On the Off-chip Memory Latency of Real-Time Systems: Is DDR DRAM Really the Best Option?

Mohamed Hassan

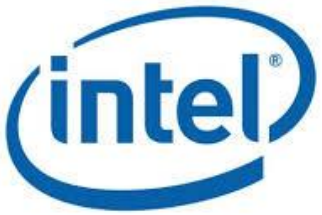


Motivation
01
DRAMs
02
PREDICTABILITY
03
RLDRAM
04
Results
05

Outline

- Historically, SRAMs have been the option for real-time safety-critical embedded systems

- With the increase in data demand
  → the cost became unaffordable

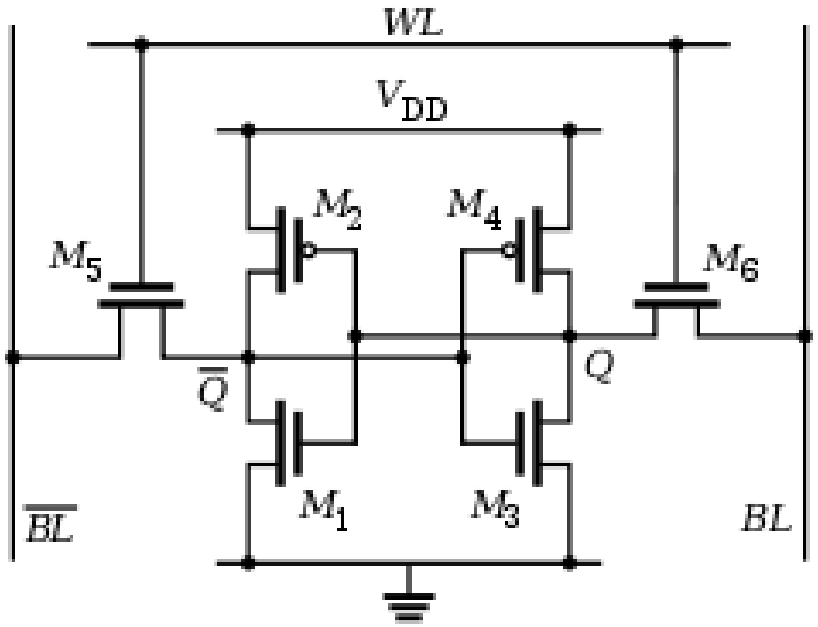

SRAMs

MOTIVATION

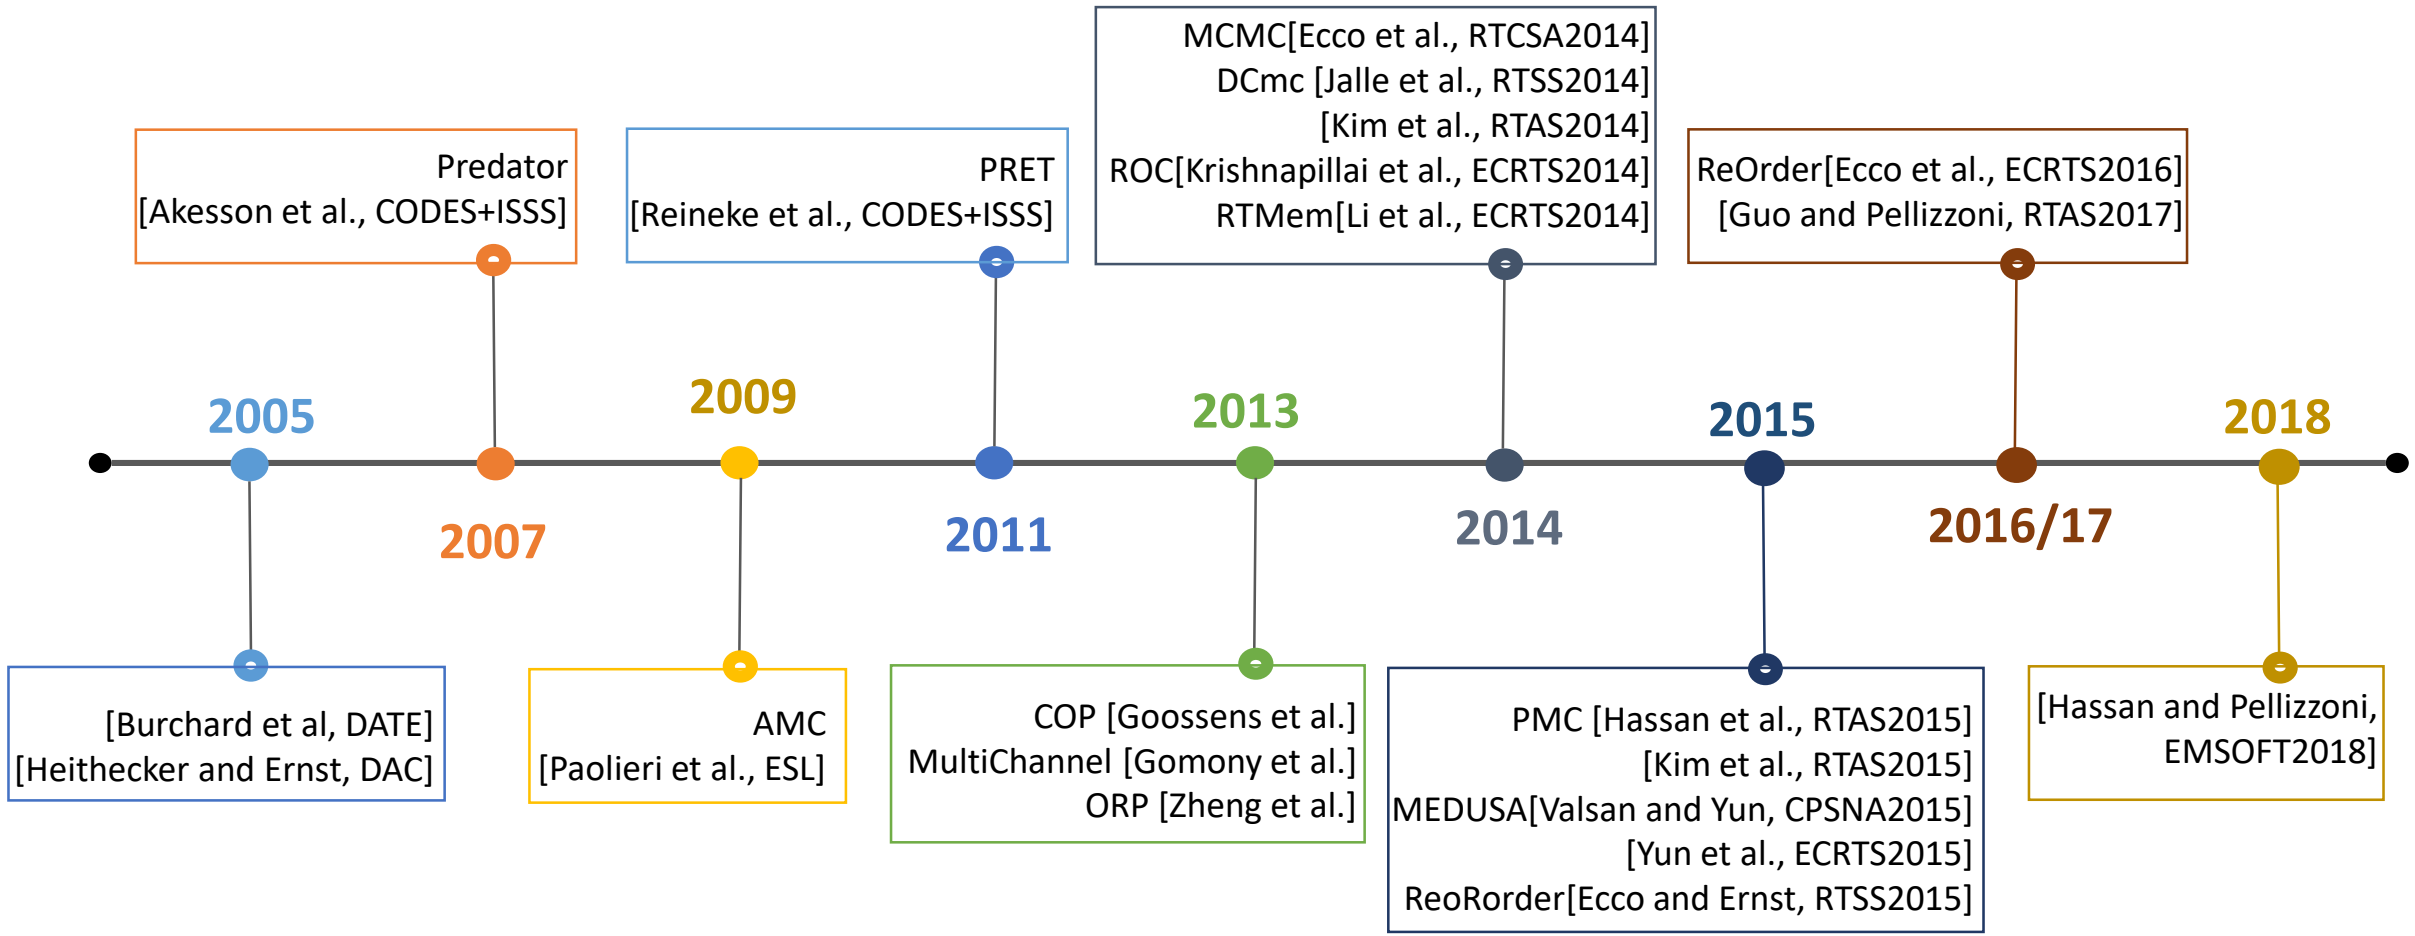

Work in Off-chip Memory

MOTIVATION

Predator
[Akesson et al., CODES+ISSS]

PRET
[Reineke et al., CODES+ISSS]

MCMC[Ecco et al., RTCSA2014]
DCmc [Jalle et al., RTSS2014]
[Kim et al., RTAS2014]
ROC[Krishnapillai et al., ECRTS2014]
RTMem[Li et al., ECRTS2014]

ReOrder[Ecco et al., ECRTS2016]
[Guo and Pellizzoni, RTAS2017]

All in Double Data Rate (DDR) DRAMs

[Burchard et al, DATE]
[Heithecker and Ernst, DAC]

AMC
[Paolieri et al., ESL]

COP [Goossens et al.]
MultiChannel [Gomony et al.]
ORP [Zheng et al.]

PMC [Hassan et al., RTAS2015]
[Kim et al., RTAS2015]
MEDUSA[Valsan and Yun, CPSNA2015]
[Yun et al., ECRTS2015]
ReoRorder[Ecco and Ernst, RTSS2015]

[Hassan and Pellizzoni, EMSOFT2018]

Work in Off-chip Memory

MOTIVATION

- DDR DRAM is the commodity off-chip memory, *Why?*
  - Low cost
  - Large capacity
  - High BW
- What is the most important requirement for real-time/safety-critical systems?
  - Yes, Predictability
- How is DDRx for predictability?
  - DDRx Random Access Memories are not Random at all!!
  - Access latency varies notably based on many factors
    - access patterns
    - transaction type (read or write)
    - DRAM state from previous accesses

A Context about DDRs

- Multiplexed address mode:
- The address bits are split into two segments provided to the device in two stages:
  1. Row address → row decoder
  2. Column address → column decoder

✓ Low cost (less pin count)

✘ High latency

✘ Huge variability

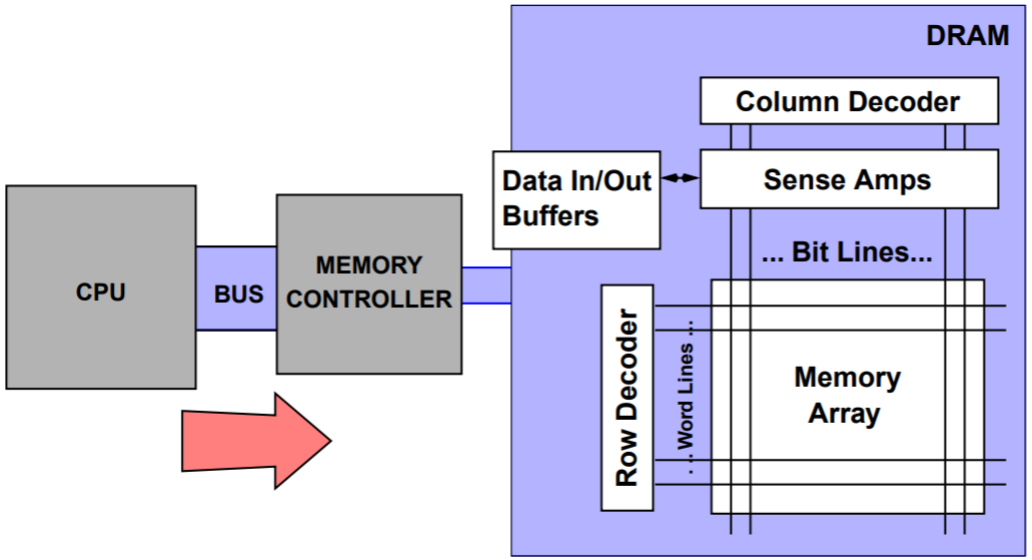

Background

DRAM

- A request in general can consist of one, two, or three commands:
  - ACTIVATE (A) command:
    - Bring data row from cells into sense amplifiers

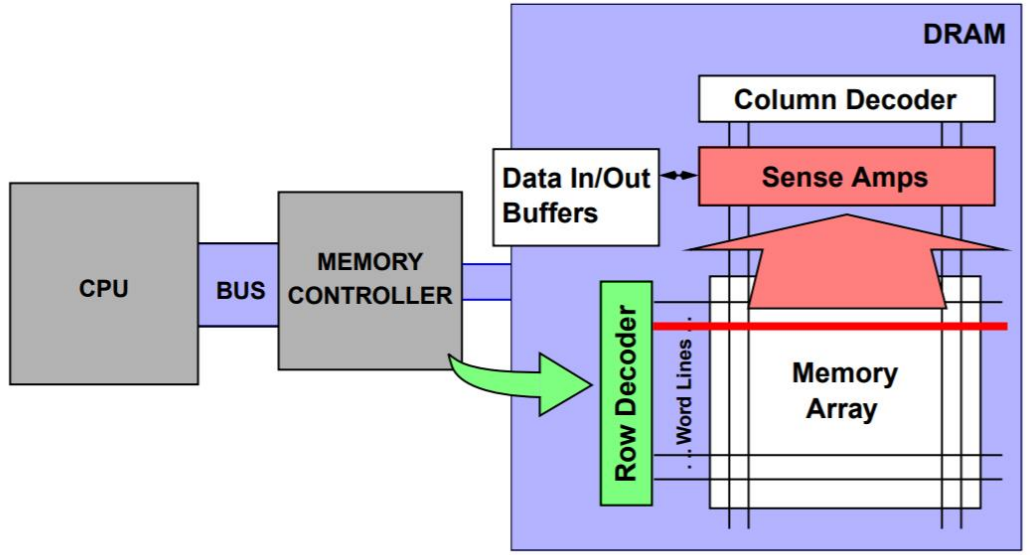

Background

DRAM

- DRAM Consists of multiple banks
- The memory controller (MC) manages accesses to DRAM
- A request in general consists of:
  - ACTIVATE (A) command:
    - Bring data row from cells into sense amplifiers
  - Read/Write (R/W) commands:
    - To read/write from specific columns in the sense amplifiers

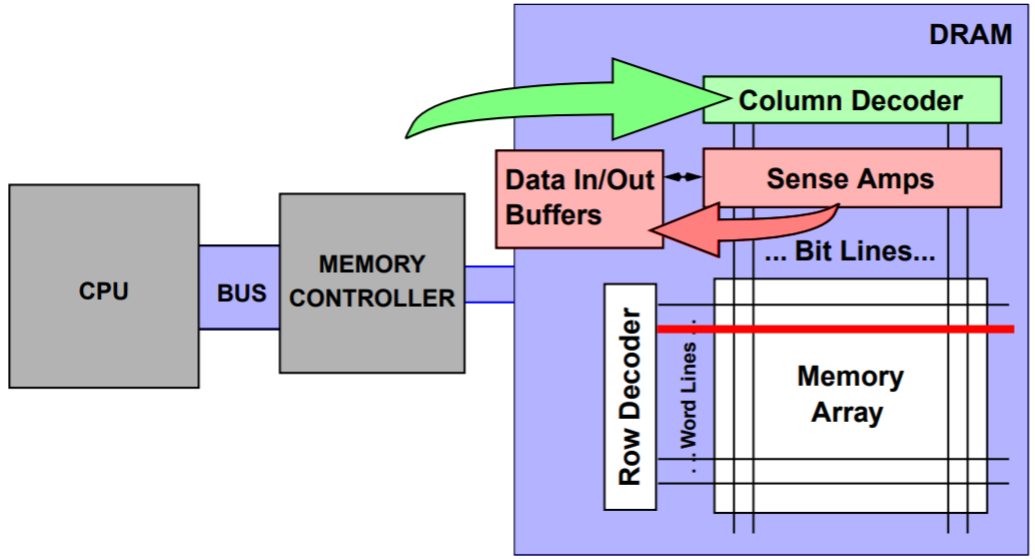

Background

DRAM

- DRAM Consists of multiple banks
- The memory controller (MC) manages accesses to DRAM
- A request in general consists of:
  - ACTIVATE (A) command:
    - Bring data row from cells into sense amplifiers
  - Read/Write (R/W) commands:
    - To read/write from specific columns in the sense amplifiers
  - PRECHARGE (P) command:
    - to write back a previous row in the sense amplifiers before bringing the new one

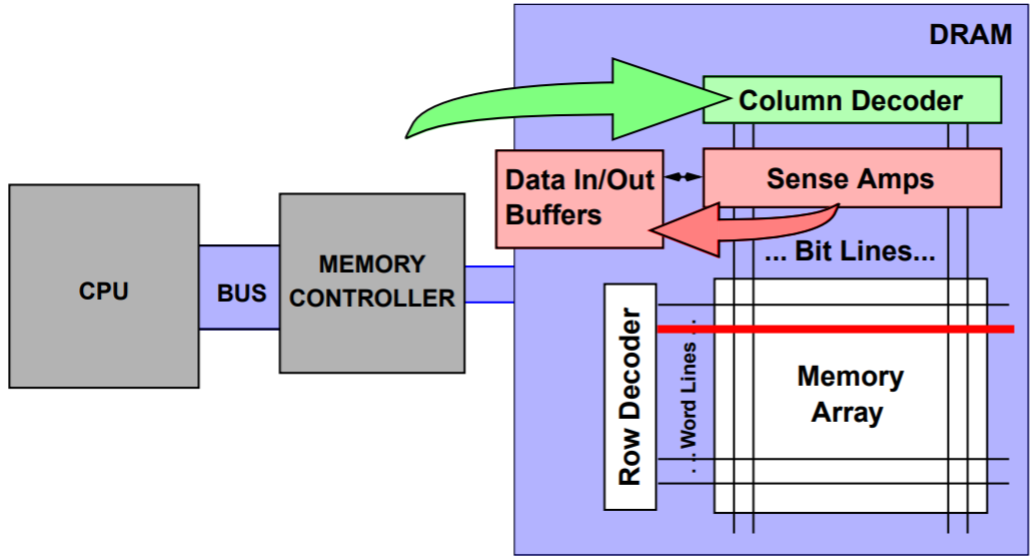

Background

DRAM

- DRAM Consists of multiple banks
- The memory controller (MC) manages accesses to DRAM
- A request in general consists of:
  - ACTIVATE (A) command:
    - Bring data row from cells into sense amplifiers
  - Read/Write (R/W) commands:
    - To read/write from specific columns in the sense amplifiers
  - PRECHARGE (P) command:
    - to write back a previous row in the sense amplifiers before bringing the new one

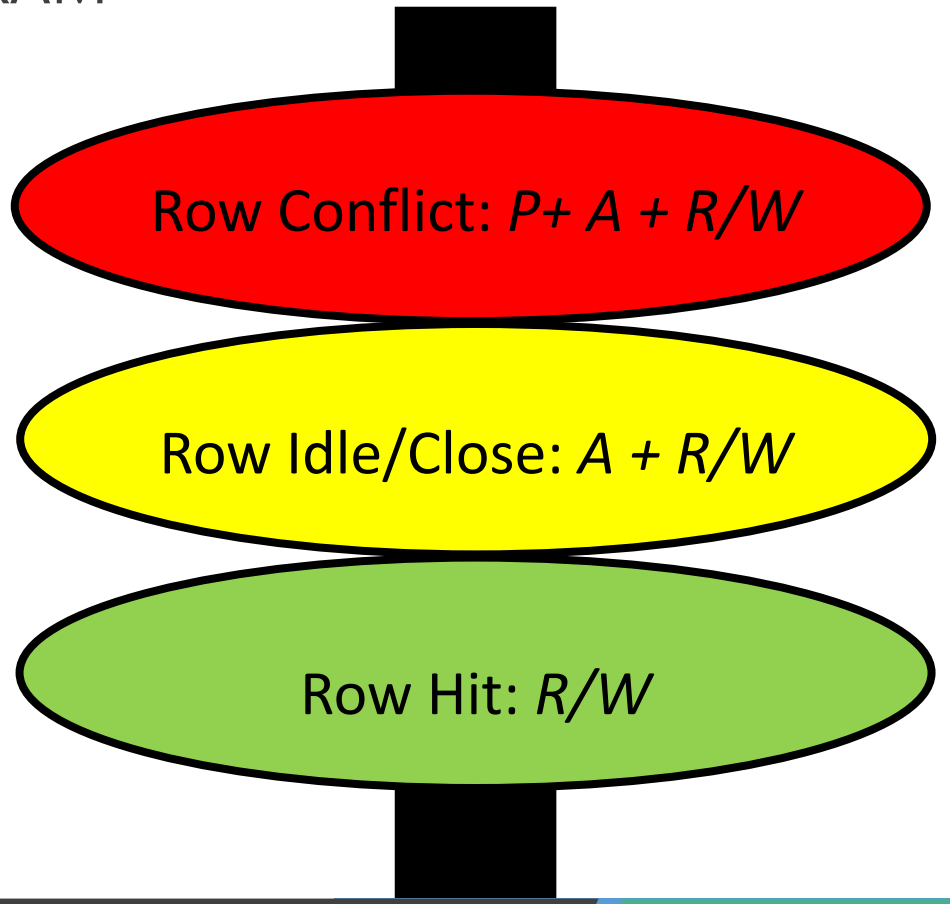

Background

DRAM

- DRAM Consists of multiple banks
- The memory controller (MC) manages accesses to DRAM
- A request in general consists of:
  - ACTIVATE command
  - R/W commands
  - PRECHARGE command
- All commands have associated timing constraints that have to be satisfied by the controller (20+ timing constraints)

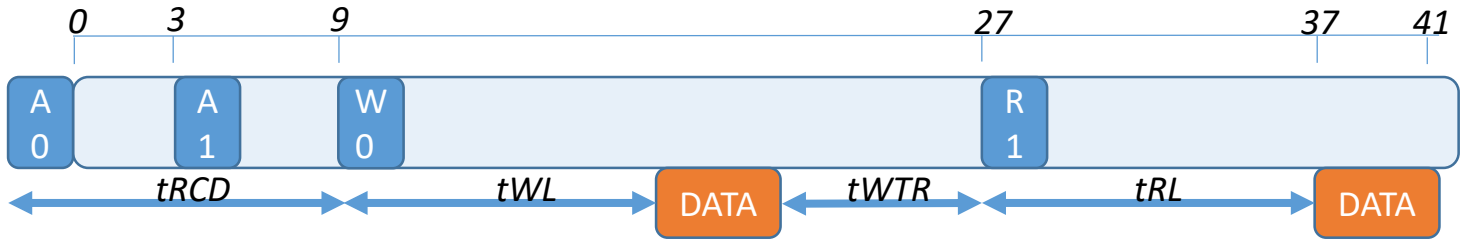

Background

DRAM

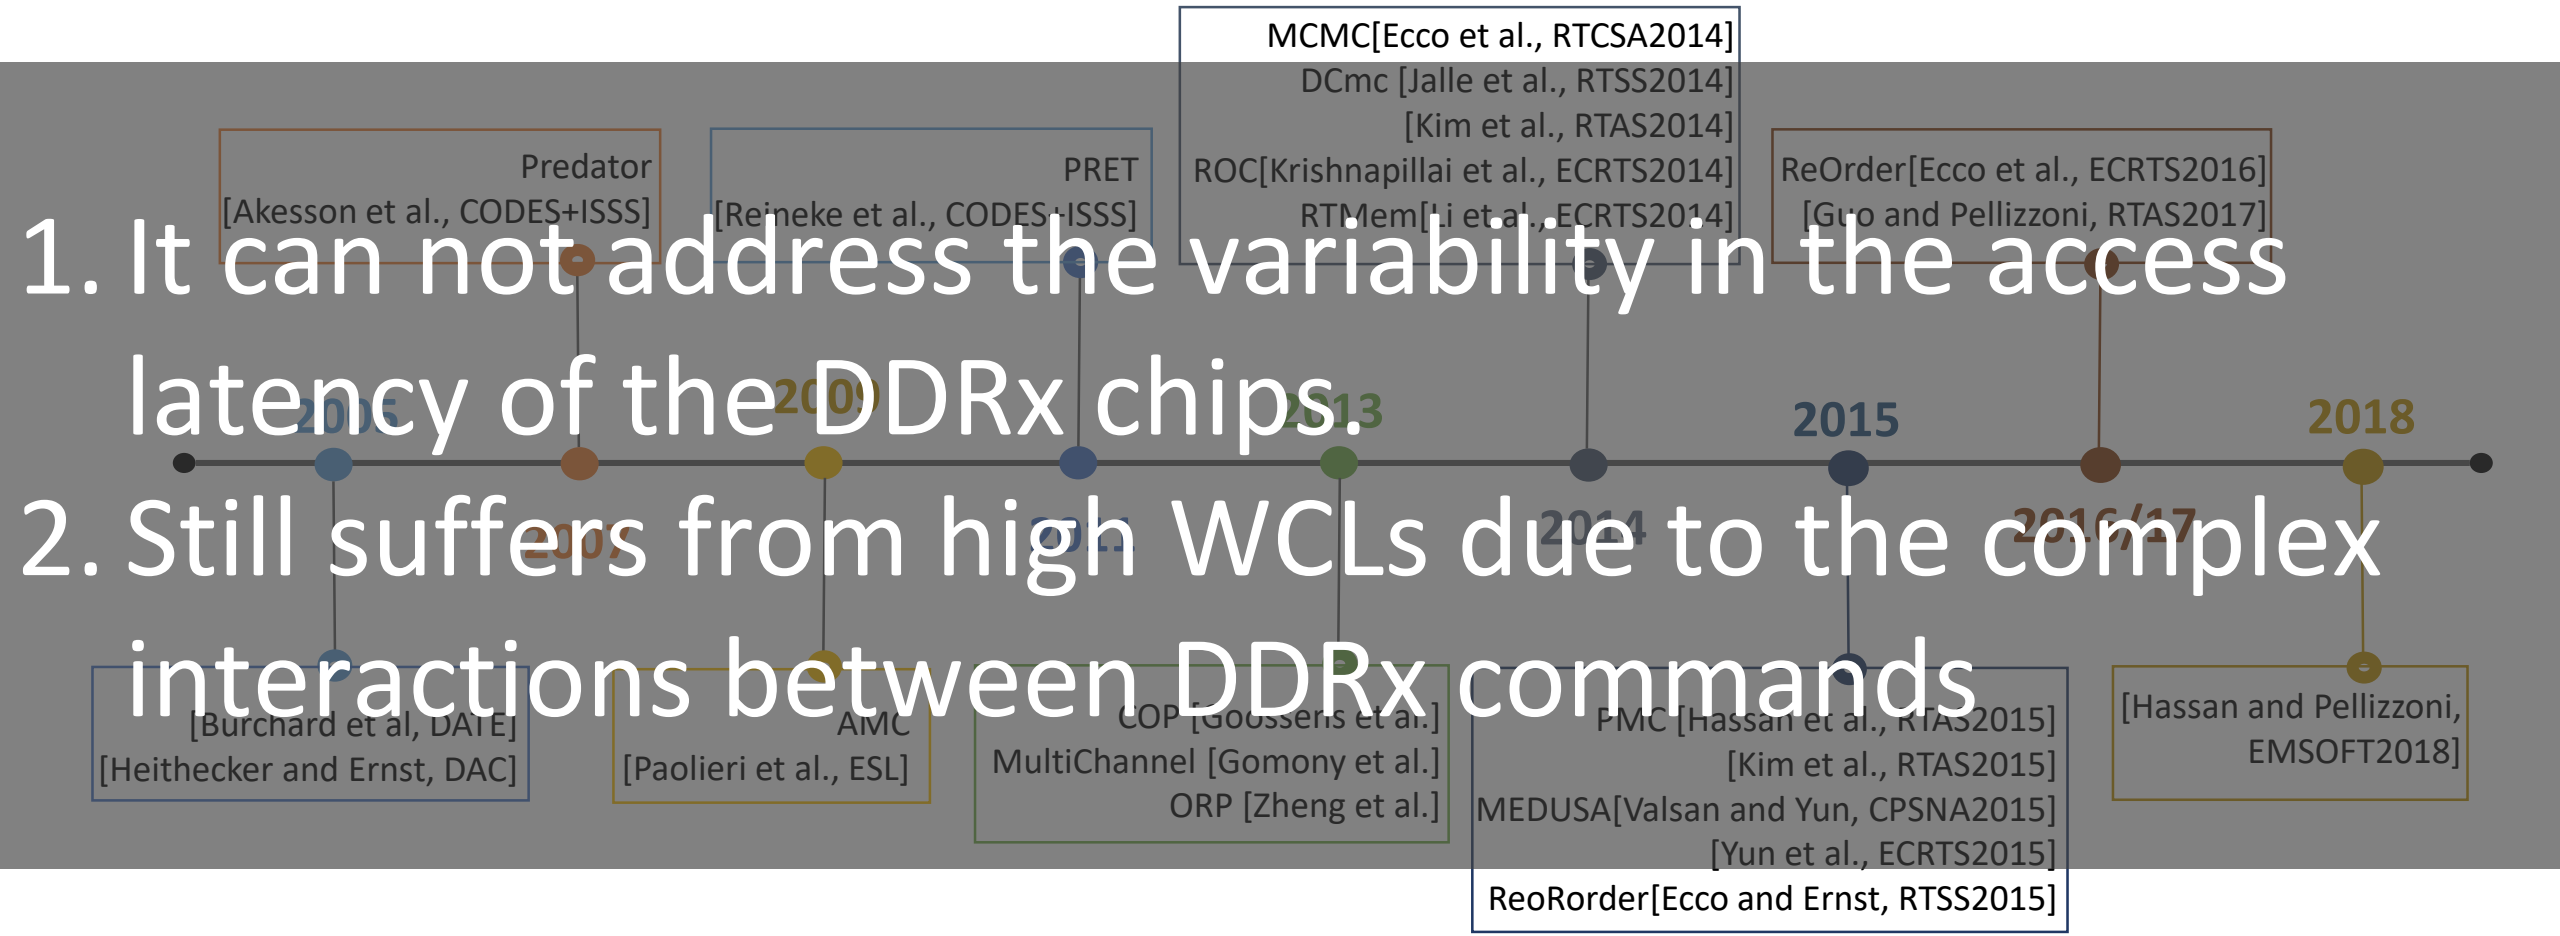

1. It can not address the variability in the access latency of the DDRx chips.
2. Still suffers from high WCLs due to the complex interactions between DDRx commands

Work in Off-chip Memory

MOTIVATION

- DDR DRAM is the commodity off-chip memory, *Why?*
  - Low cost
  - high capacity
  - High BW
- What is the most important requirement for real-time/safety-critical systems?
  - Yes, Predictability
- How is DDRx for predictability?
  - DDRx Random Access Memories are not Random at all!!
  - access latency varies notably based on many factors
    - access patterns
    - transaction type (read or write)
    - DRAM state from previous accesses

} ← ***Comprehensively study these factors***

Assessing DDRs for Predictability

MOTIVATION

- How is DDRx for predictability?

  - Predictability has different definitions in the real-time literature
  - One important measure is ***the relative difference between best- and worst-case execution times*** (or latencies in case of memories) [Wilhelm et al, TECS08]
  - We define "**Variability Window**" (VW) to quantitatively measure the DRAM predictability

$$VW = \frac{WCL - BCL}{WCL} \times 100$$

Assessing DDRs for Predictability

PREDICTABILITY

- Targets an open row (only R command)
- (a) is best-case
- (e) arrives after a write to same rank

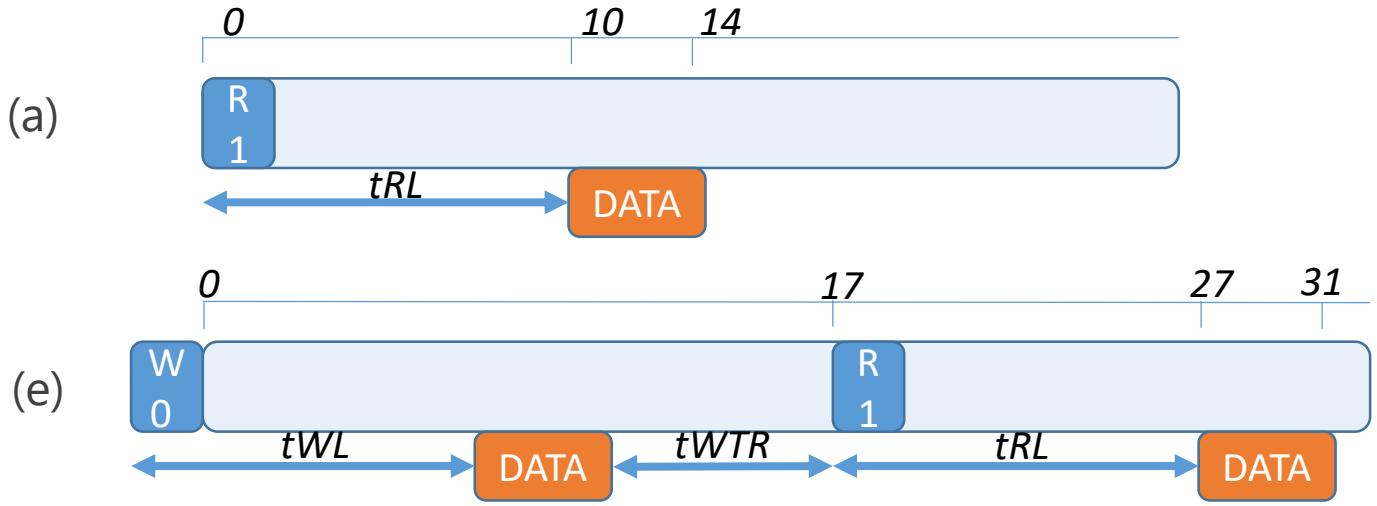

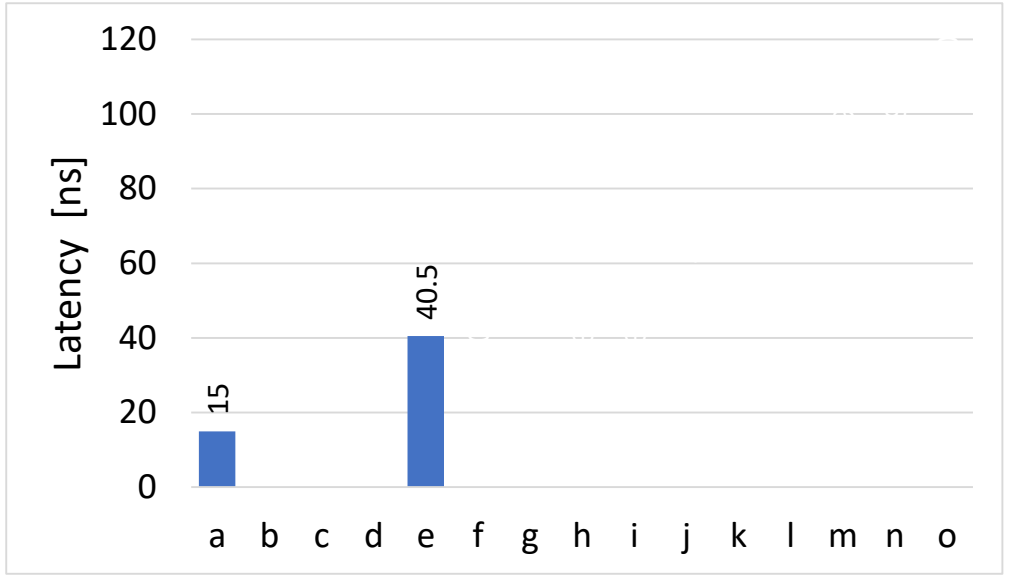

Assessing DDRs for Predictability

PREDICTABILITY

- Targets an open row (only R command)
- (a) is best-case
- (e) arrives after a write to same rank

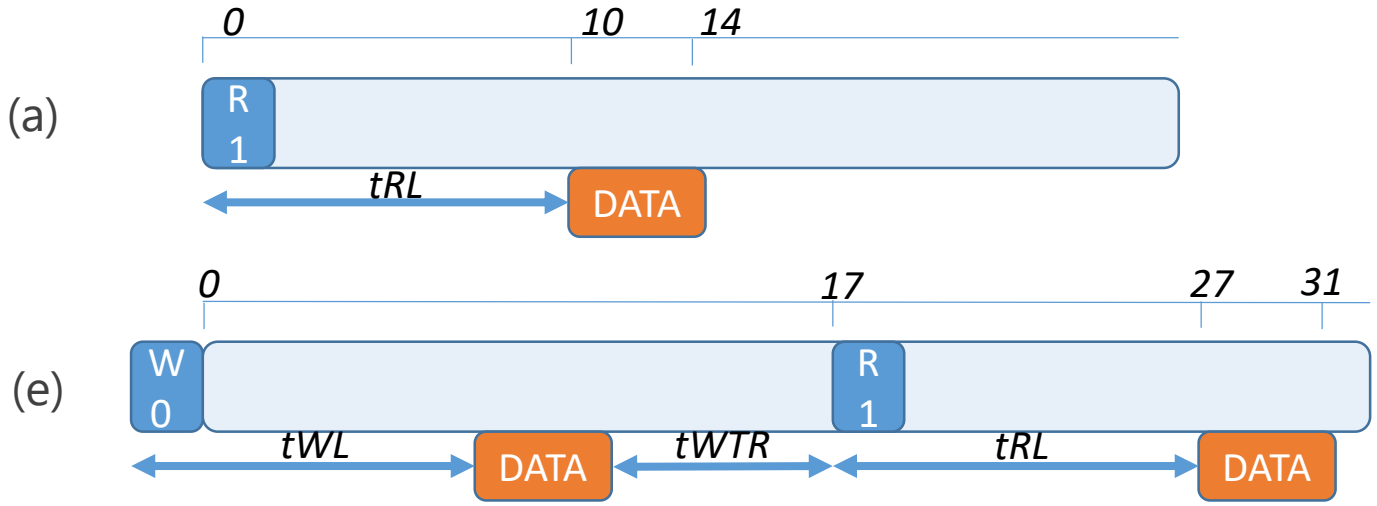

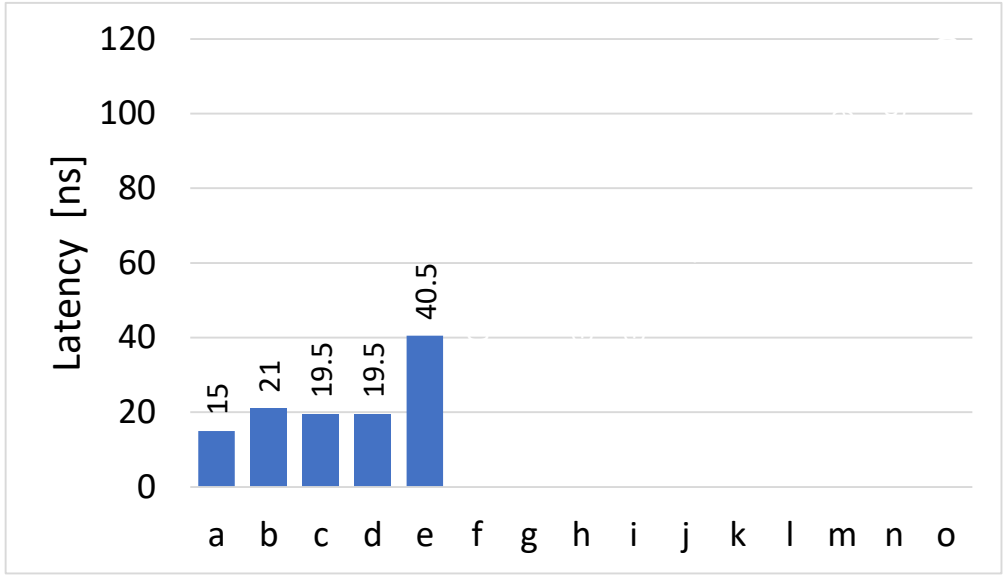

Assessing DDRs for Predictability

PREDICTABILITY

- Targets a close row (A + R command)
- (f) is best-case
- (j) arrives after a closed write to same rank

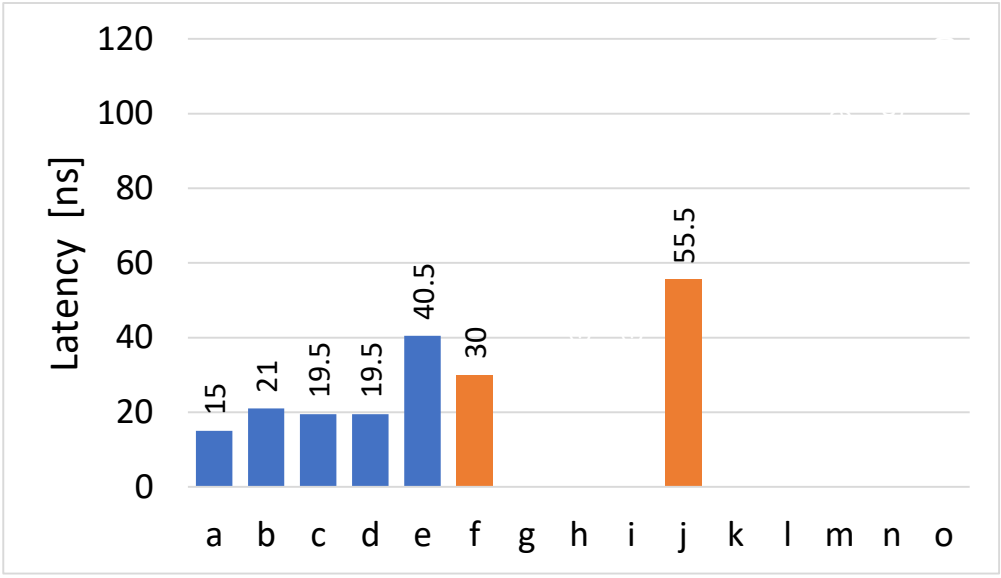

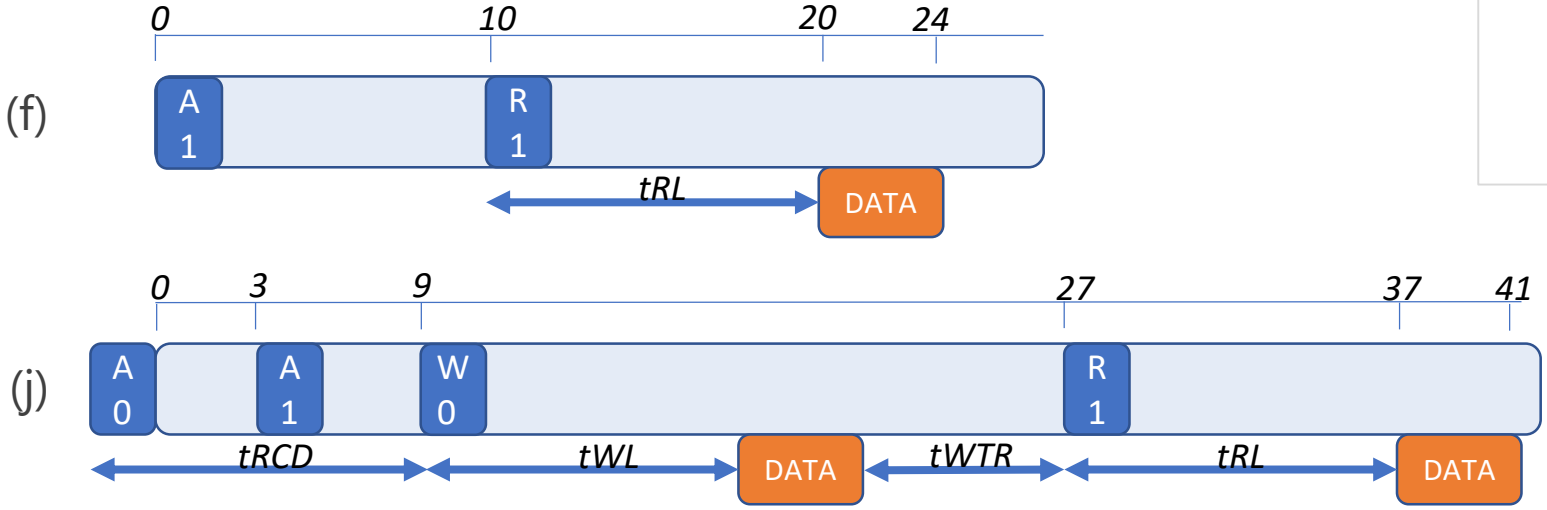

Assessing DDRs for Predictability

PREDICTABILITY

- Targets a close row (A + R command)
- (f) is best-case
- (j) arrives after a closed write to same rank

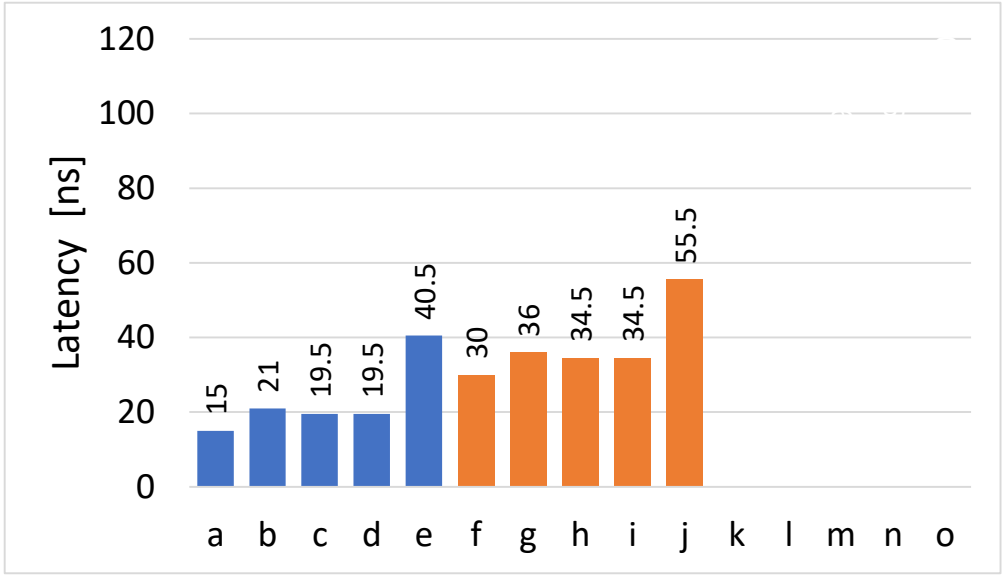

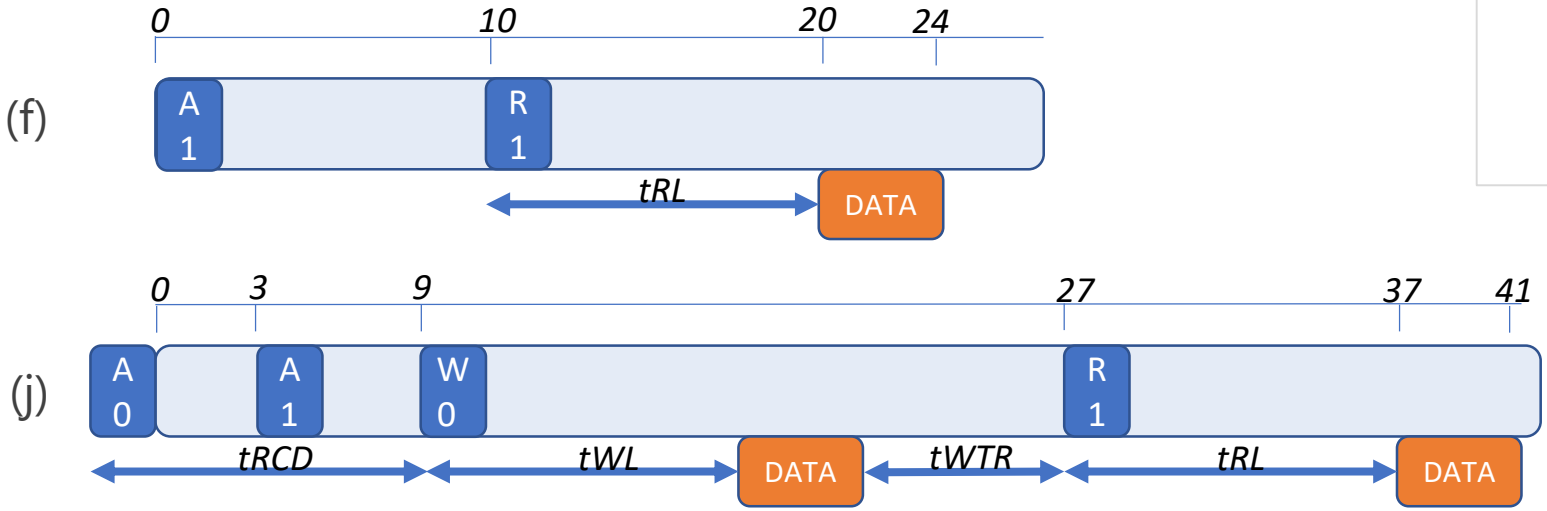

Assessing DDRs for Predictability

PREDICTABILITY

- Targets a conflict row (P + A + R command)
- (k) is best-case
- (o) arrives after a conflict write to same rank

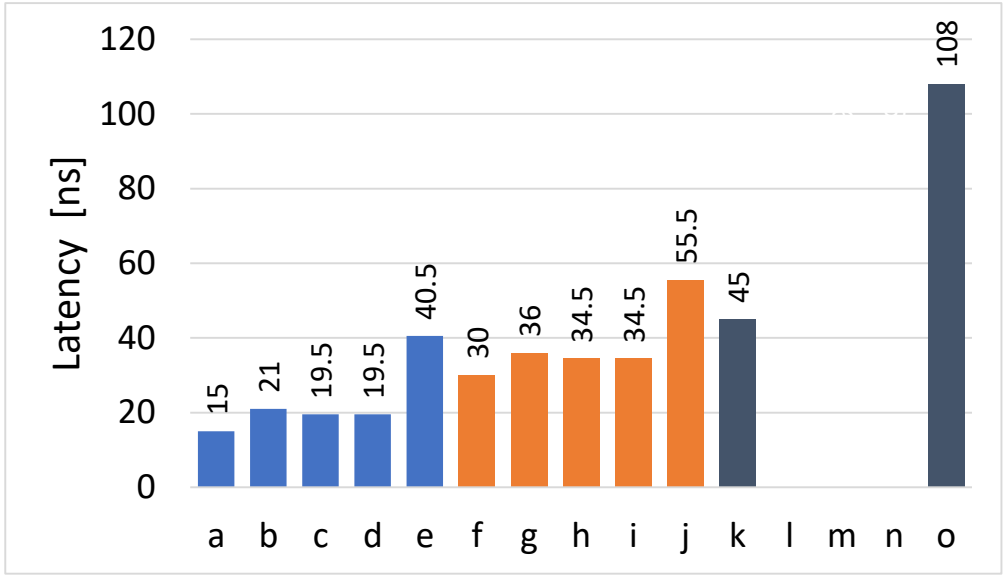

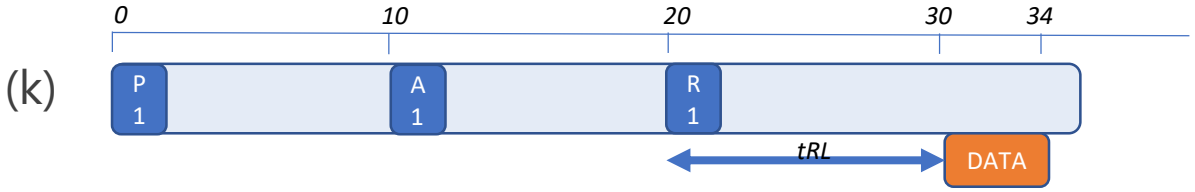

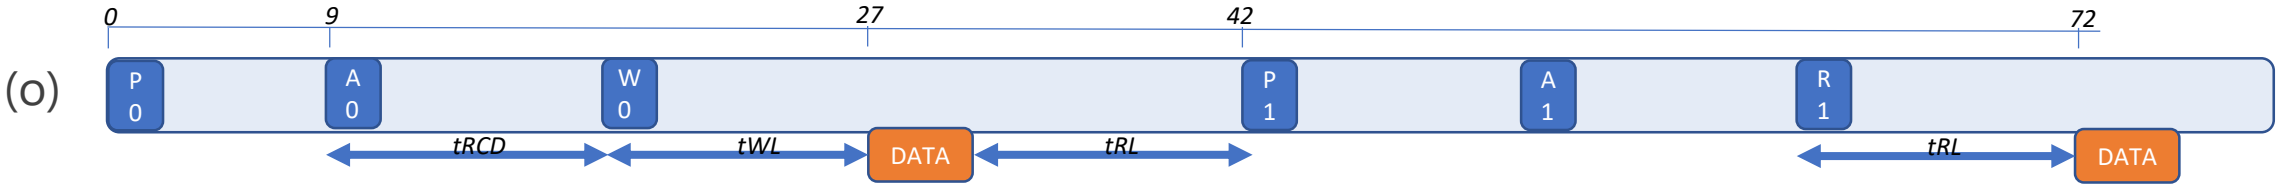

Assessing DDRs for Predictability

PREDICTABILITY

- Targets a conflict row (P + A + R command)
- (k) is best-case
- (o) arrives after a conflict write to same rank

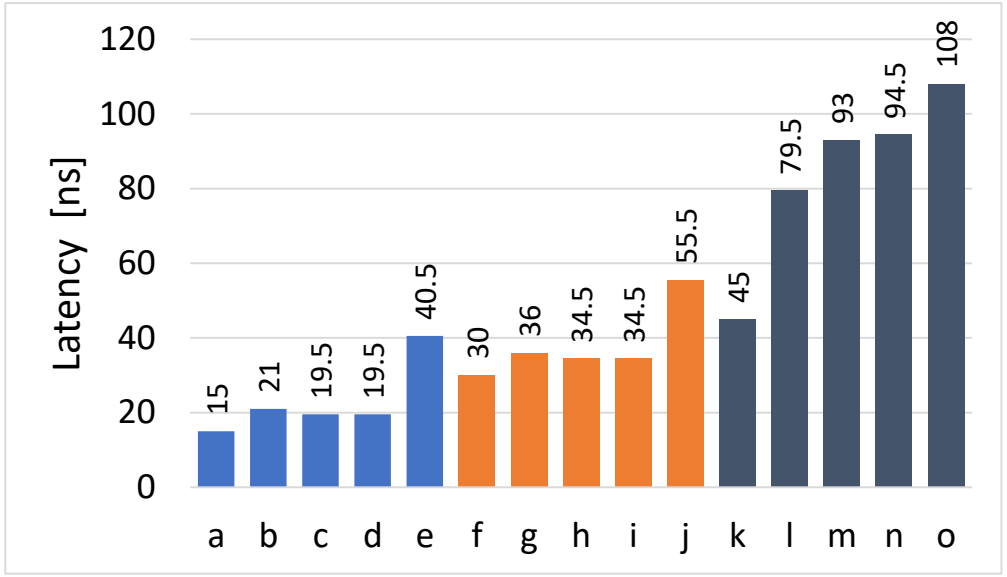

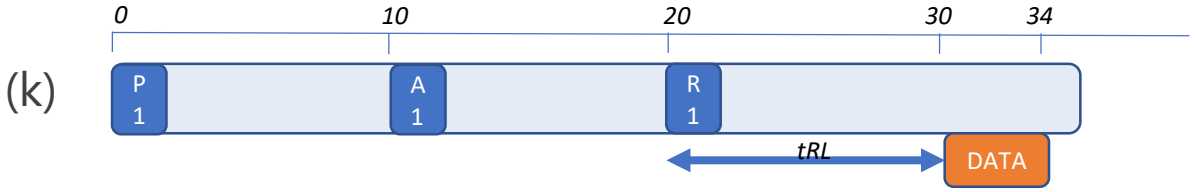

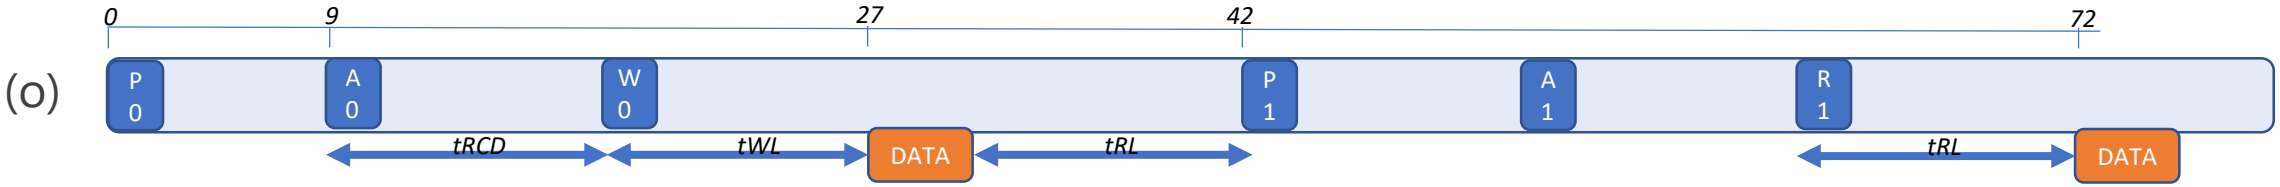

Assessing DDRs for Predictability

PREDICTABILITY

- 15 cases for a read request
- Another 15 for a write request

VW= 620%
WCL=108ns

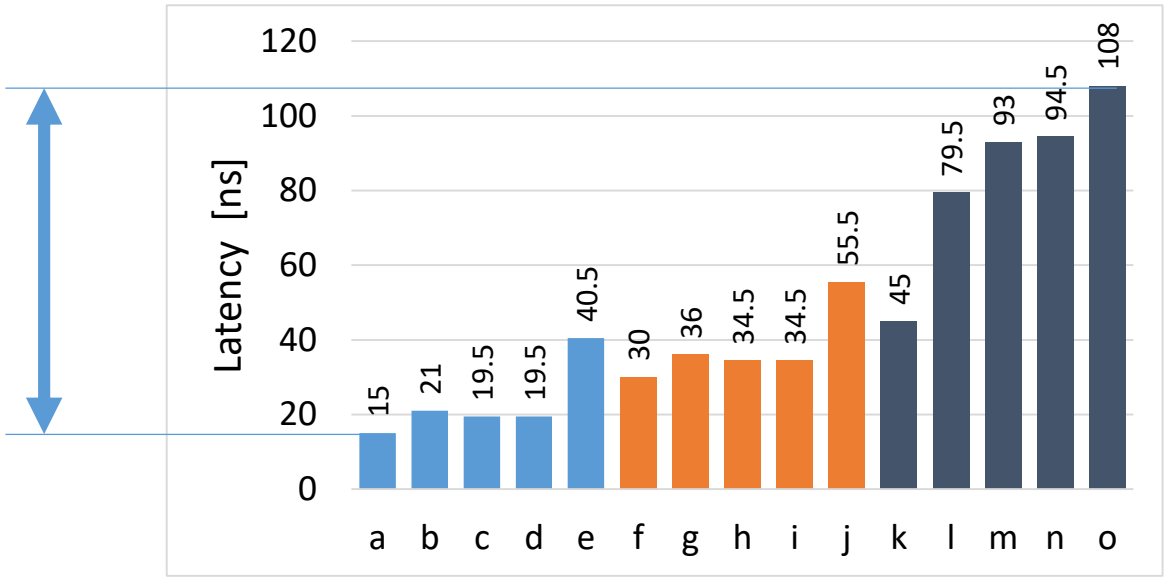

Assessing DDRs for Predictability

- We calculate the VW for 8 of the state-of-the-art DDRx DRAM Controllers

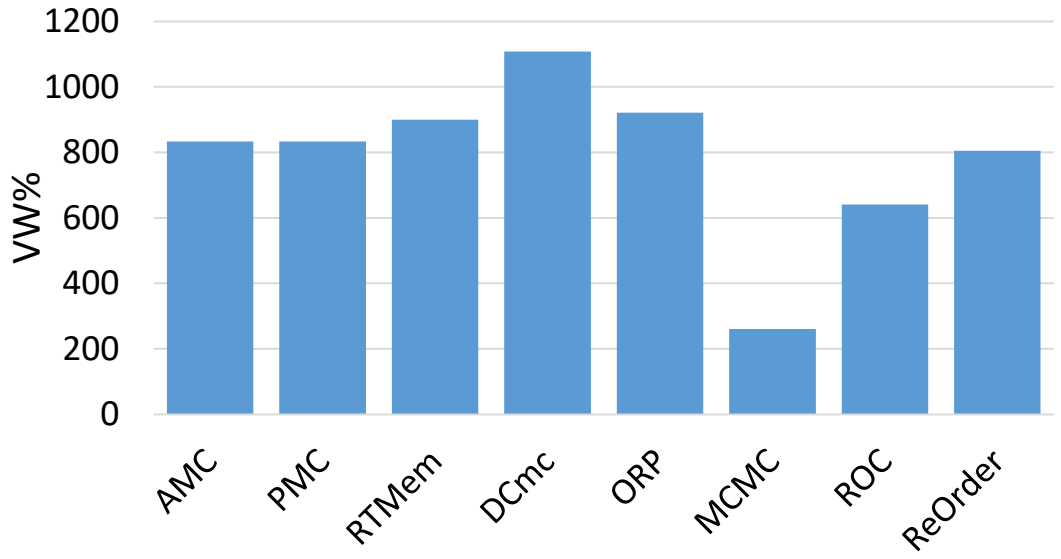

Assessing DDR Controllers for Predictability

- We calculate the VW for 8 of the state-of-the-art DDRx DRAM Controllers

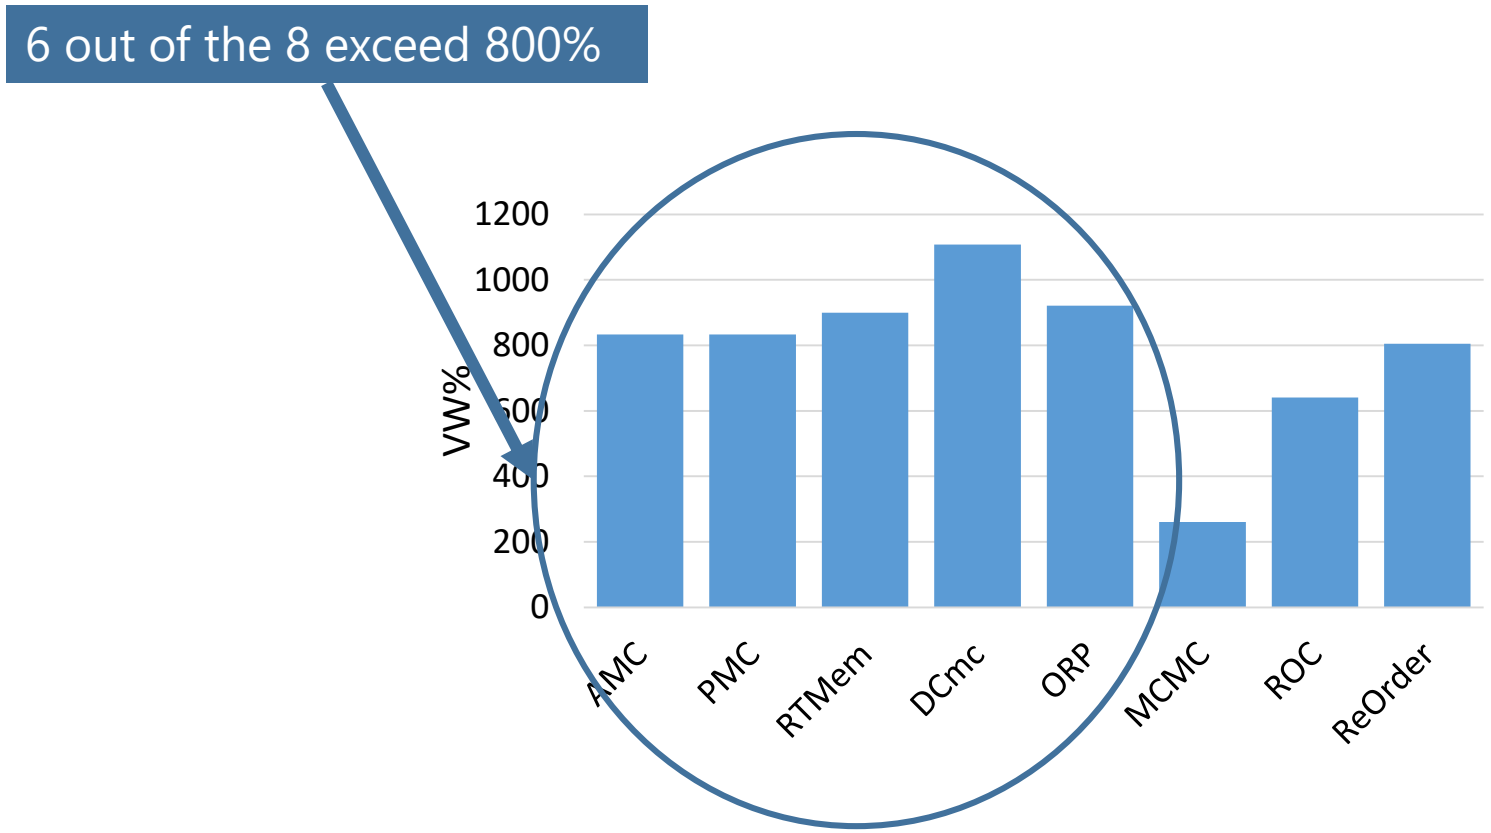

Assessing DDR Controllers for Predictability

PREDICTABILITY

- We calculate the VW for 8 of the state-of-the-art DDRx DRAM Controllers

6 out of the 8 exceed 800%

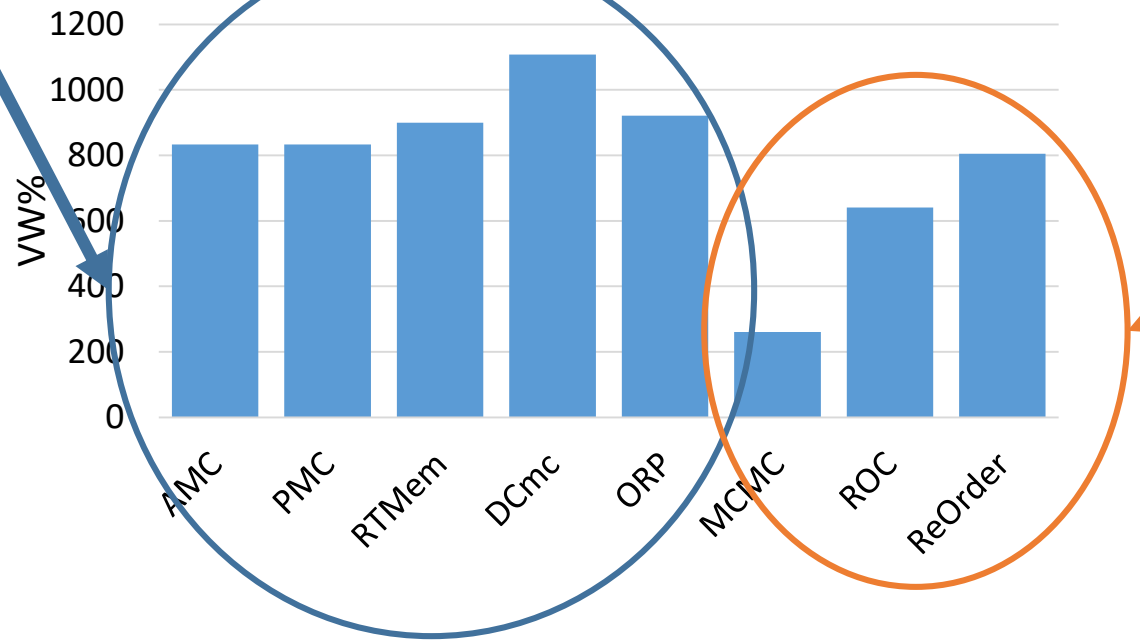

Achieve less variability at the expense of
1. complexity:
   - Bank partitioning
   - Rank switching
2. Conservatism:
   - e.g. using close-page (MCMC)

Assessing DDR Controllers for predictability

Even with the pessimism and complexity, 261% is still a significant variability for safety-critical systems

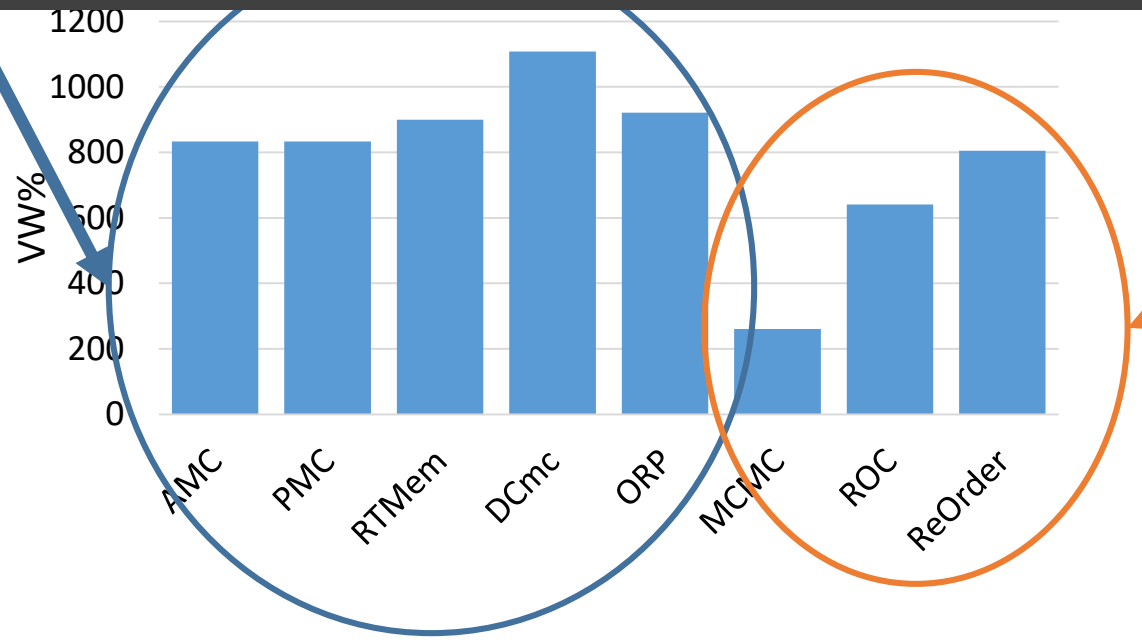

2. Conservatism:
   - e.g. using close-page (MCMC)

Assessing DDR Controllers for predictability

PREDICTABILITY

Even with the pessimism and complexity, 261% is still a significant variability for safety-critical systems

2. Conservatism:

Exploring other types of memories that address these limitations is unavoidable towards providing more predictable memory performance with less variability and tighter bounds

Assessing DDR Controllers for predictability

PREDICTABILITY

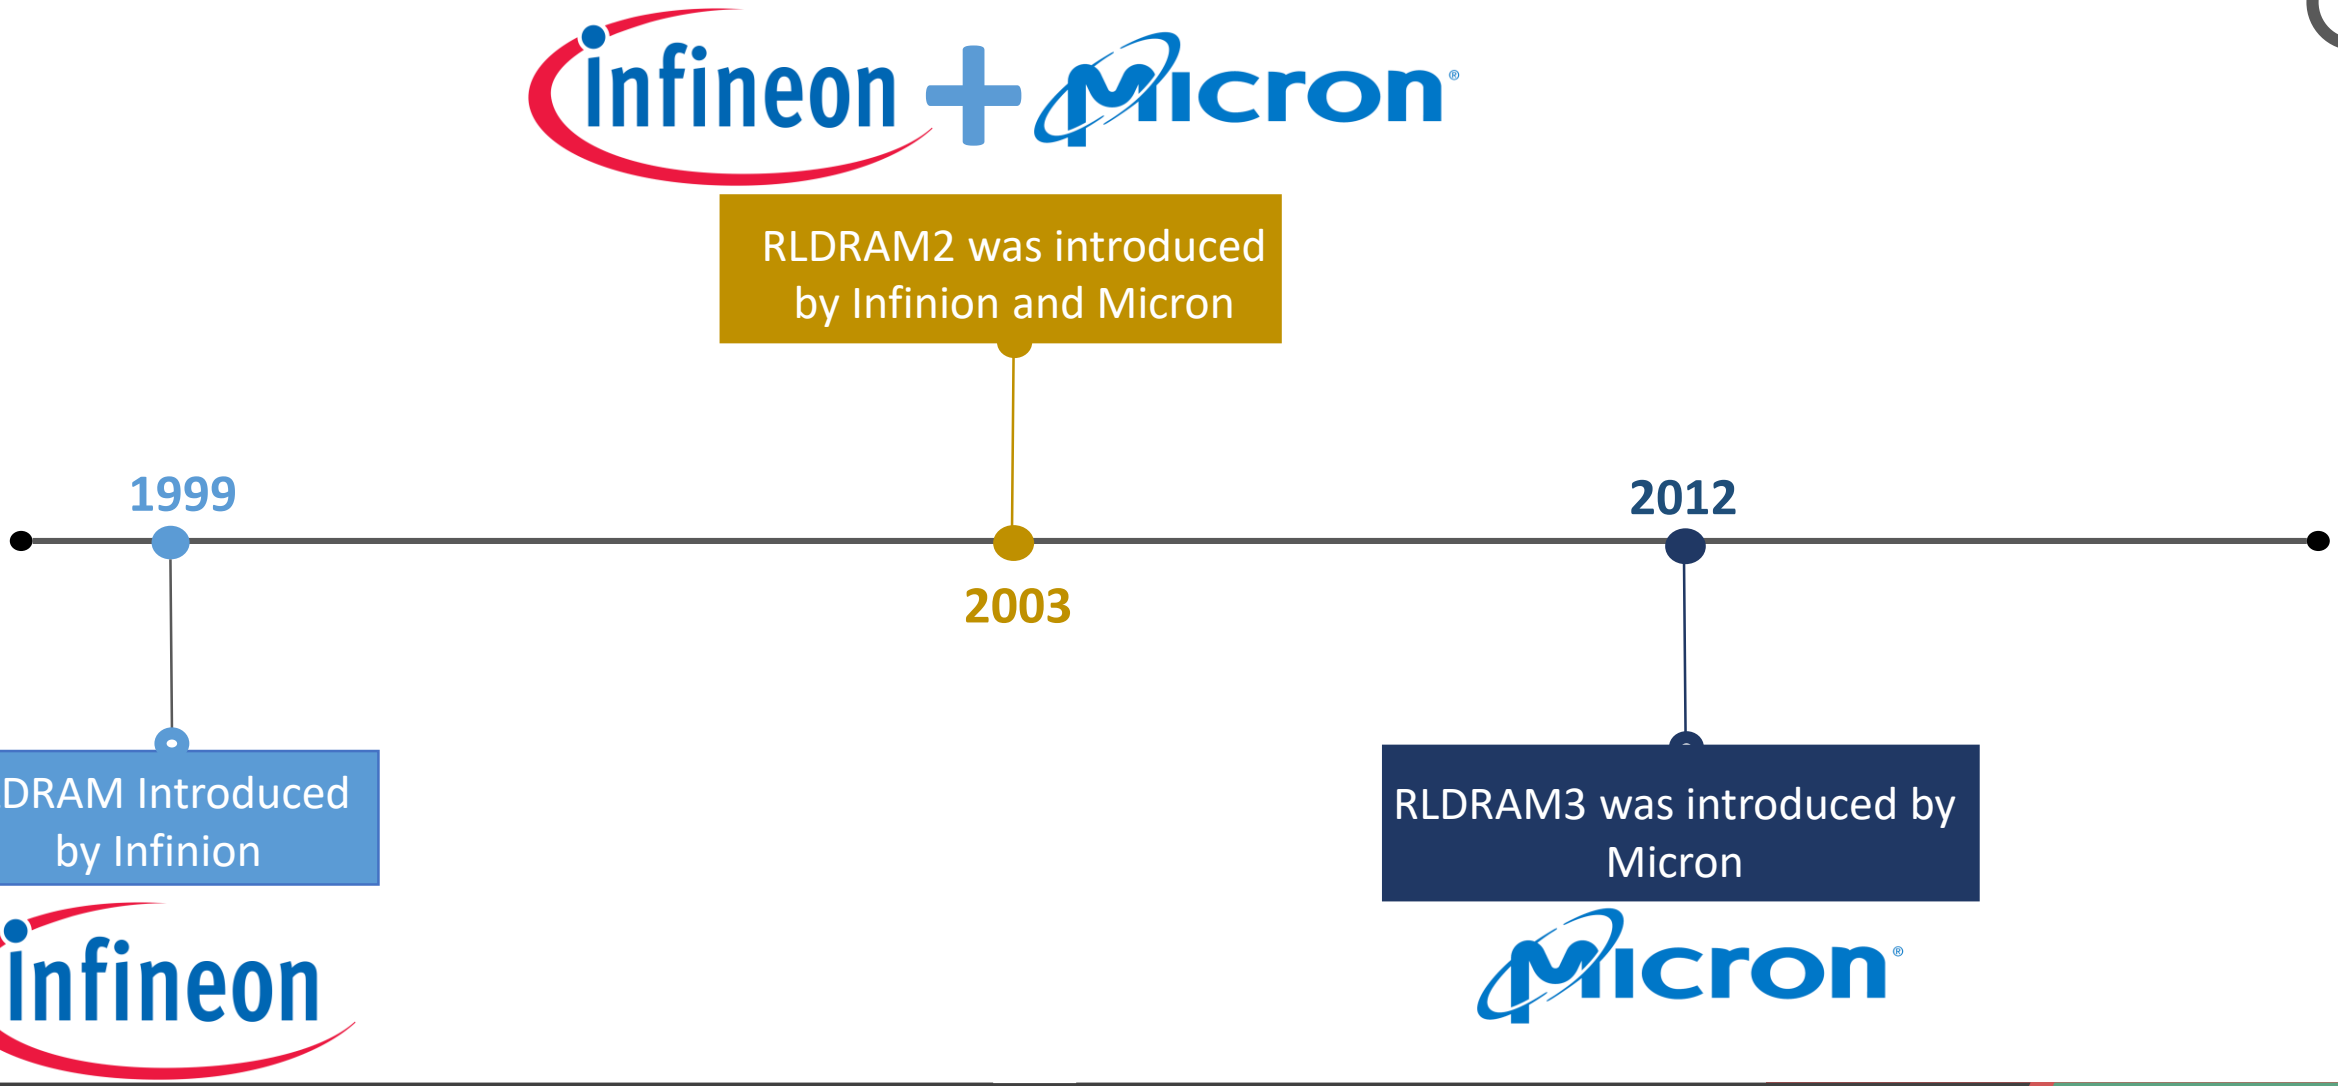

RLDRAM: An Alternative

Multiplexed Address Mode

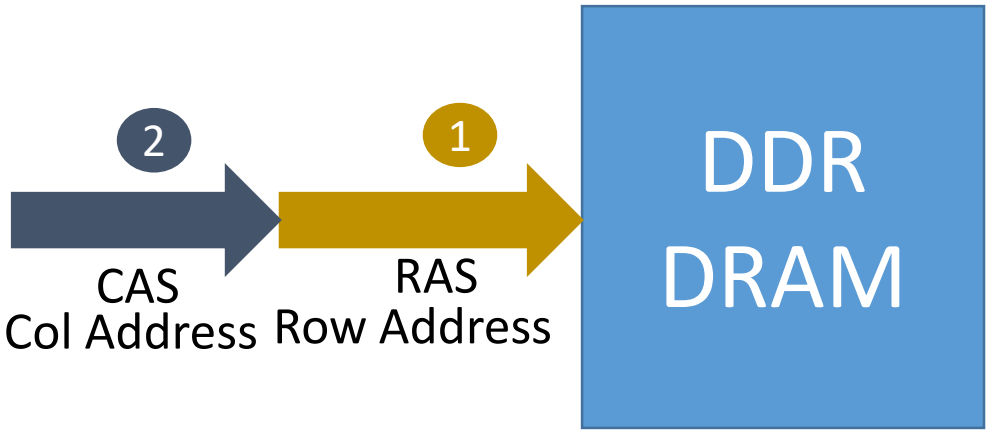

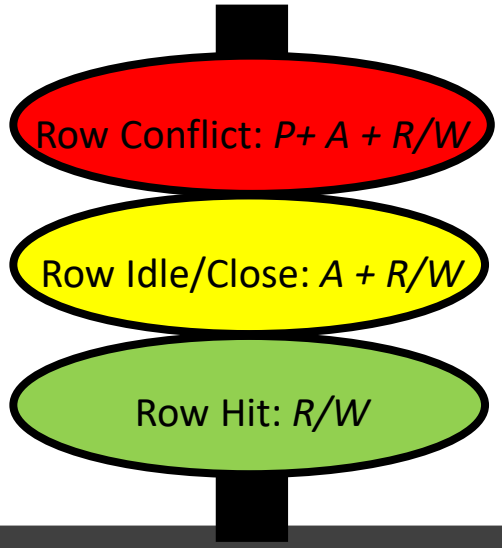

Non-Multiplexed Address Mode

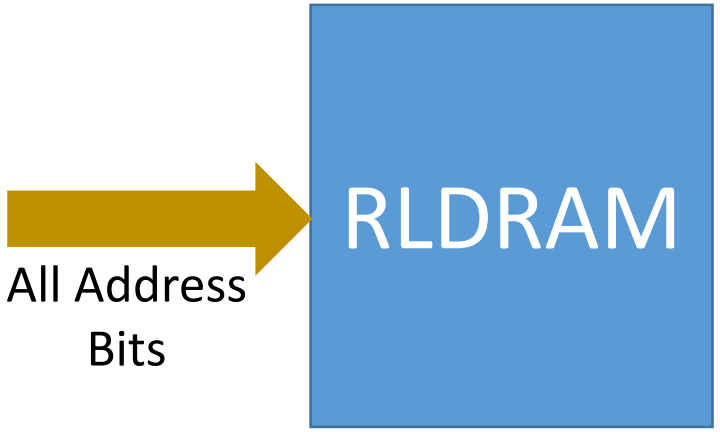

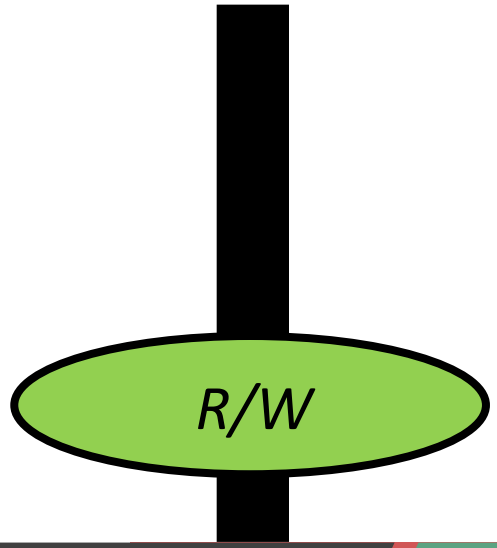

Why RLDRAM?

RLDRAM

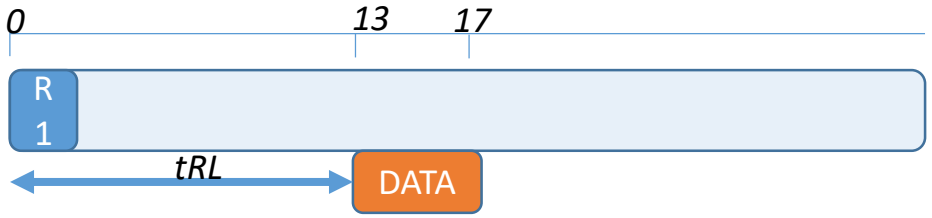

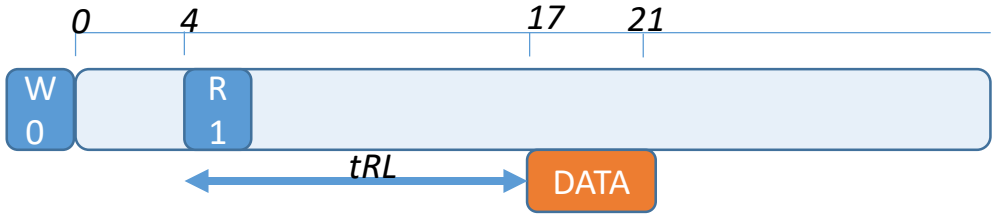

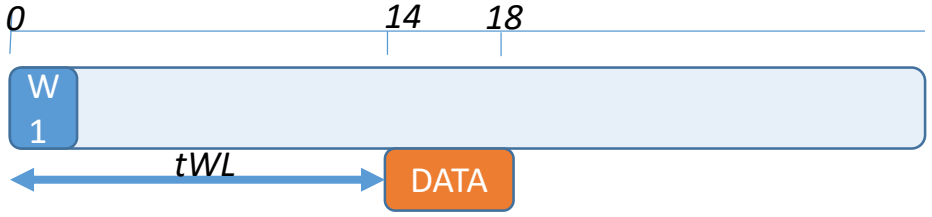

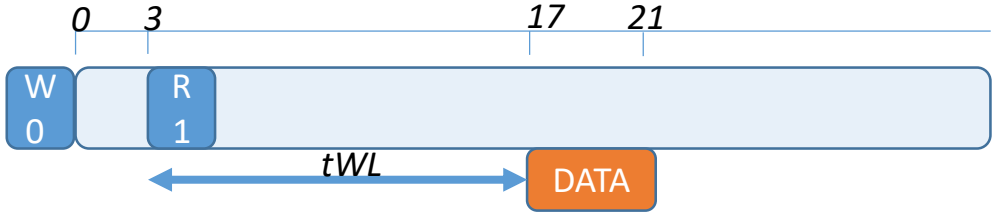

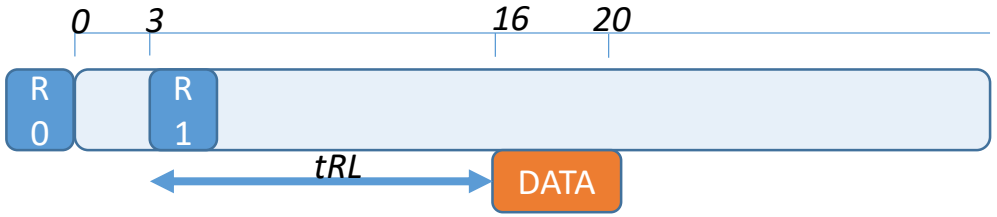

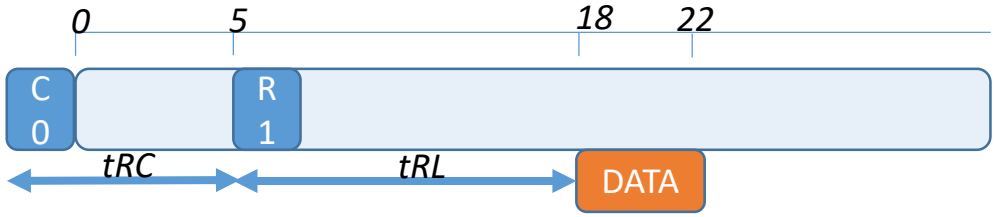

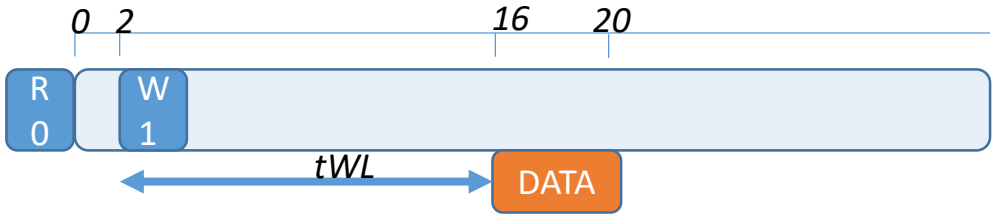

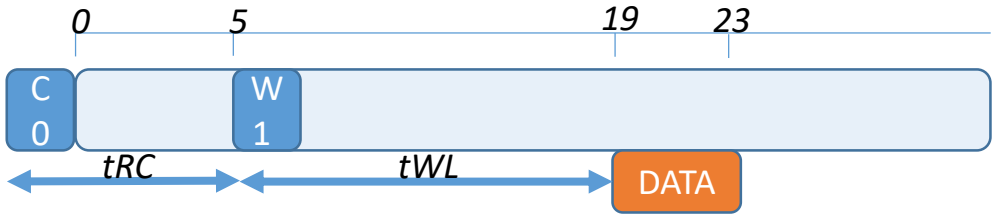

Assessing RLDRAM for Predictability

RLDRAM

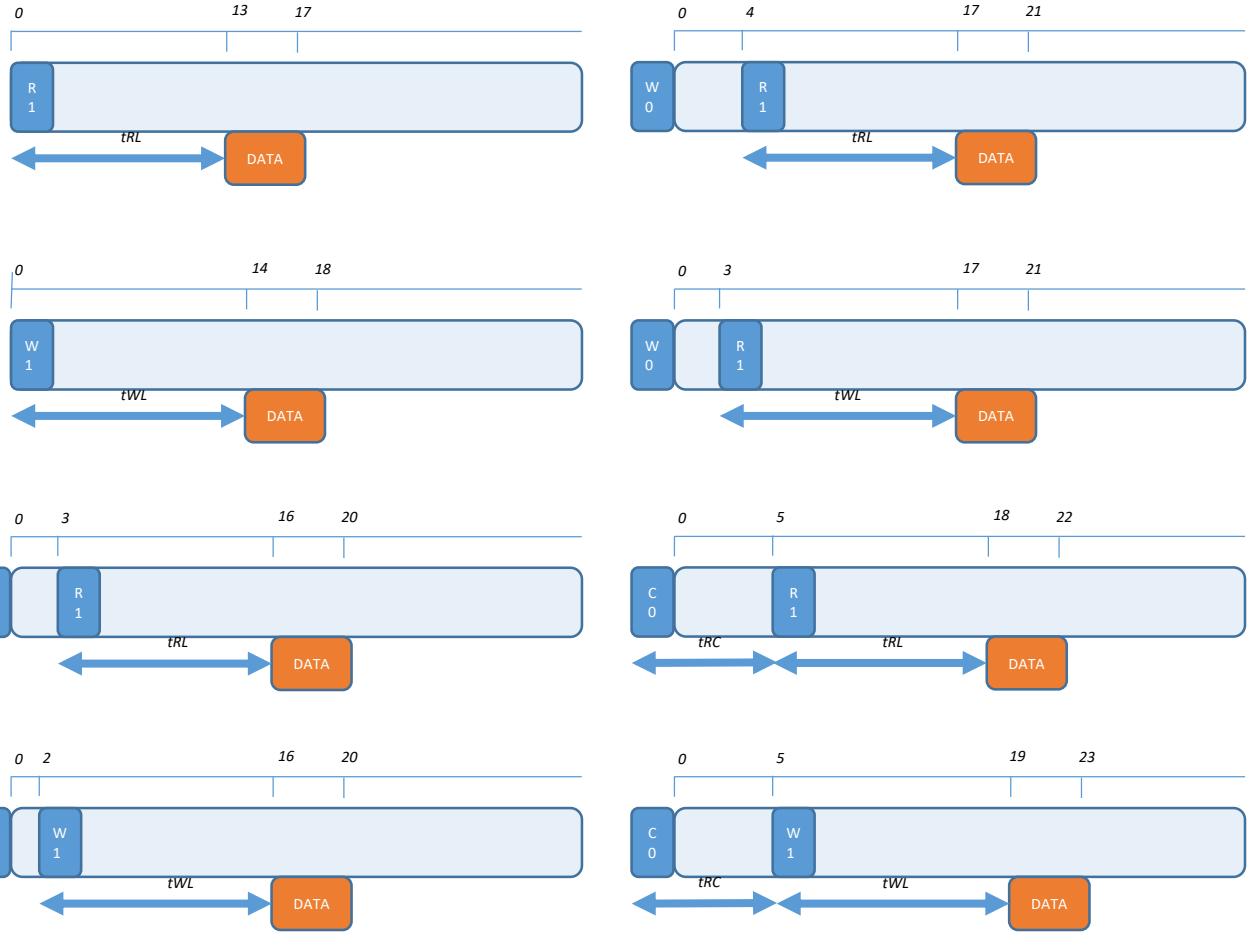

- A total of 8 cases
- Variability window is 46.2% (13.4x reduction)
- WCL=28.5ns (3.79x reduction)

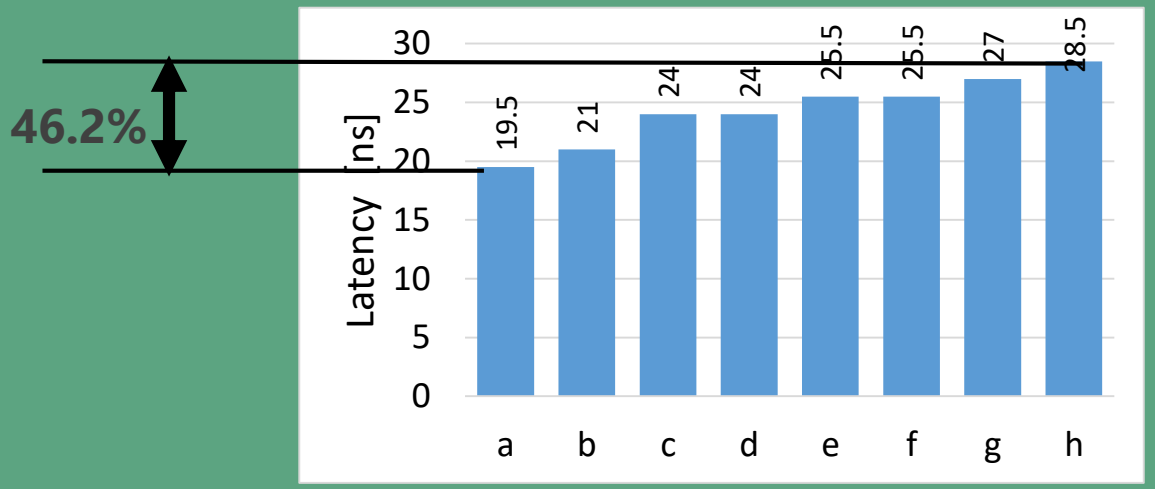

Assessing RLDRAM for Predictability

RLDRAM

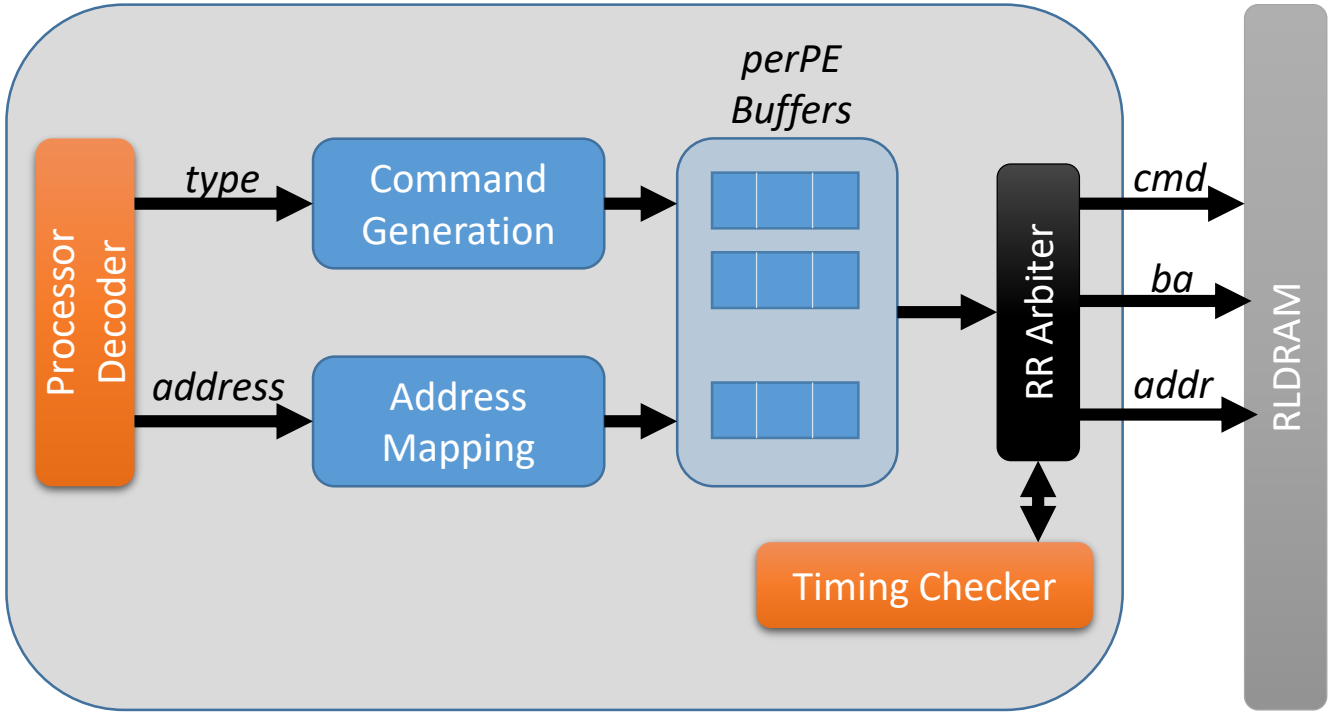

- RLDC to predictably manage accesses to RLDRAM
- Round Robin
- Support both bank sharing and bank partitioning
- Simple timing checker → good for analyzability, V&V, Certification

RLDC: A Predictable Controller for RLDRAM

RLDRAM

- Bank Sharing Scheme:

$$WCL^{share} = (N-1) \times tRC + tCL$$

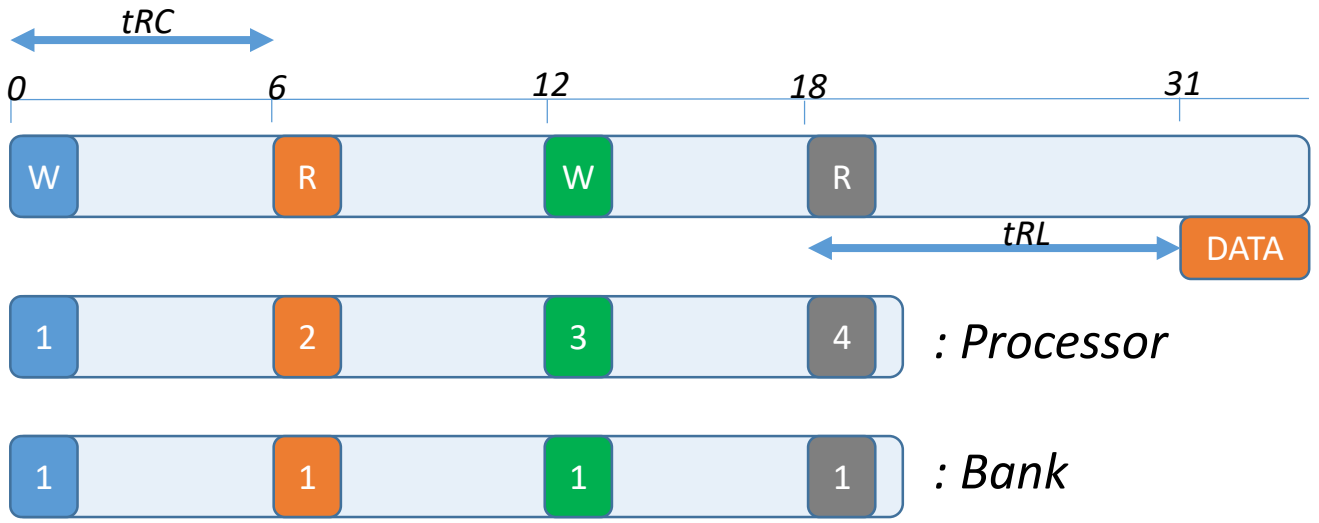

N is number of processing elements

Bounding Memory Latency

RLDRAM

- Bank Partitioning Scheme:

$$WCL^{part} = \left\lceil \frac{N-1}{2} \right\rceil \times \left(tWL - tRL + \frac{BL}{2} \right) + \left\lceil \frac{N-1}{2} \right\rceil \times \left(tRL - tWL + \frac{BL}{2} \right) + tCL$$

N is number of processing elements

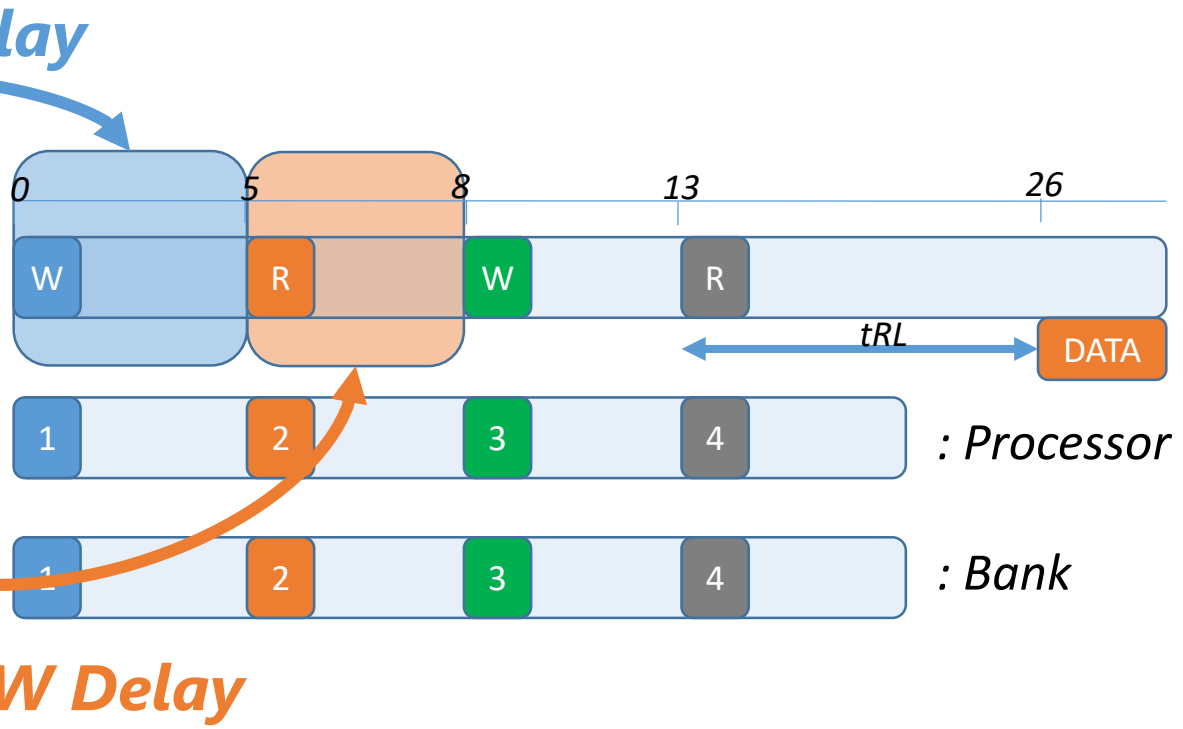

Bounding Memory Latency

RLDRAM

PEs	4 Processors in-order pipeline a private 16KB L1 a shared 1MB L2 cache
DRAM	Either RLDRAM or DDRx
RLDRAM	RLDRAM3-1600 RLDC manages accesses to RLDRAM
DDRx	DDR3-1600 AMC, PMC, RTMem, DCmc, ORP, MCMC, ROC, or ReOrder manages access to DDR3
Bank Management	We experiment with both bank partitioning and bank sharing among PEs for RLDC
Benchmarks	EEMBC Automotive

Evaluation Setup

RESULTS

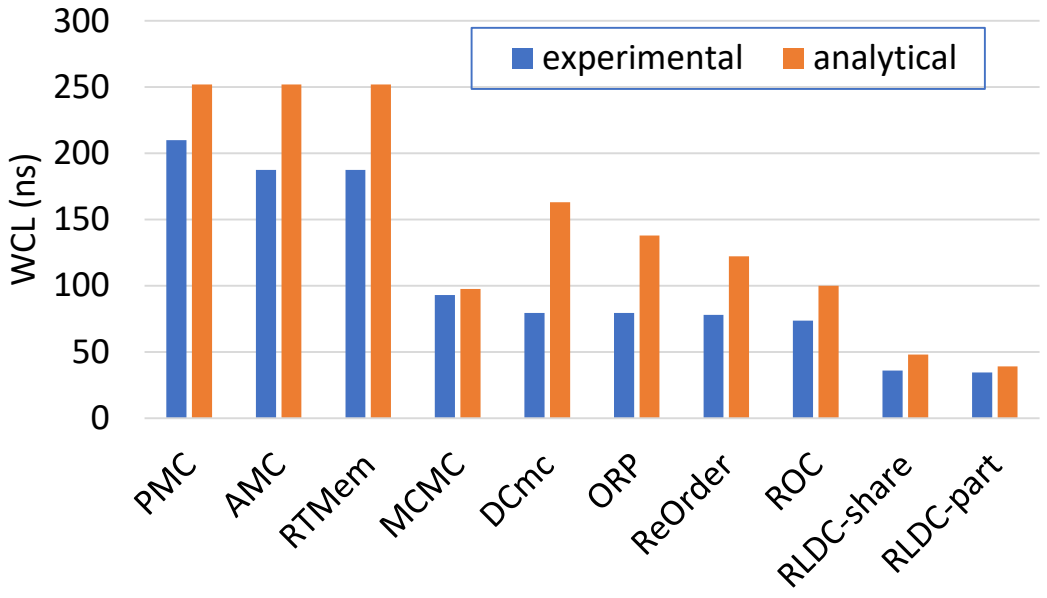

1. DDRx MC has **2.5x** to **6.46x** worse analytical WCL than RLDC
   → *Very similar numbers for exp.*
2. Relatively low WCL of MCMC, ROC, ReOrder is due to 4 ranks!
3. For RLDC: bank partitioning provide tighter WCL than sharing
   → *at the expense of flexibility*
4. Gap between exp. vs analytical WCL is much higher for DDR
   → *again due to inherent variability*

Worst-Case Latency

RESULTS

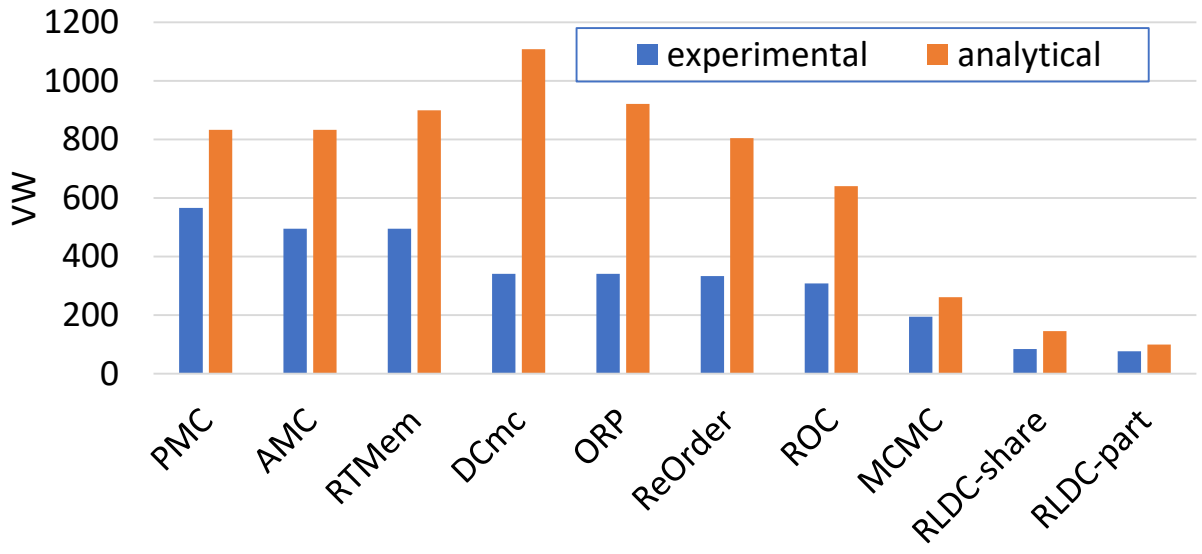

1. Already discussed analytical VW
2. *Exp. VW for DDRx MCS:*
   - *>400% for 4 MCs,*
   - *300%-400% for 3 MCs and*
   - *~200% for 1 MC.*
3. *for RLDC:*
   - *76.9% for partitioned banks*
   - *84.6% for shared banks*

Variability Window

RESULTS

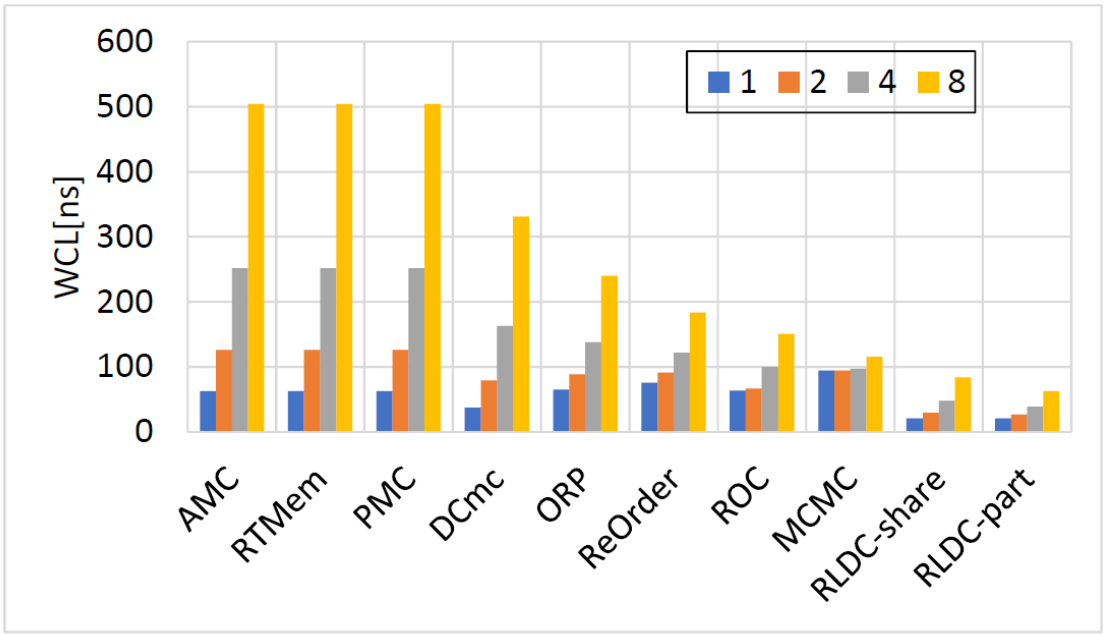

1. The WCL latency gap between RLDC and majority of DDRx MCs increases drastically

Scalability: # Processors

RESULTS

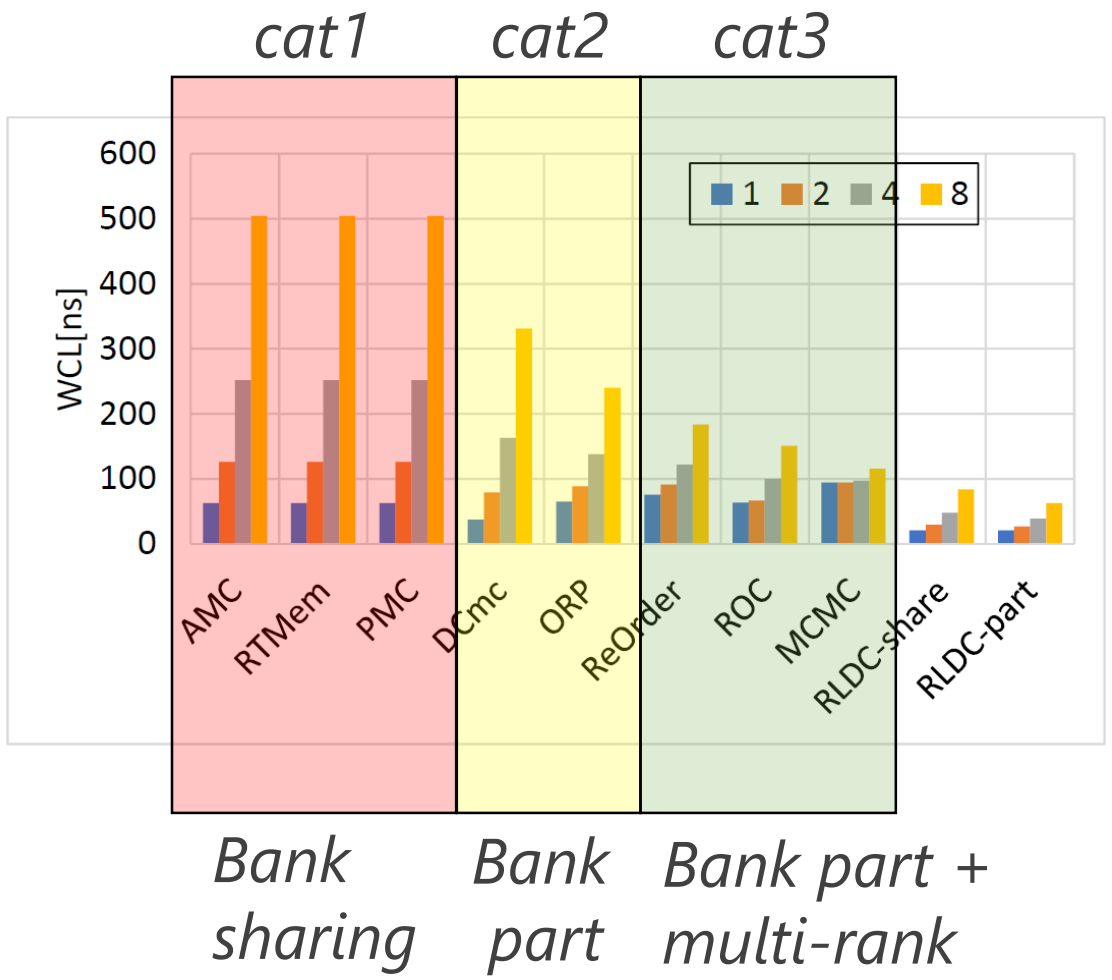

1. The WCL latency gap between RLDC and majority of DDRx MCs increases drastically
2. DDRx MCs can be categorized into three categories

Scalability: # Processors

RESULTS

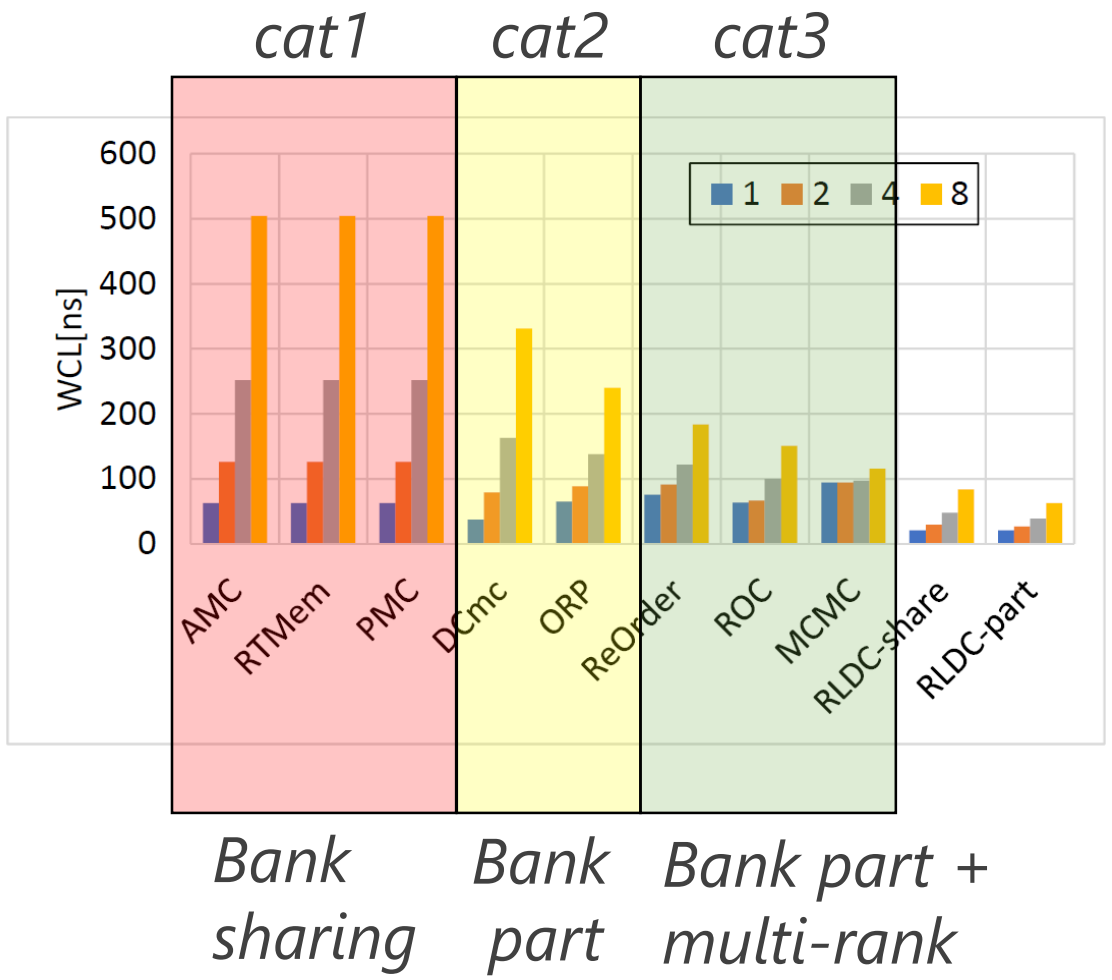

1. The WCL latency gap between RLDC and majority of DDRx MCs increases drastically
2. DDRx MCs can be categorized into three categories
3. RLDC's WCL is less than all categories for all #PEs
   → for both part and sharing
   → without complex arbitration/ reorderings (better analyzability and composability)

Scalability: # Processors

RESULTS

1. How mature is RLDRAM?

- Has been there since 1999
- Long-term Supported by Micron
- The main off-chip memory in networking and other low-latency needs

2. Can I buy RLDRAM off-the-shelf?

- Definitely
- They are sold as discrete components (which is the norm for embedded systems anyway)
- There are also specialized DDR-compatible sockets for RLDRAM

3. Are there any platforms/boards/test-beds?

XILINX®

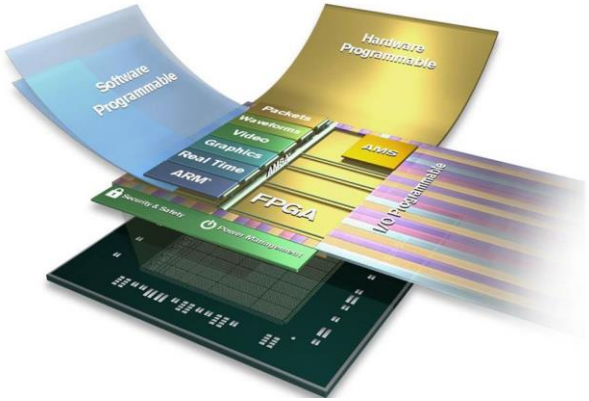

Zynq UltraScale/+ MPSoC

intel®

Arria10 SoC

ARTESYN™
EMBEDDED TECHNOLOGIES

ATCA-9305-NSP
Processing blade

NORTHWEST LOGIC

LATTICE SEMICONDUCTOR™

Discussion

RESULTS

4. If it is that good, why it did not take off then?

- Take off = replaces DDR? → It is not meant to be!
- It is a specialized type of memory for low-latency guarantees
- Commodity general-purpose market is looking for high BW, capacity, cost
- We should ask ourselves, what are we looking for?
- An analogy: FPGAs have been there for long-time, why are they taking off now? → they satisfy a "new" need in the AI market
  - Will they replace CPUs then?
  - It is not an either-or decision

5. Does that mean we no longer need DRAMs in real-time systems?

- No
- Use-case dependent
- A heterogenous memory system?
- Mixed Criticality with different requirements?

6. As a scheduling researcher, why should I care?

- Taking the shared resources interference into account is unavoidable to provide more accurate numbers
- DDR DRAM is very complex to account for its details (rd vs wr, conf vs hit, ..etc) → explodes the analysis
- RLDRAM alleviates this complexity..which can make the analysis more feasible

7. Limitations/trade-off

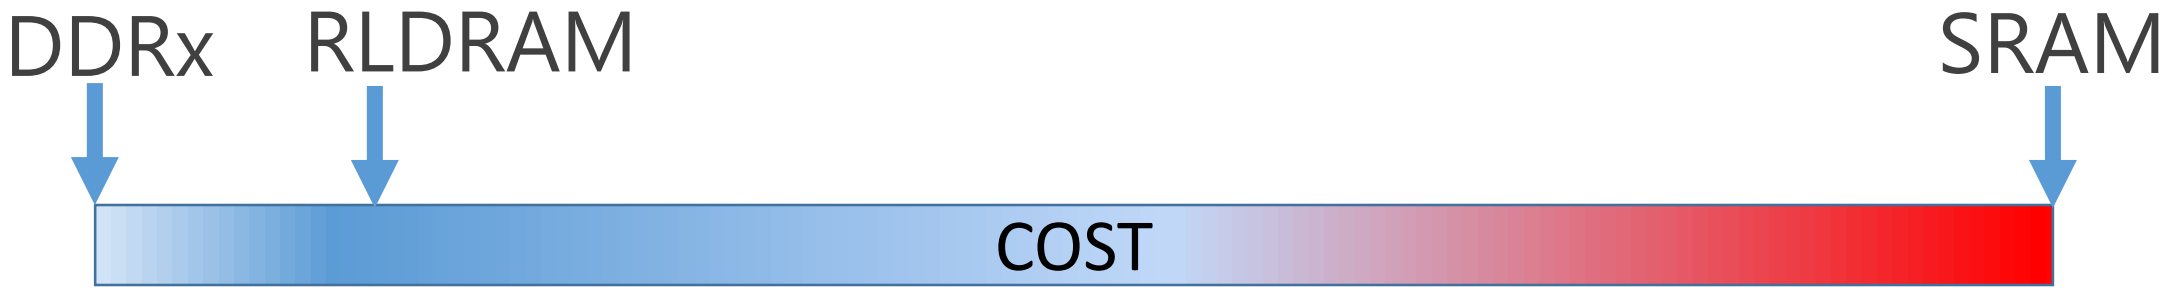

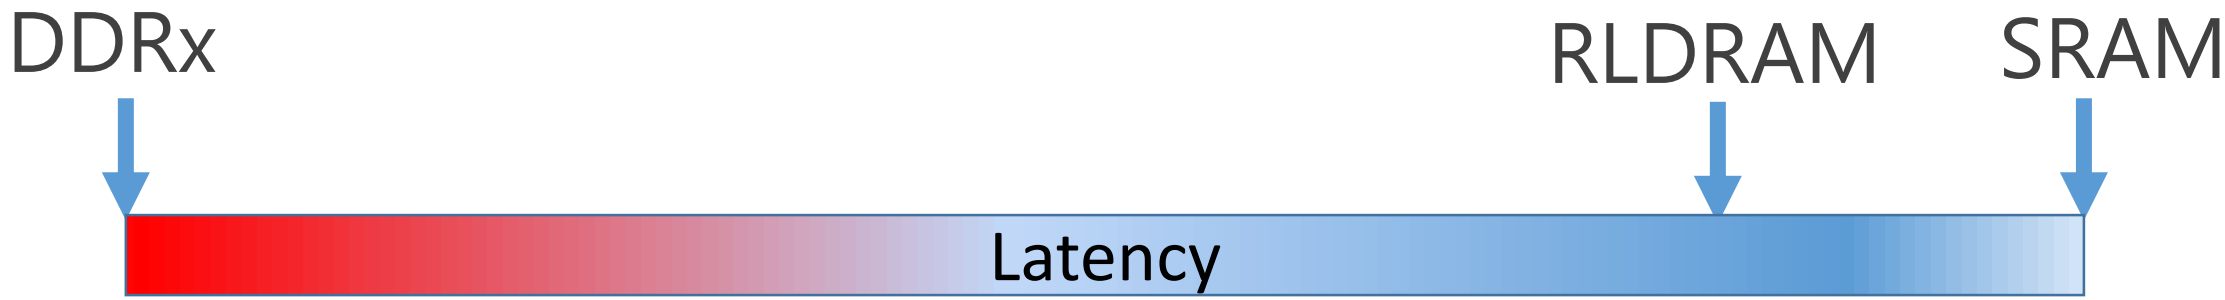

Discussion

RESULTS

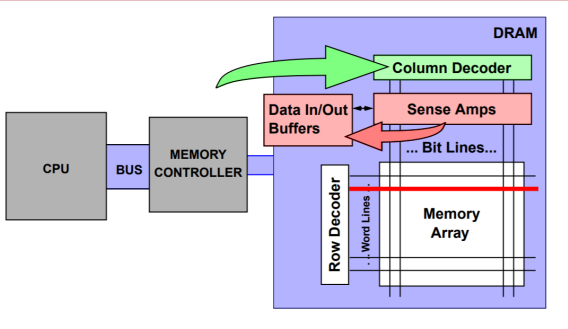

11× less timing variability
6.4× less WCL

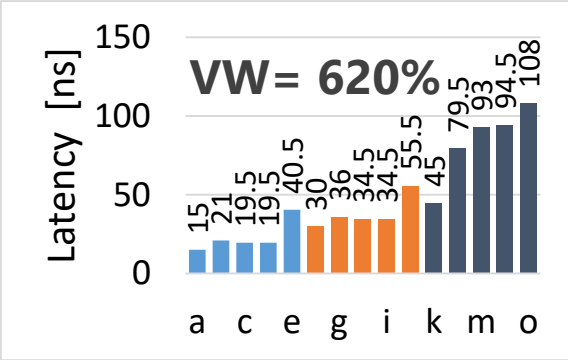

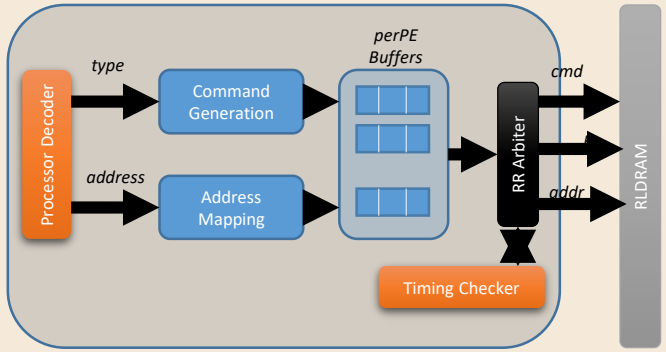

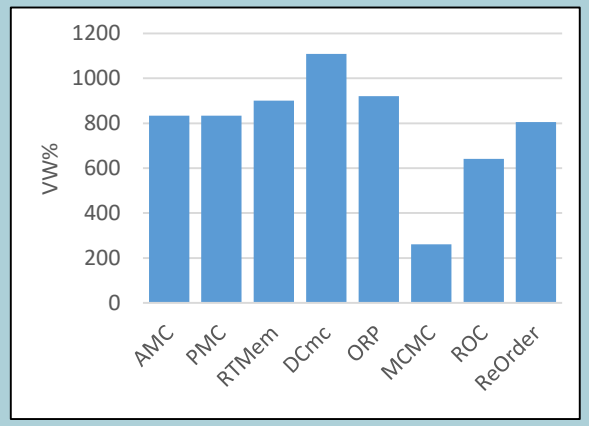

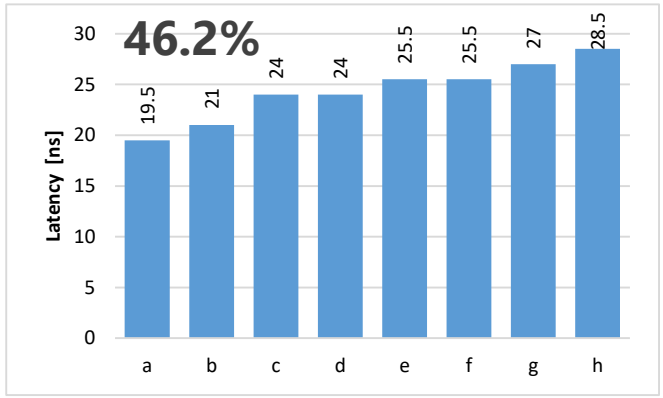

Main Lessons:

1. DDR DRAM is not designed for predictability, RLDRAM is.
2. Looking for solns that address our needs instead of starting from the mainstream soln?
3. Not either-or:
   A heterogenous memory to address conflicting needs of MCS

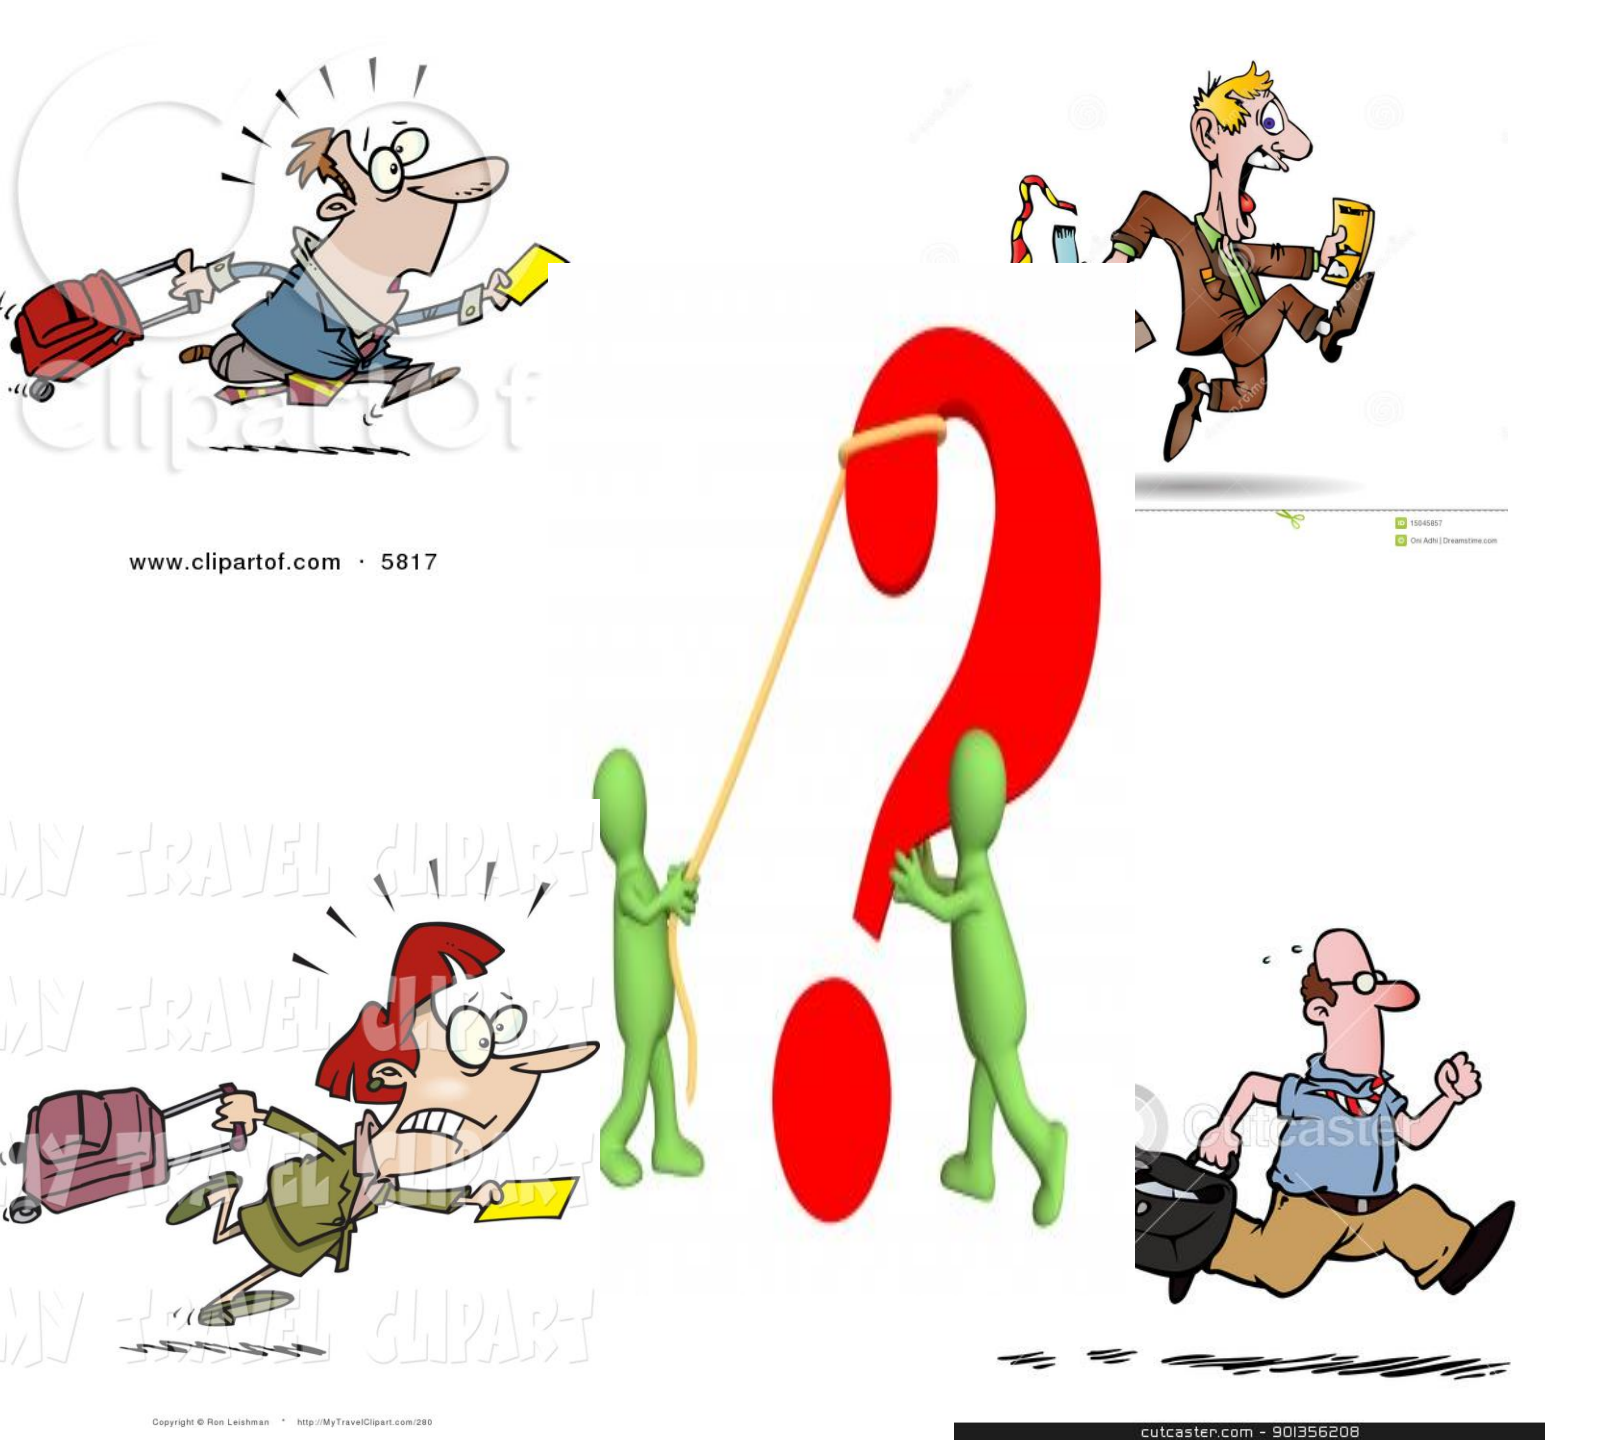

Main Lessons:

1. DDR DRAM is not designed for predictability, RLDRAM is.
2. Looking for solns that address our needs instead of starting from the mainstream soln?
3. Not either-or:
   A heterogenous memory to address conflicting needs of MCS