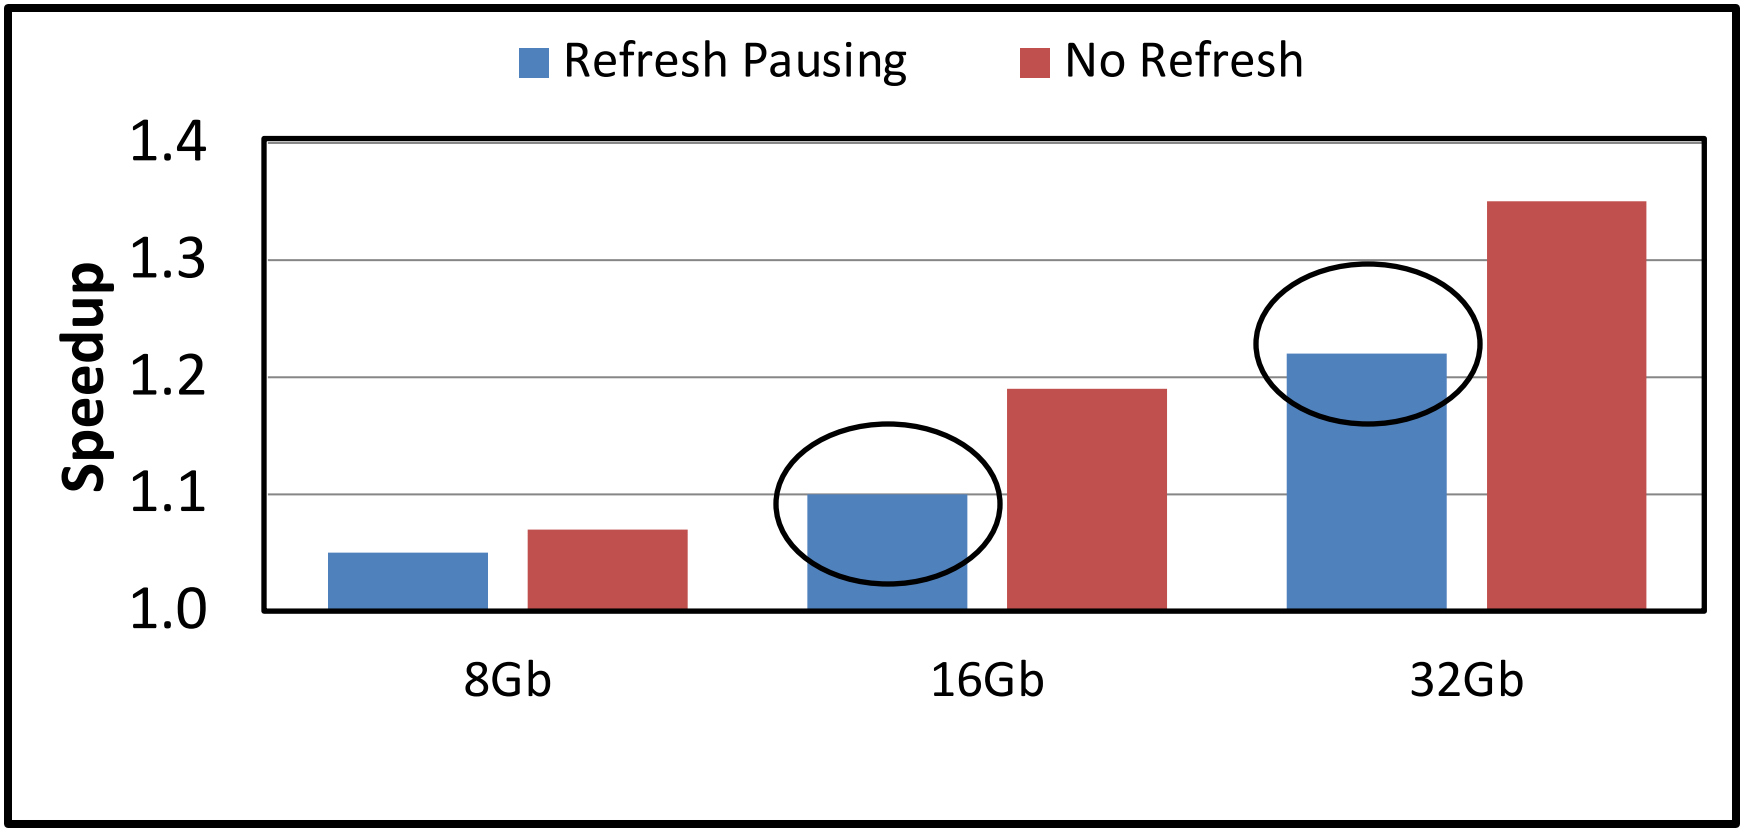

Results: Impact of Chip Density

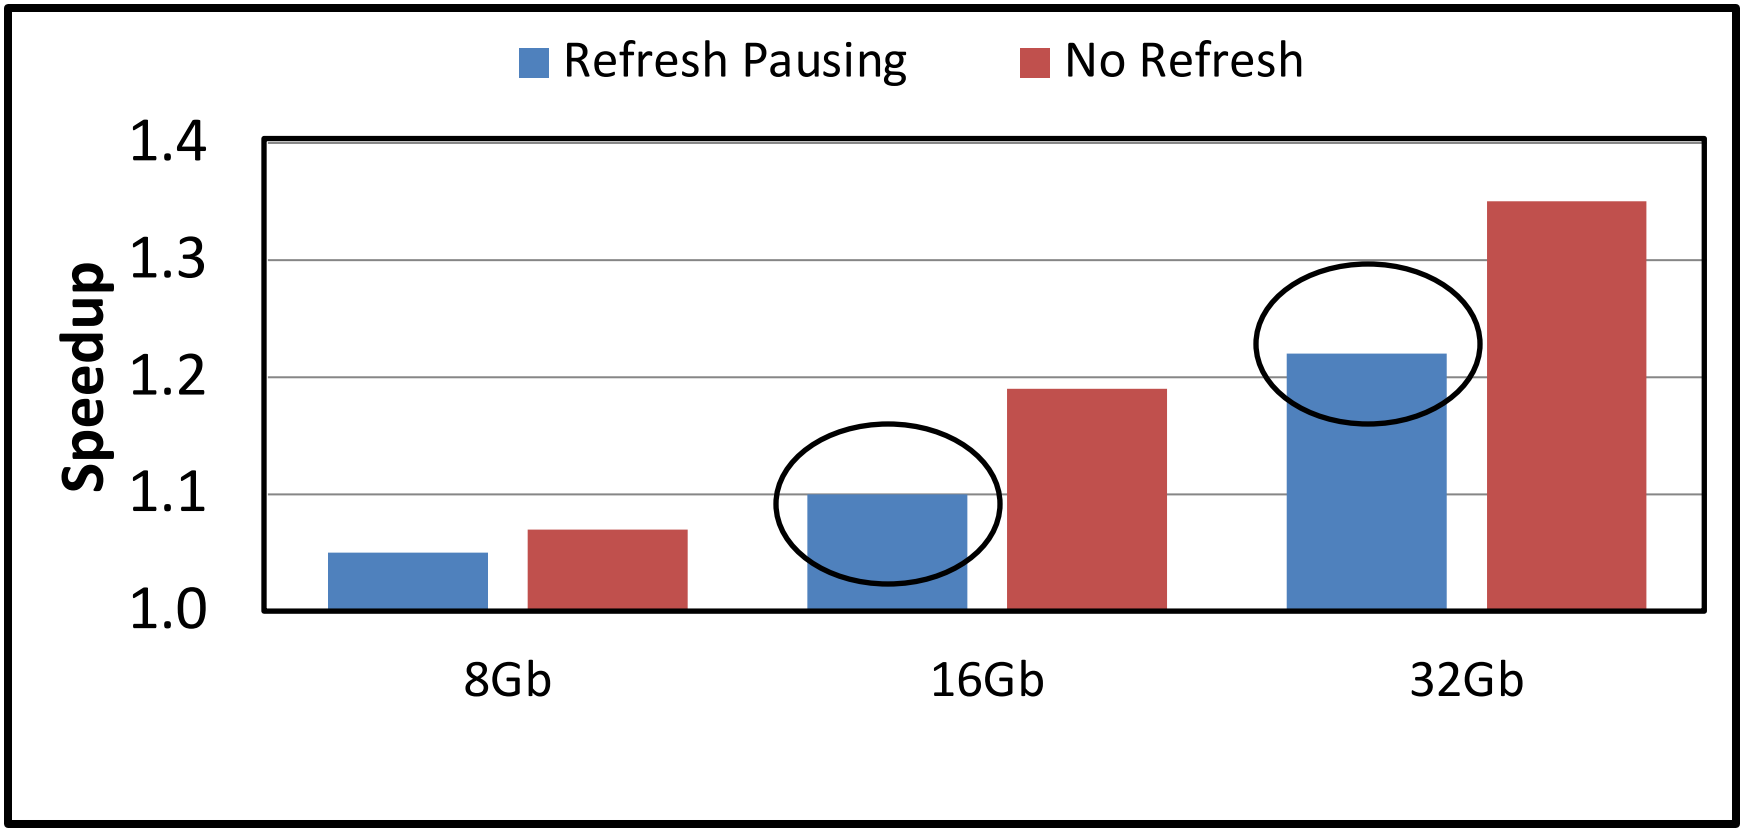

Refresh Pausing more effective as chips density increases

Outline

- Introduction & Motivation
- Refresh Operation: Background
- Refresh Pausing
- Evaluation
- **Alternative Proposals**
- Summary

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

No Refreshes

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

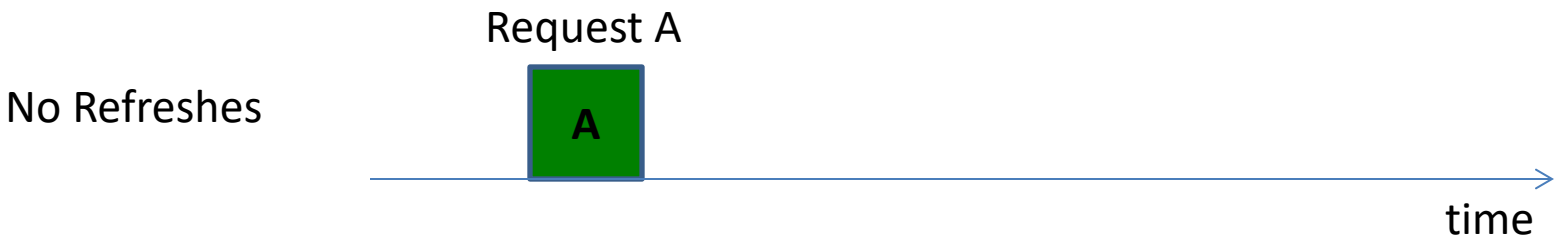

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

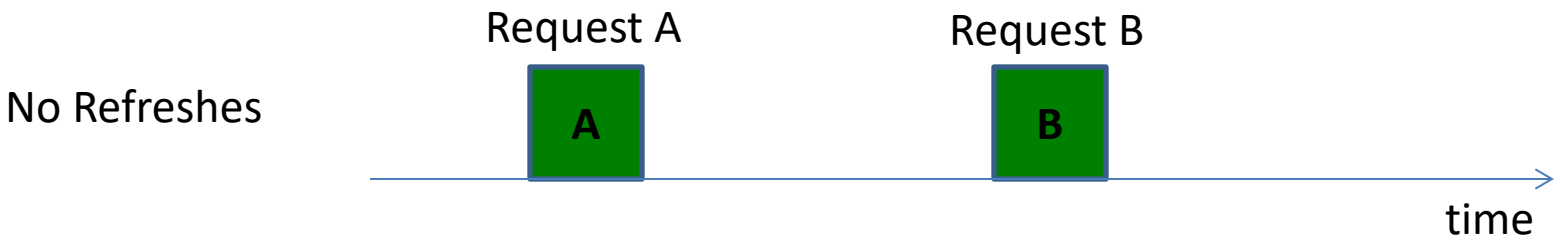

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

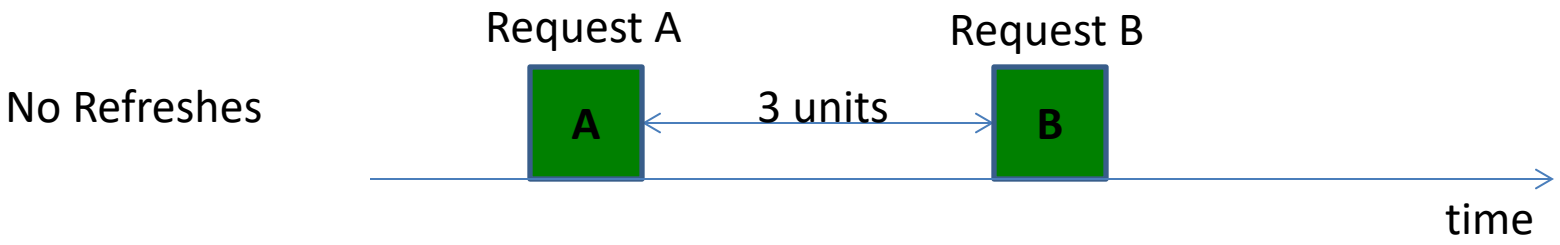

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

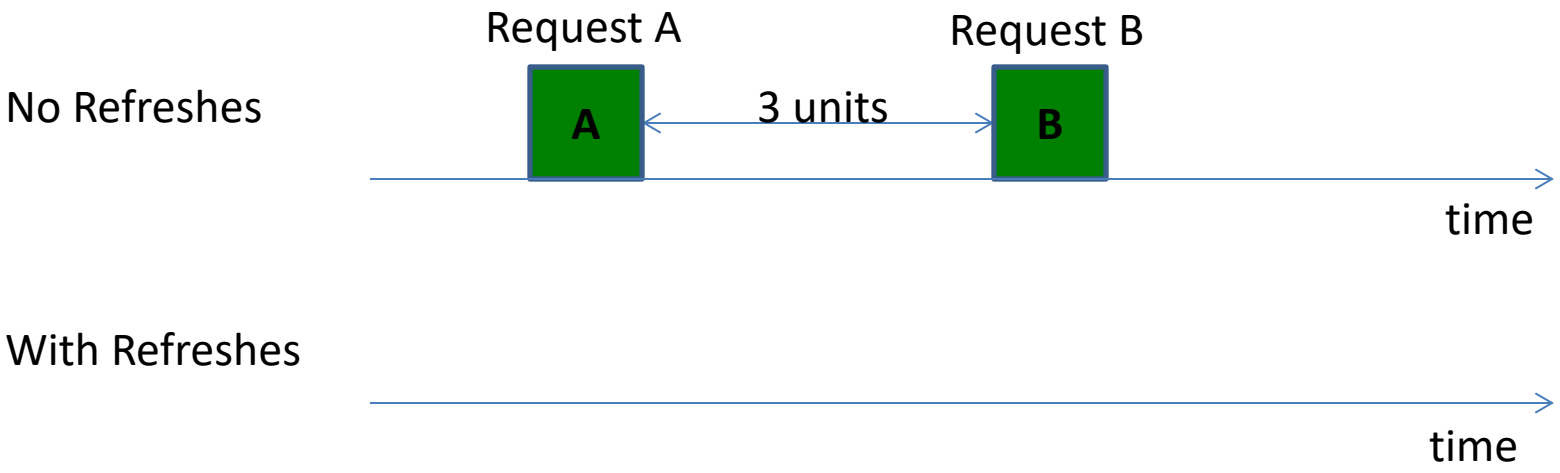

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

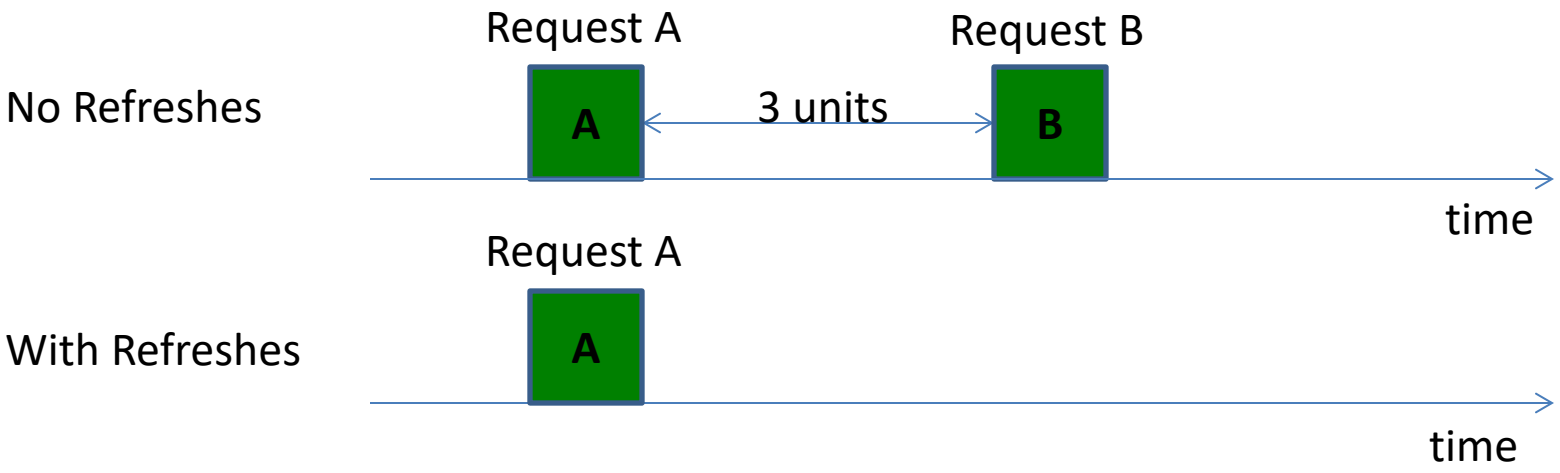

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

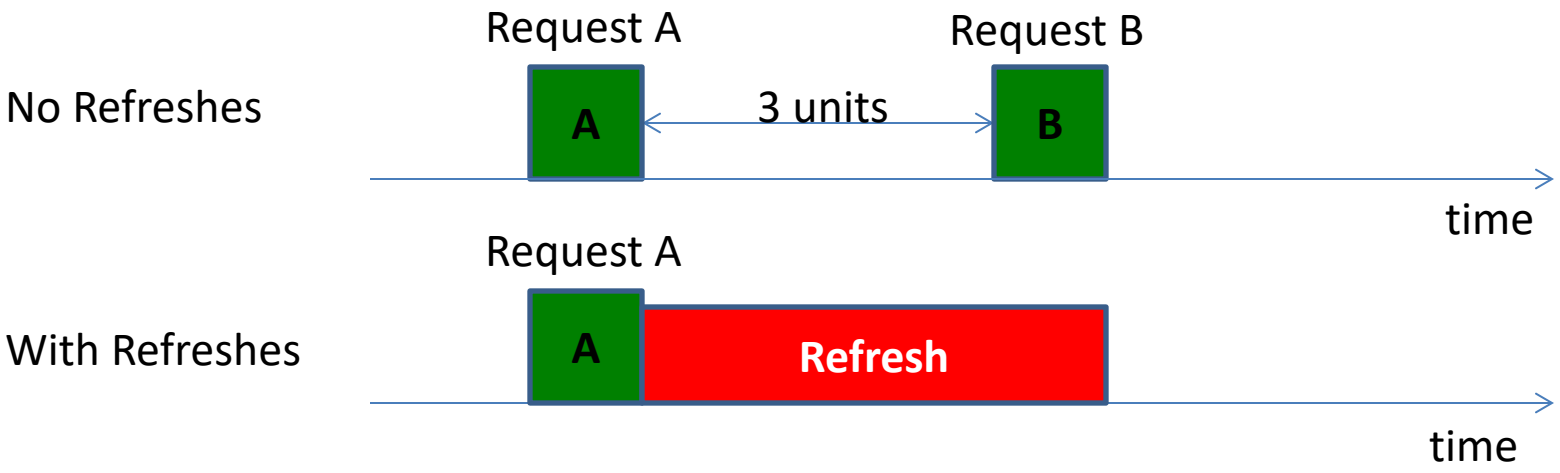

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

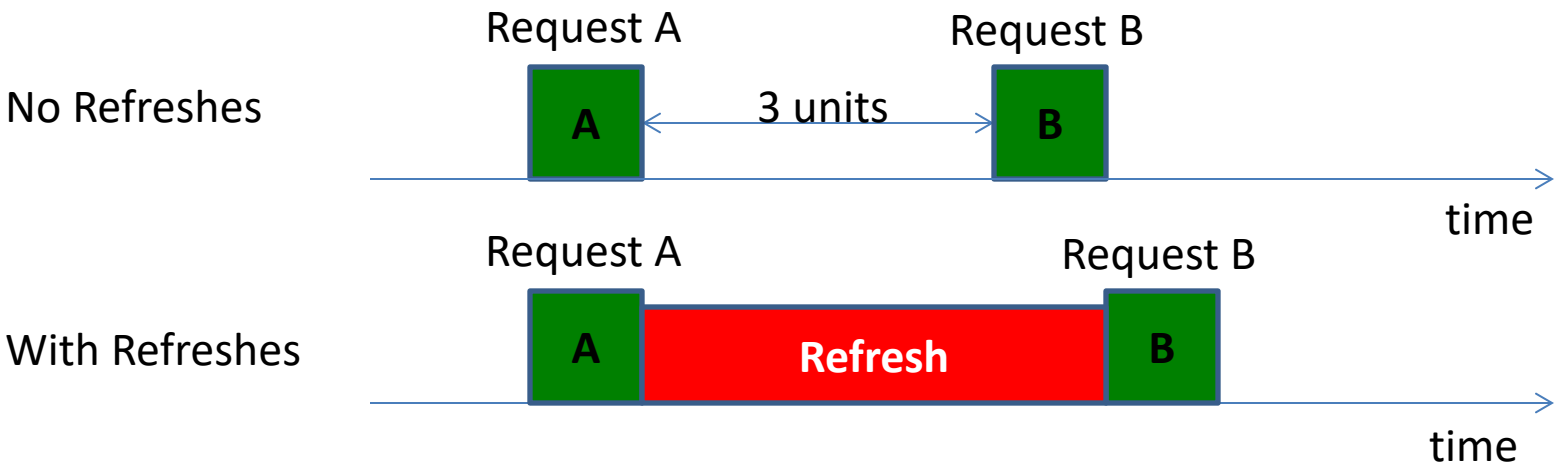

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

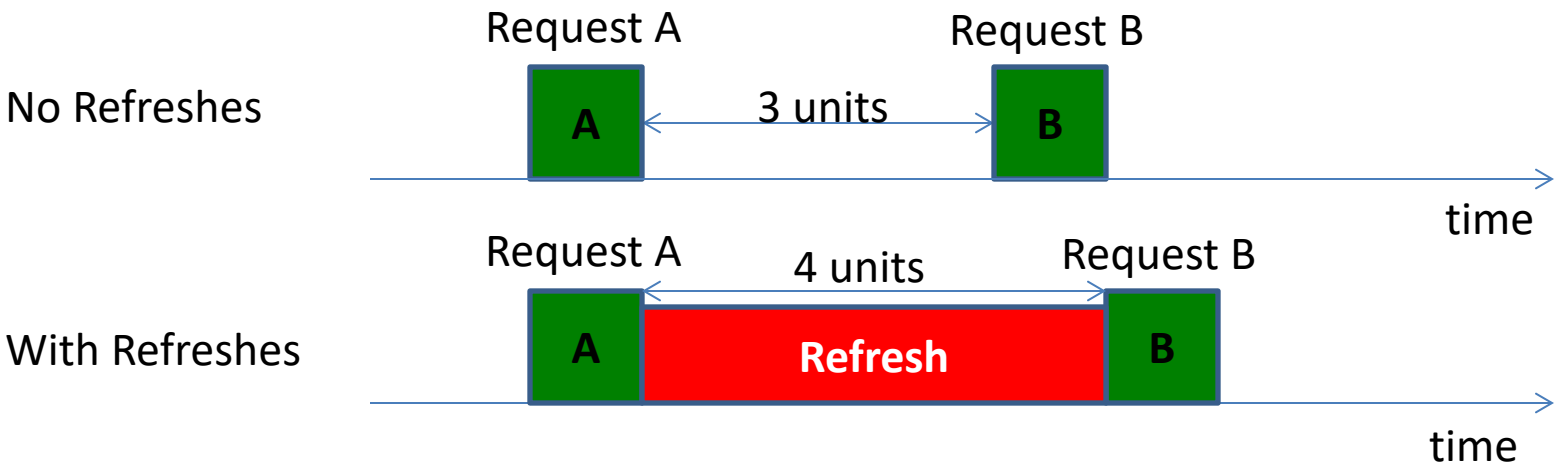

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

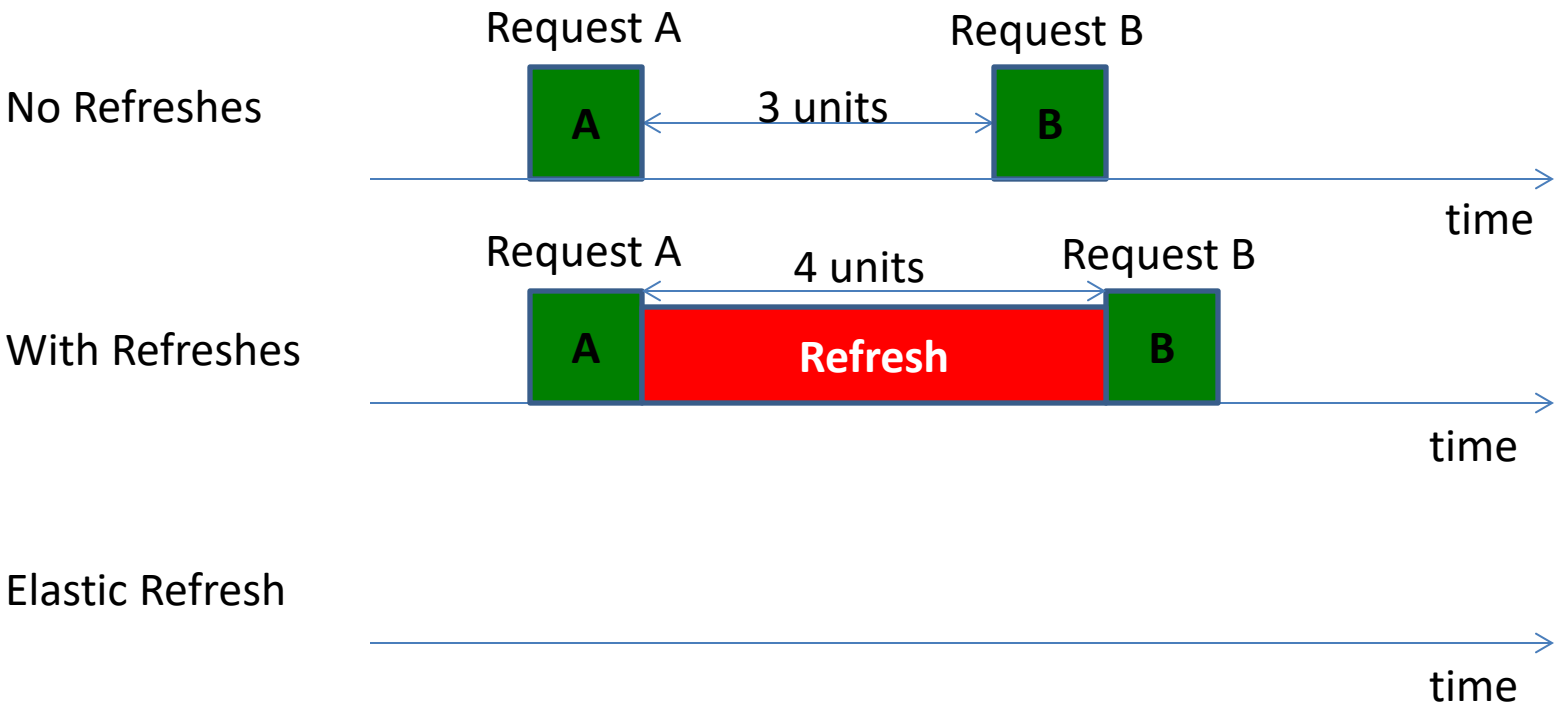

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

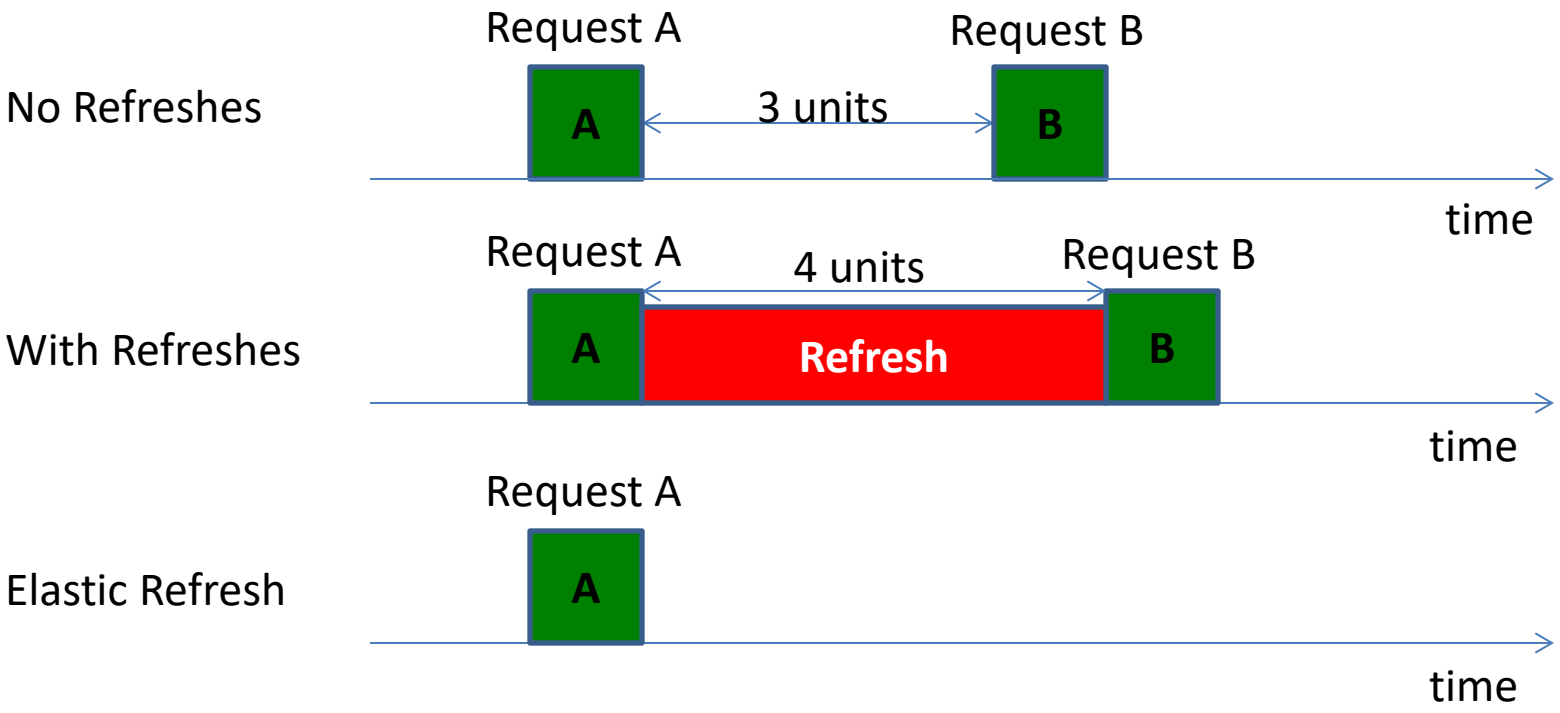

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

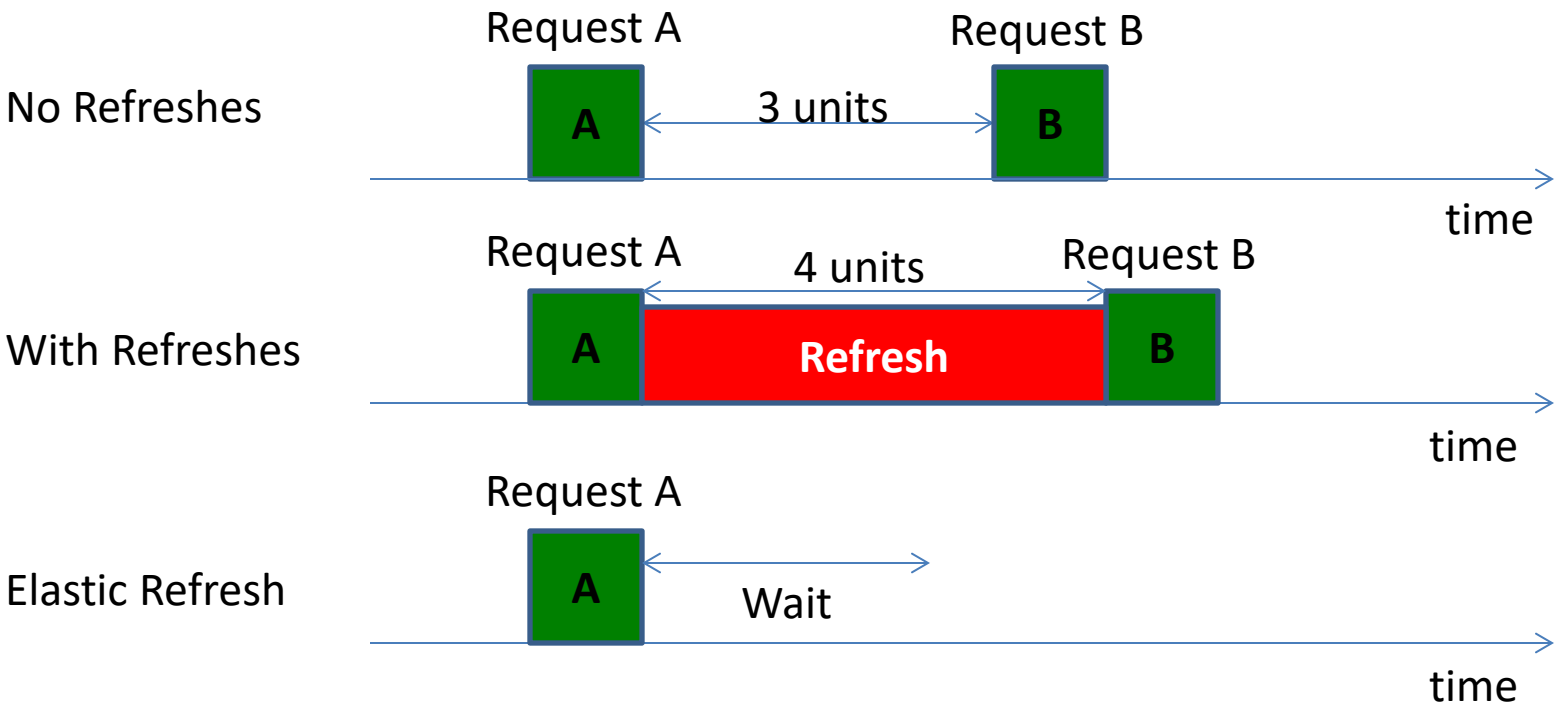

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

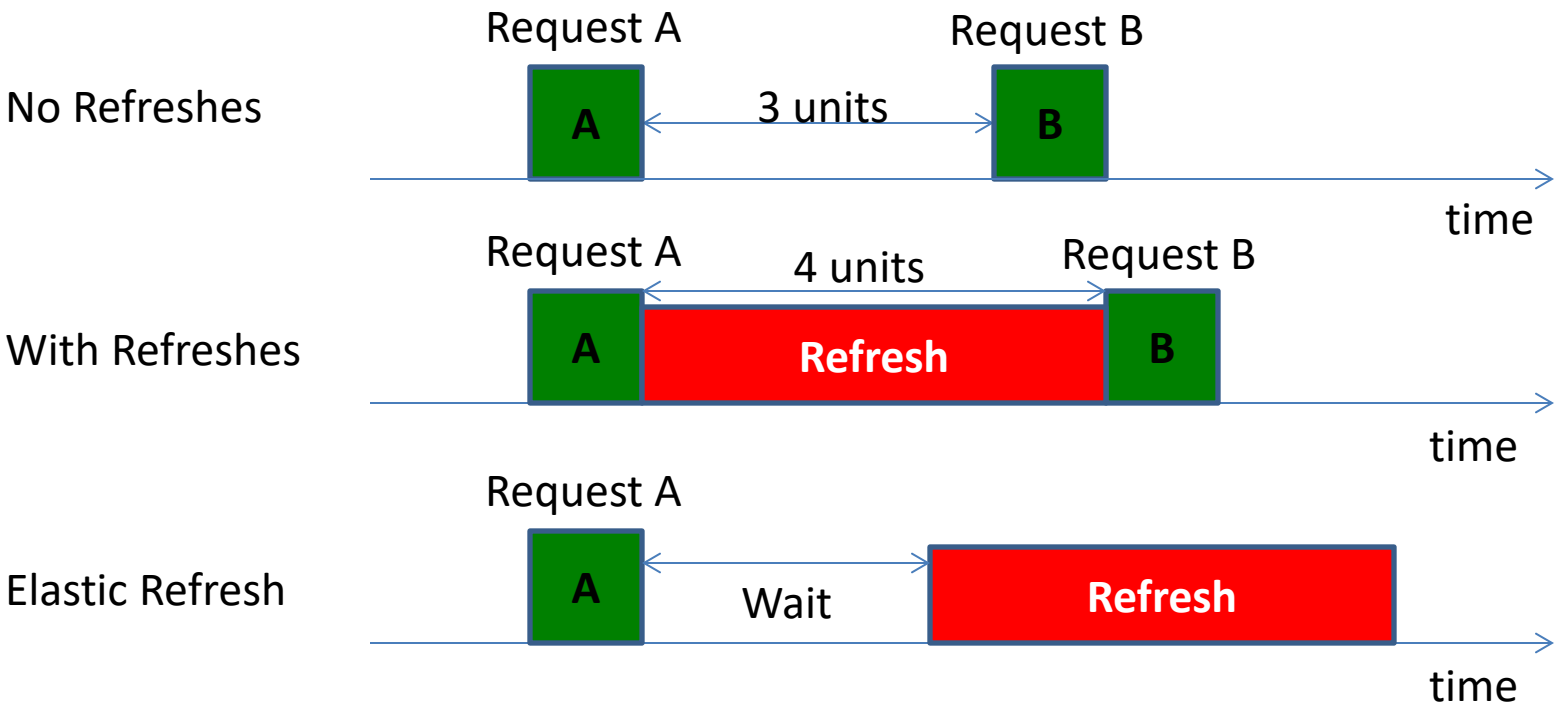

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

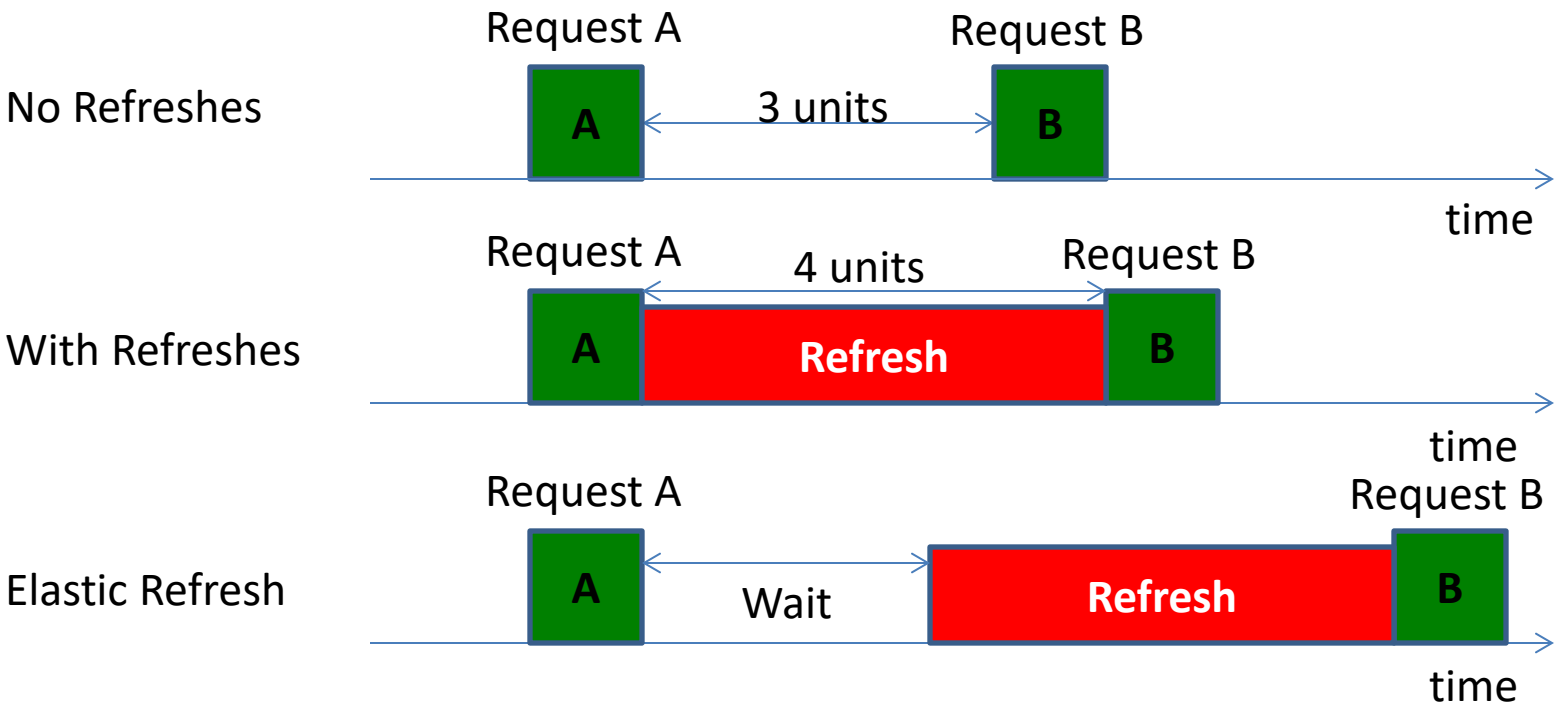

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

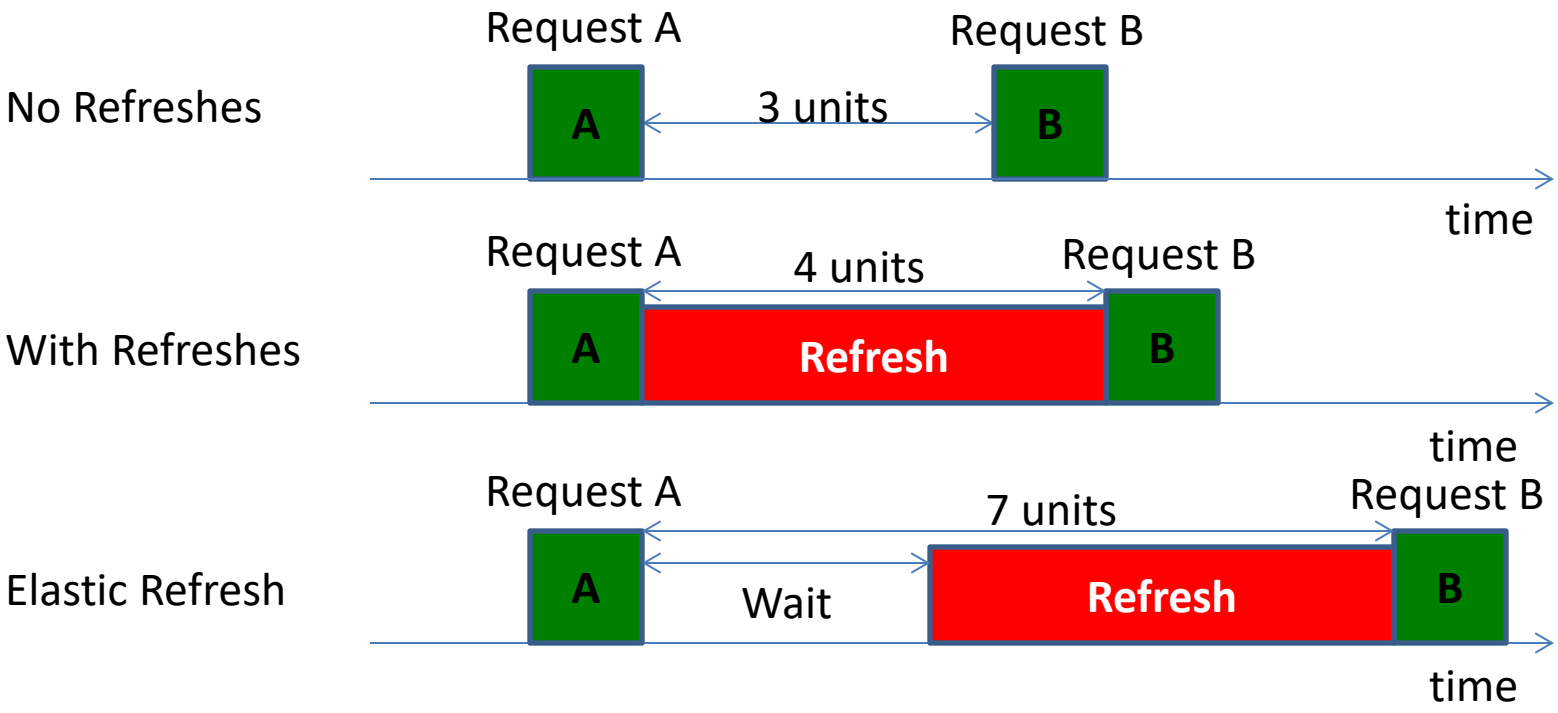

Elastic Refresh for Scheduling Refresh [MICRO'10]

- Elastic Refresh waits for idle period before issuing a refresh
- Estimates average inter-arrival time of memory request

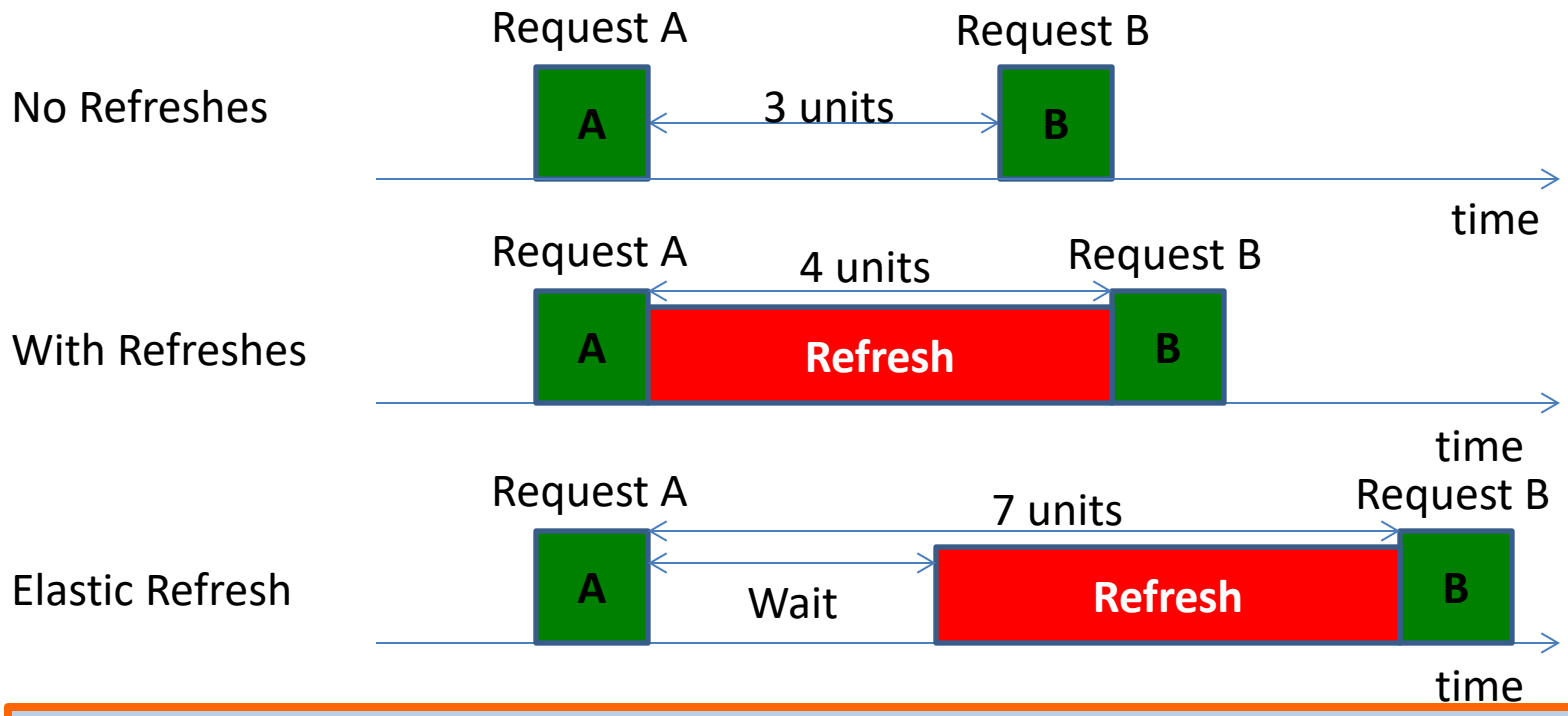

The "Wait and Watch" policy can increase wait times

Comparison with Elastic Refresh

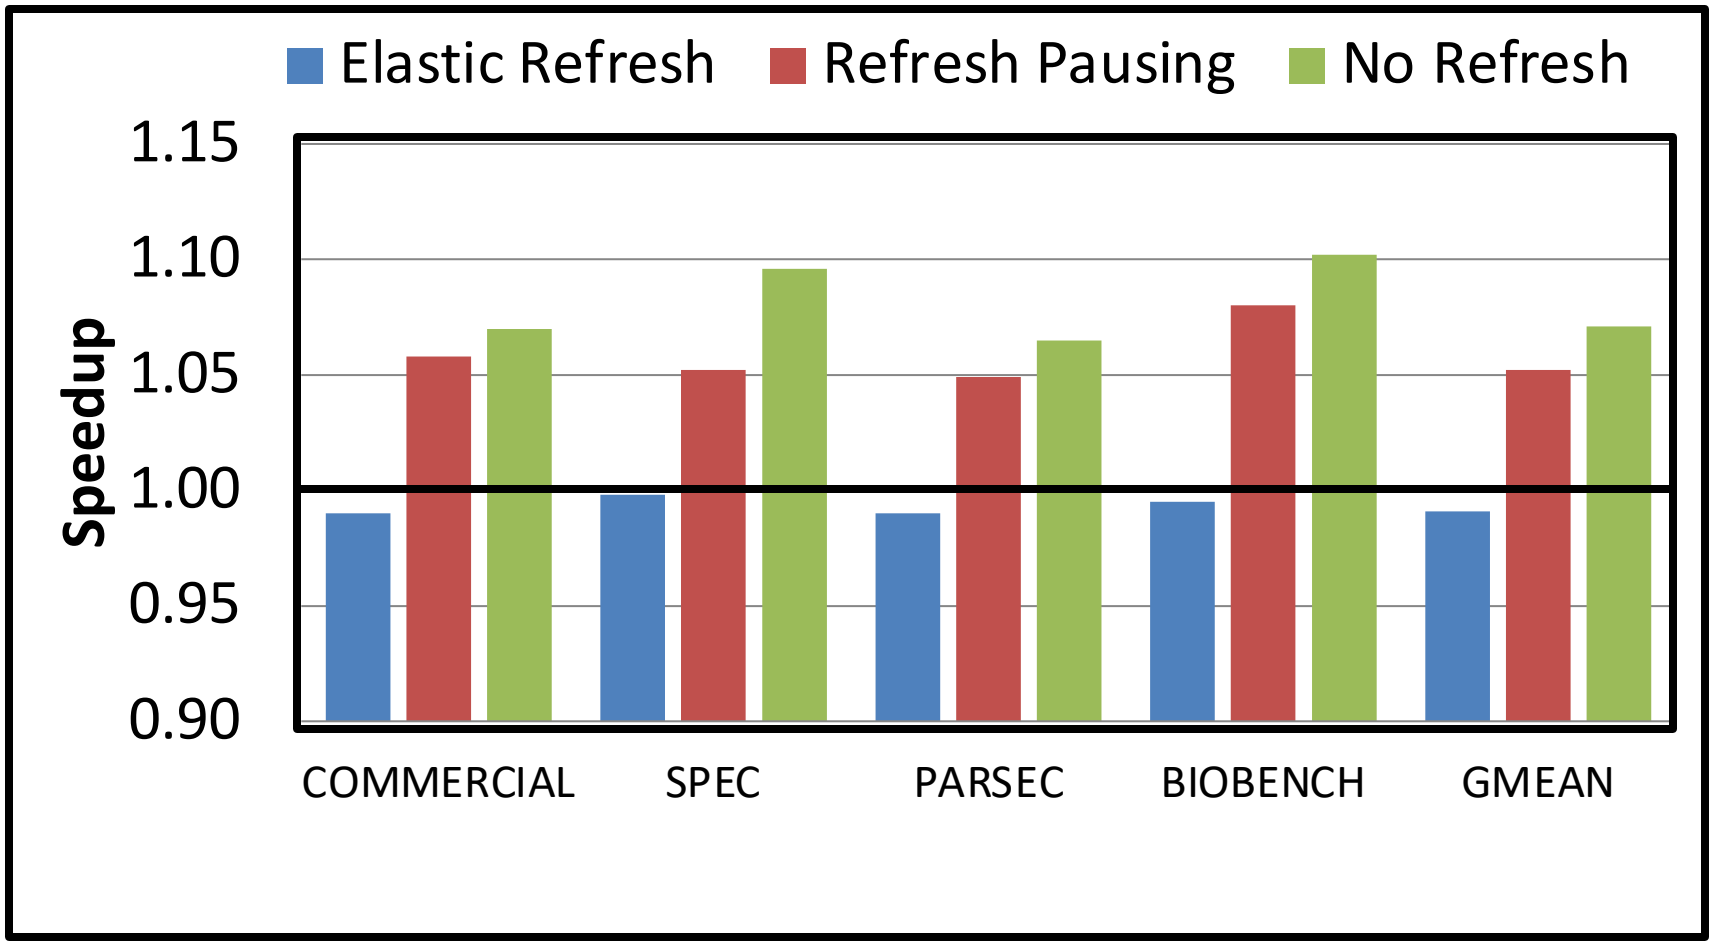

Refresh Pausing outperforms Elastic Refresh

DDR4 proposals: x2 and x4 modes

Reduce bundles size and have more bundles

DDR4 proposals: x2 and x4 modes

Reduce bundles size and have more bundles

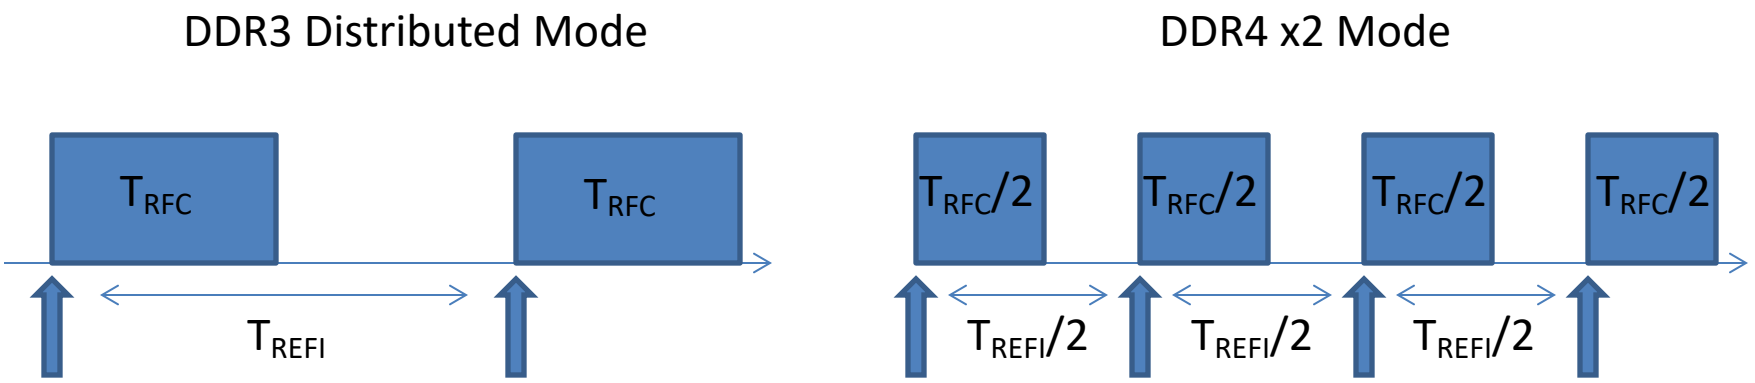

- In x2 mode, T_{REFI} is reduced by 2 (x4 mode by 4)
- In x2 mode T_{RFC} is reduced by 2 (x4 mode by 4)

DDR4 proposals: x2 and x4 modes

Reduce bundles size and have more bundles

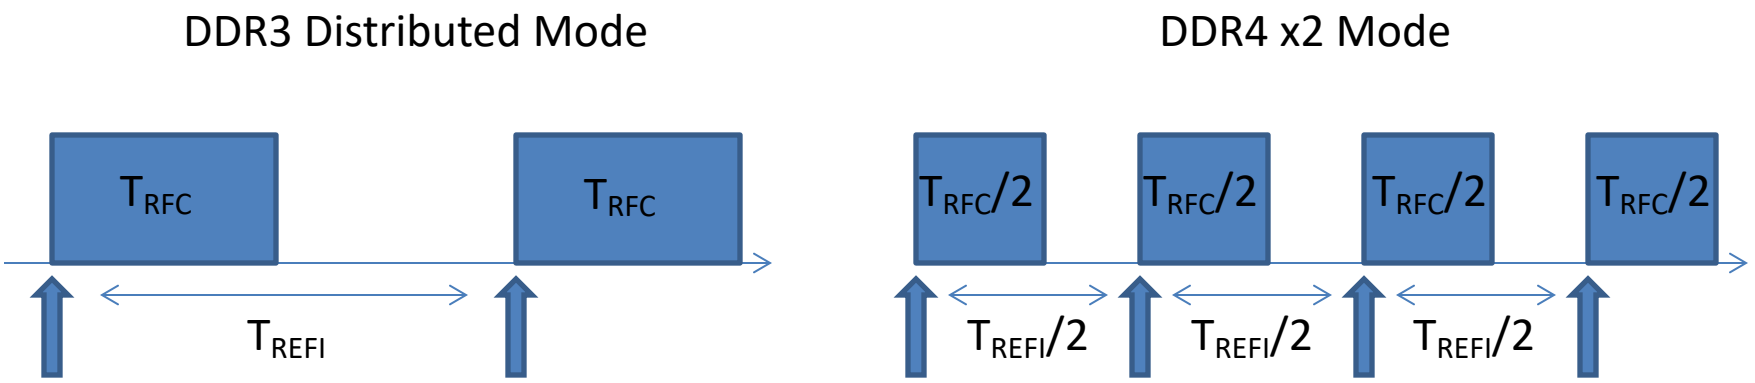

- In x2 mode, T_{REFI} is reduced by 2 (x4 mode by 4)
- In x2 mode T_{RFC} is reduced by 2 (x4 mode by 4)

Fine Grained Refresh to reduce contention of Refresh

Comparison with DDR4

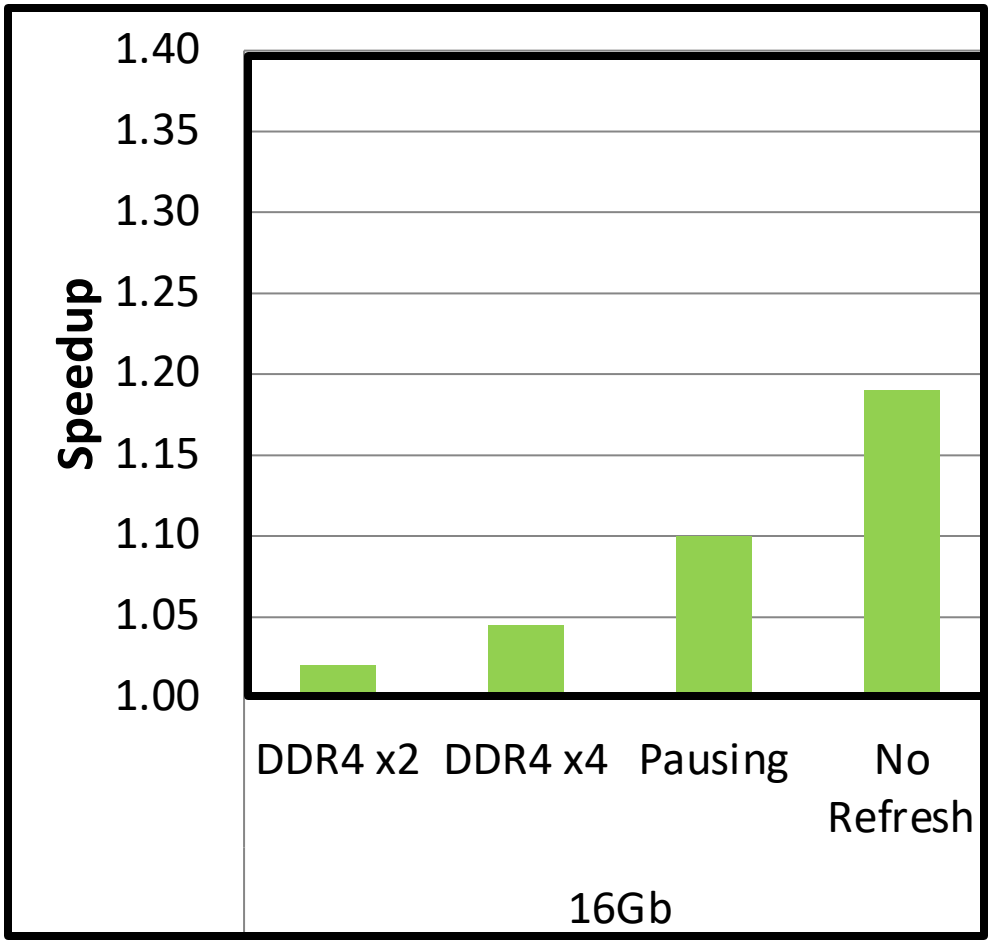

DDR4 modes (x2 and x4) useful but not enough

Comparison with DDR4

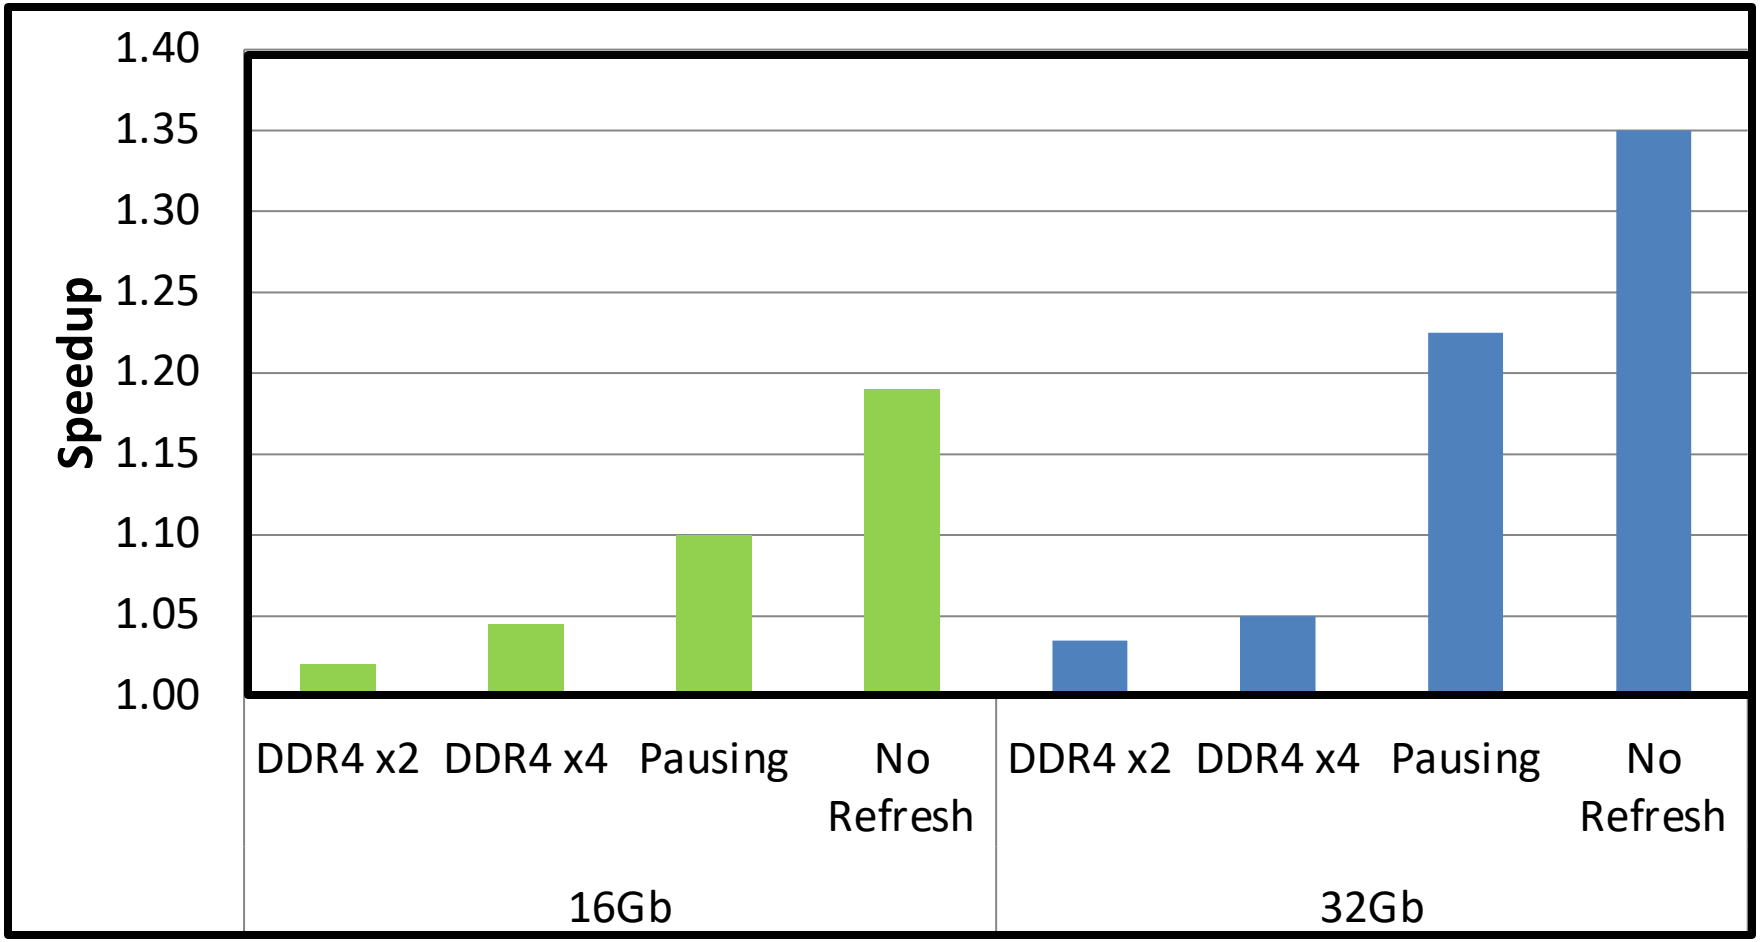

DDR4 modes (x2 and x4) useful but not enough

Outline

- Introduction & Motivation
- Refresh Operation: Background
- Refresh Pausing
- Evaluation
- Alternative Proposals
- **Summary**

Summary

- DRAM relies on Refresh for data integrity
- Time for Refresh increases with chip density
- Refresh blocks read, increases read latency
- Refresh Pausing: make Refresh Interruptible
- Pausing provides 5% improvement for 8Gb, increases with higher density
- Applicable also to DDR4 (fine grained refresh)