An MLP-Aware Leakage-Free Memory Controller

ANDREW VUONG, A. SHAFIEE, M. TAASSORI,
R.BALASUBRAMONIAN

Timing Channels in a Shared Memory Controller

- Spectre attack relied on cache timing channels
  - Information can be leaked through a variety of channels including the memory controller
- Discover private keys, crack passwords, leak secrets, etc.
- Side channel and covert channel attacks

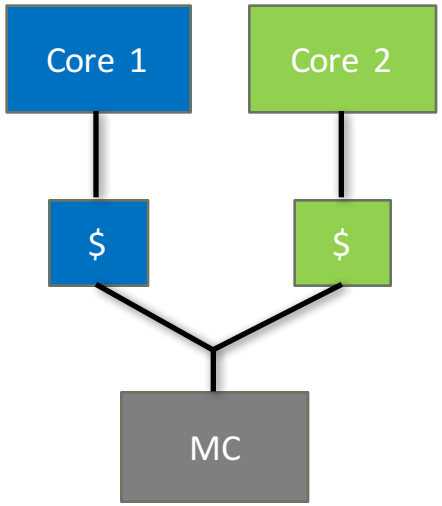

UTAH ARCH

Covert Channel Attacks

- Sender changes amount of memory requests
  - Low # of requests – 0 bit
  - High # of requests – 1 bit
- Receiver sends constant amount of requests.
  - Uses throughput to decipher information sent.

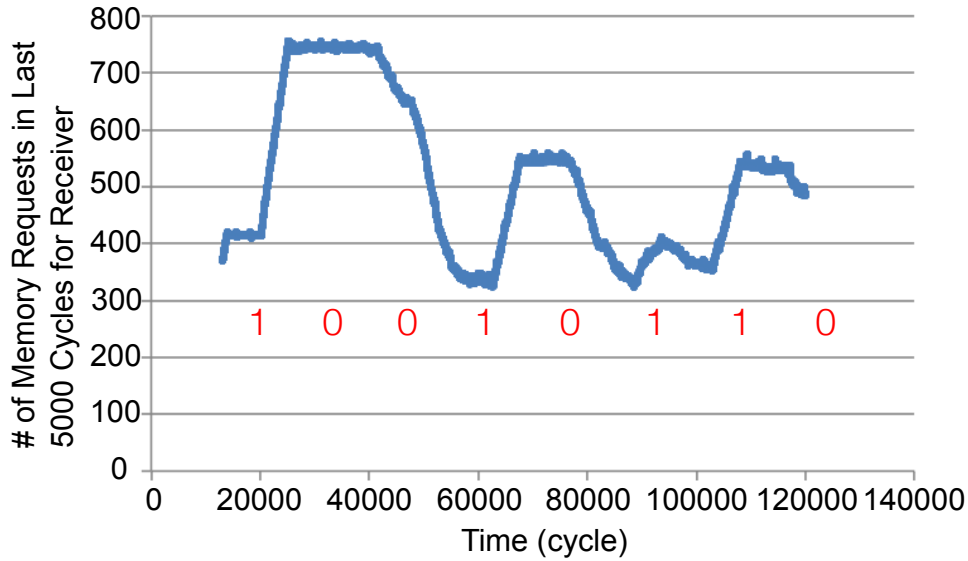

Yao Wang, Andrew Ferraiuolo, and Edward Suh, "Timing Channel Protection for a Shared Memory Controller", HPCA 2014

UTAH ARCH

Covert Channel Attack Example

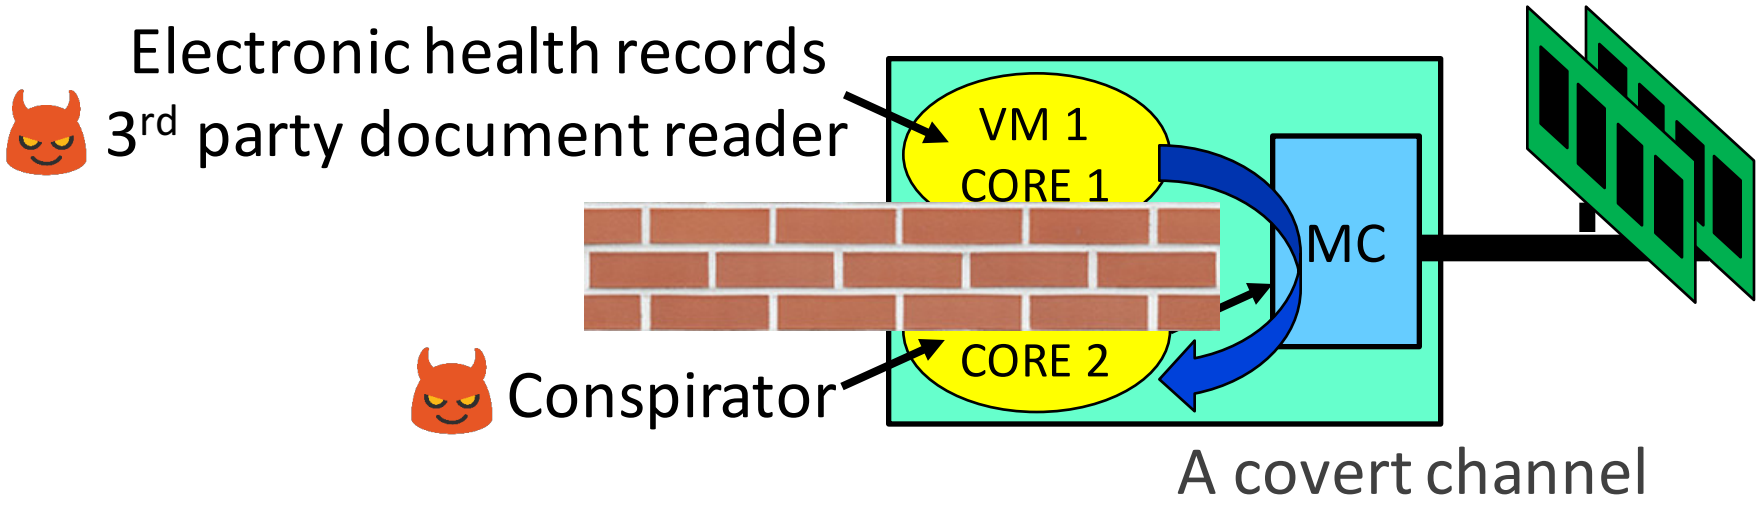

Timing Restrictions

- DRAM cycles between consecutive accesses
  - Different ranks: **7 cycles**
  - Same rank, different banks: **15 cycles**
  - Same rank, same bank, different rows: **43 cycles**

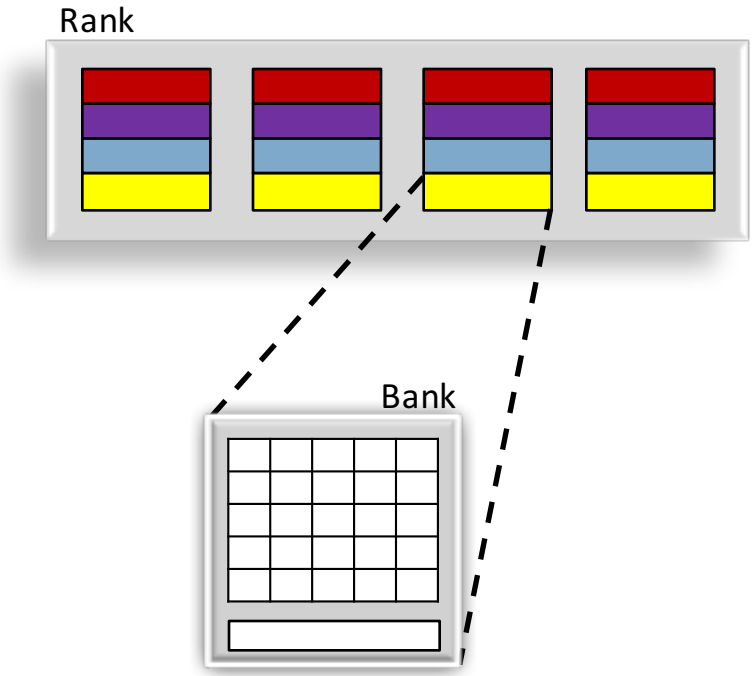

Temporal Partitioning (TP) Wang et al., HPCA'14

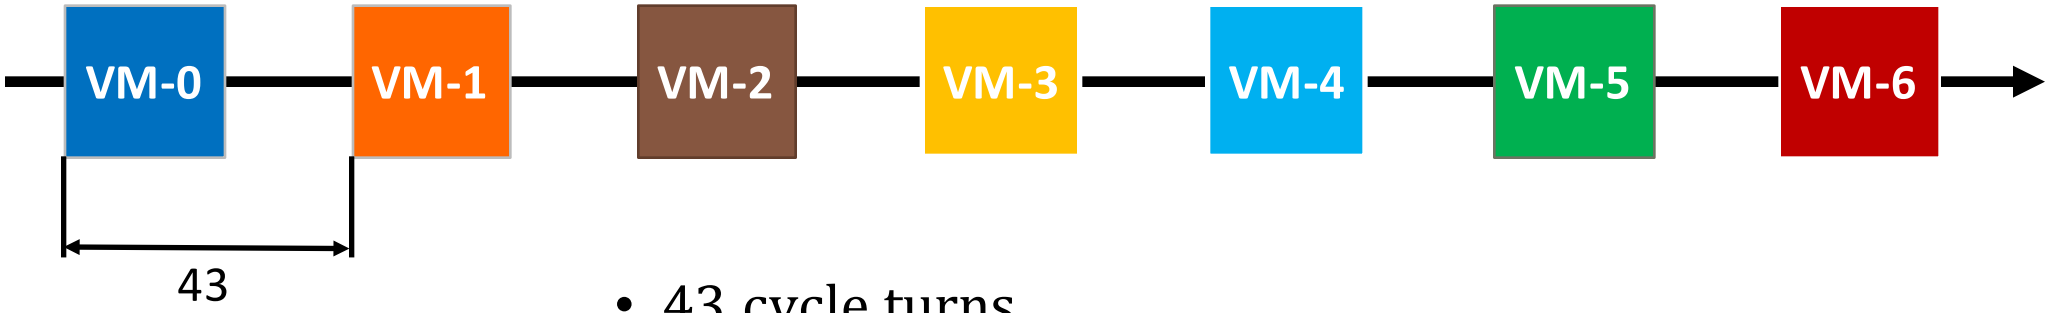

- 43 cycle turns
- Round-robin scheduling
- Large overhead from dead time

UTAH ARCH

Fixed Service (FS) Shafiee et al., MICRO'15

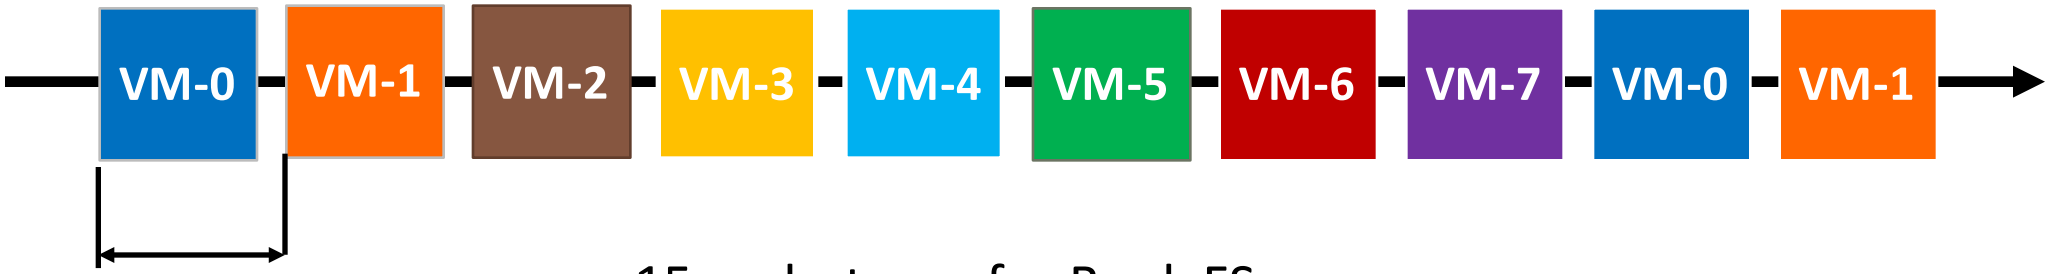

Bank FS: 15
Rank FS: 7

- 15 cycle turns for Bank FS
- 7 cycle turns for Rank FS
- **Spatial partitioning**
- Low scalability

UTAH ARCH

Bank Triple Alternation (BTA) Shafiee et al., MICRO'15

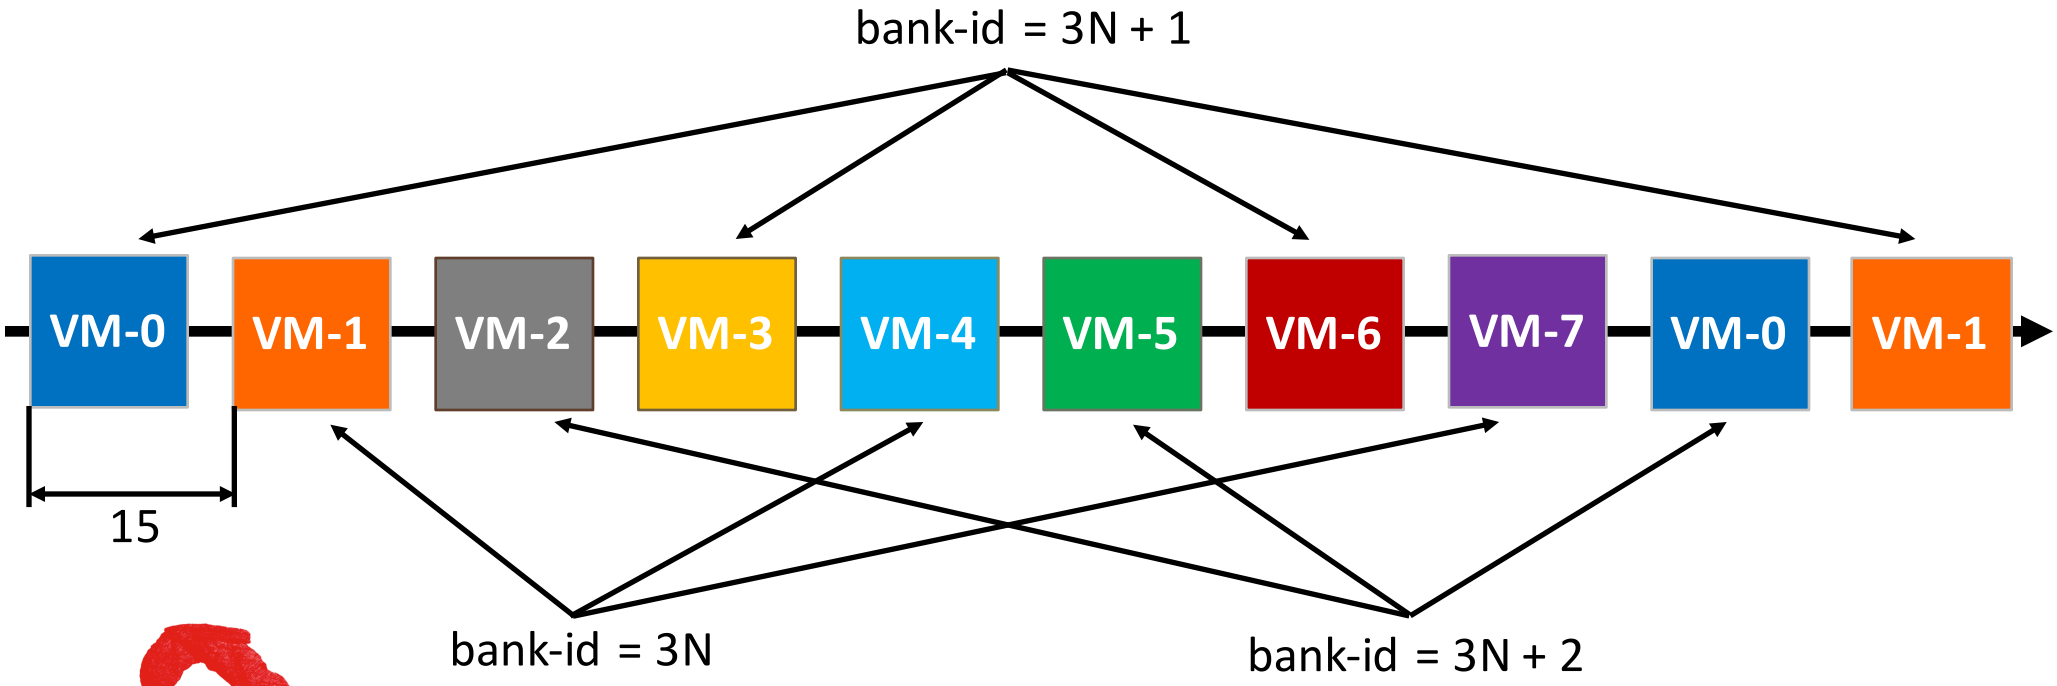

UTAH ARCH

SecMC-NI (3R/2B) Wang et al., HPCA'17

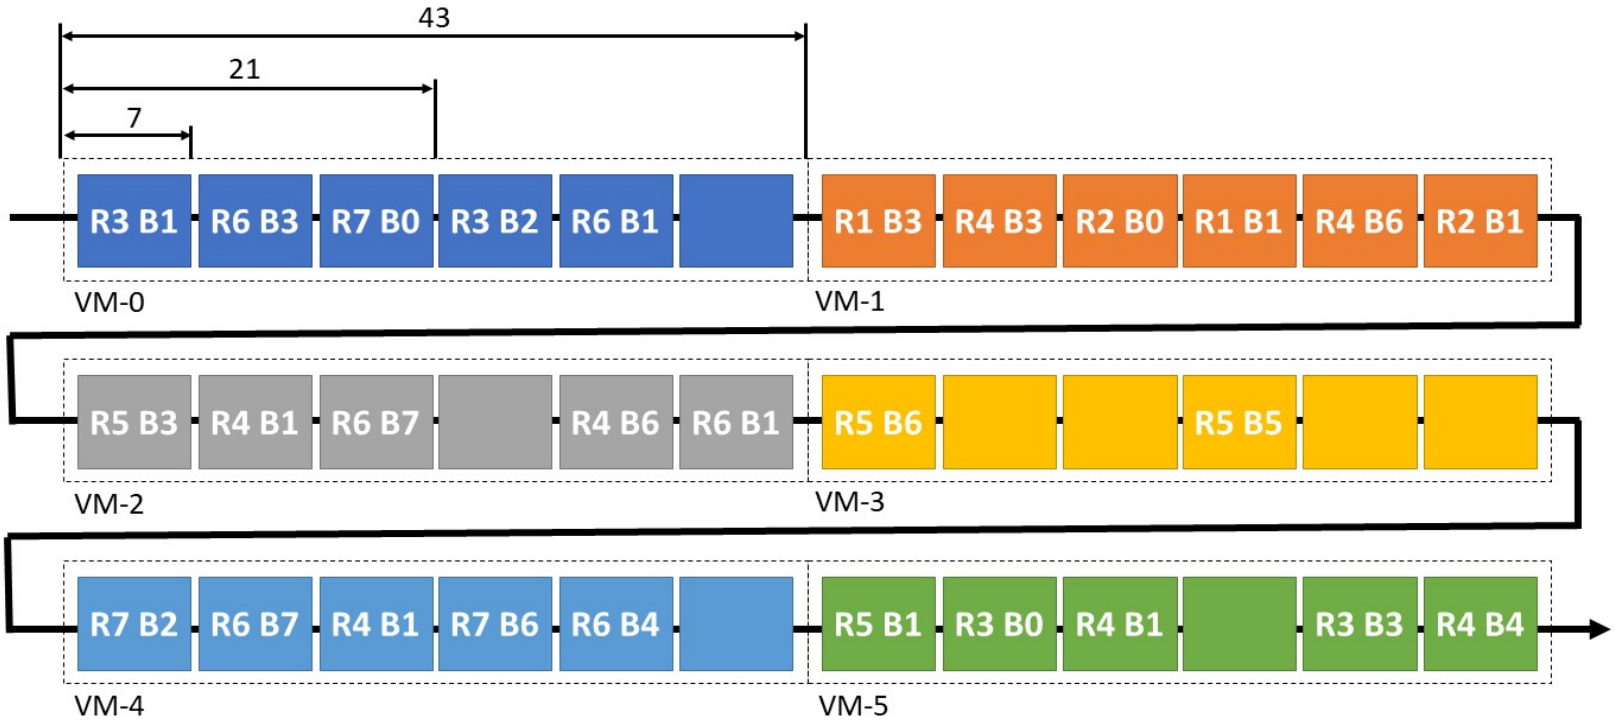

UTAH ARCH

Drawbacks of SecMC-NI (3R/2B)

- Underutilization of turns
- 82% of slots go unused
- Long-wait times between turns

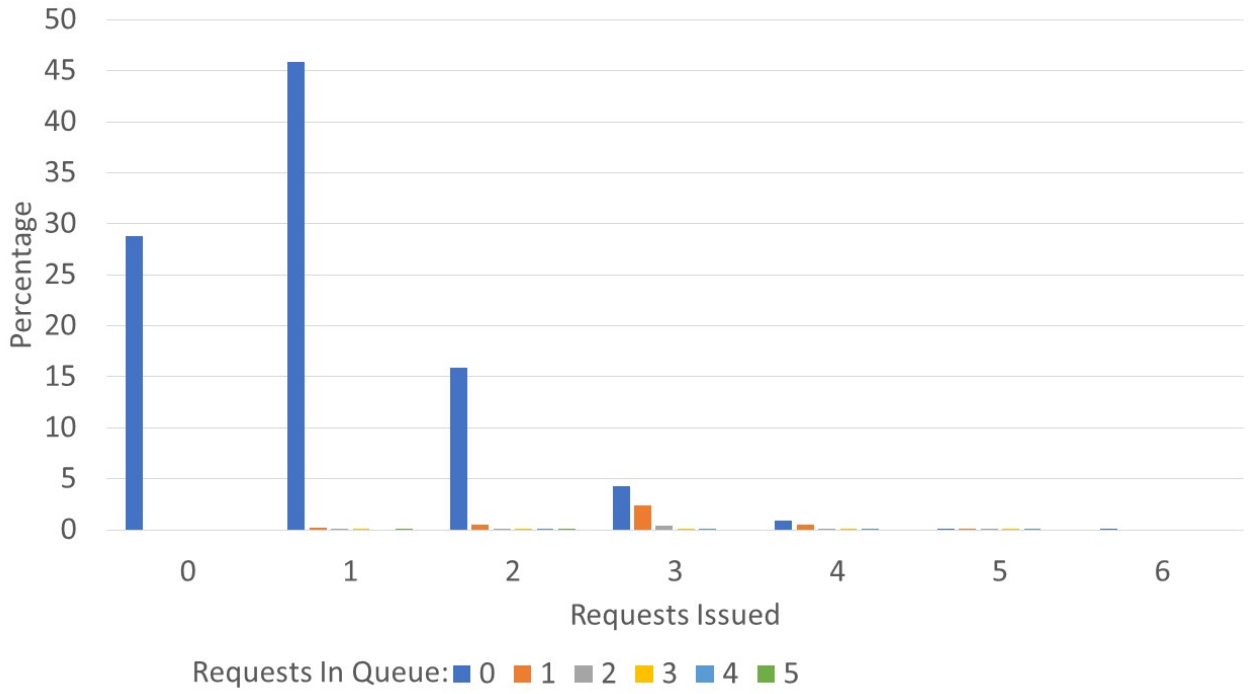

UTAH ARCH

Ranked Triple Alternation (RTA)

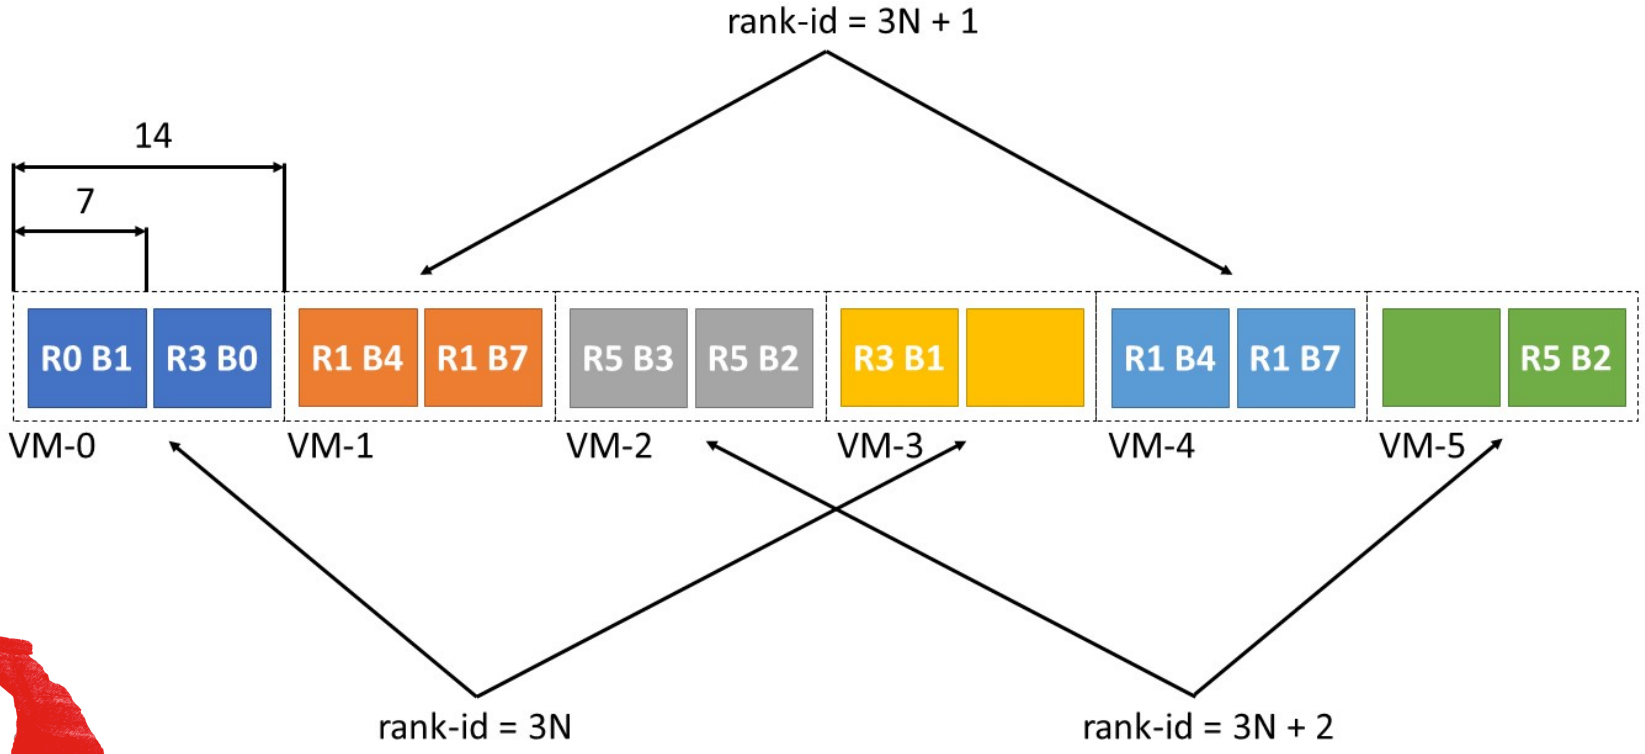

UTAH ARCH

Ranked Subsetting

- Divides memory ranks into subsets
- Increased data transfer time
- Reduced row buffer size
- Increased parallelism
- Improved energy efficiency

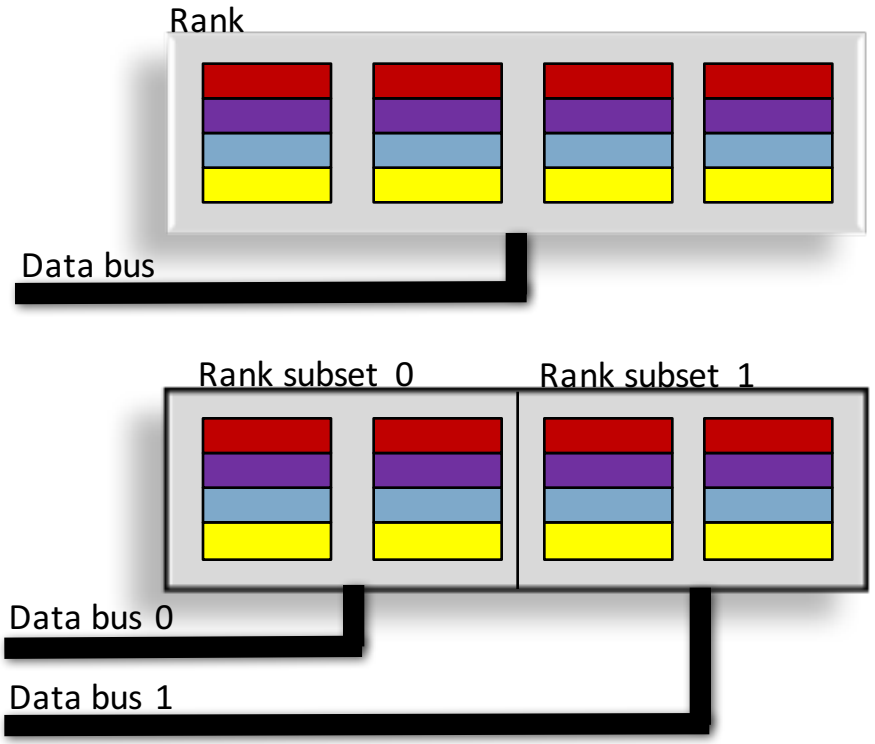

UTAH ARCH

Methodology

- Simics
  - 8 4-way superscalar cores
  - L1I (32KB)/L1D (32KB)/L2 (0.5MB) per core
- USIMM
  - 1 channel, 8 ranks, 8 banks
- SPEC 2006 benchmark
- FR-FCFS baseline

Processor	
ISA	UltraSPARC III ISA
CMP size and Core Freq.	8-core, 3.2 GHz
ROB size per core	128 entry
Fetch, Dispatch, Execute, and Retire	Maximum 4 per cycle
Cache Hierarchy	
L1 I-cache	32KB/2-way, 1-cycle
L1 D-cache	32KB/2-way, 1-cycle
L2 Cache	4MB/8-way, shared,10-cyc
DRAM Parameters	
DRAM Frequency	1600 Mbps
Channels, ranks, banks	1 ch, 8 ranks/ch, 8 banks/rank
DRAM chips	4 Gb capacity
DRAM Timing Parameters (DRAM cycles) t_{RC} = 39, t_{RCD} = 11,t_{RAS} = 28, t_{FAW} = 24 t_{WR} = 12,t_{RP} = 11,t_{RTRS} = 2, t_{CAS} = 11 t_{RTP} = 6, t_{BURST} = 4,t_{CCD} = 4, t_{WTR} = 6 t_{RRD} = 5, t_{REFI} = 7.8μs, t_{RFC} = 260ns	

UTAH ARCH

Performance of RTA

- RA – modulo 7 alternation
- RA outperforms 3R/2B by 5%
- RTA outperforms 3R/2B by 11%

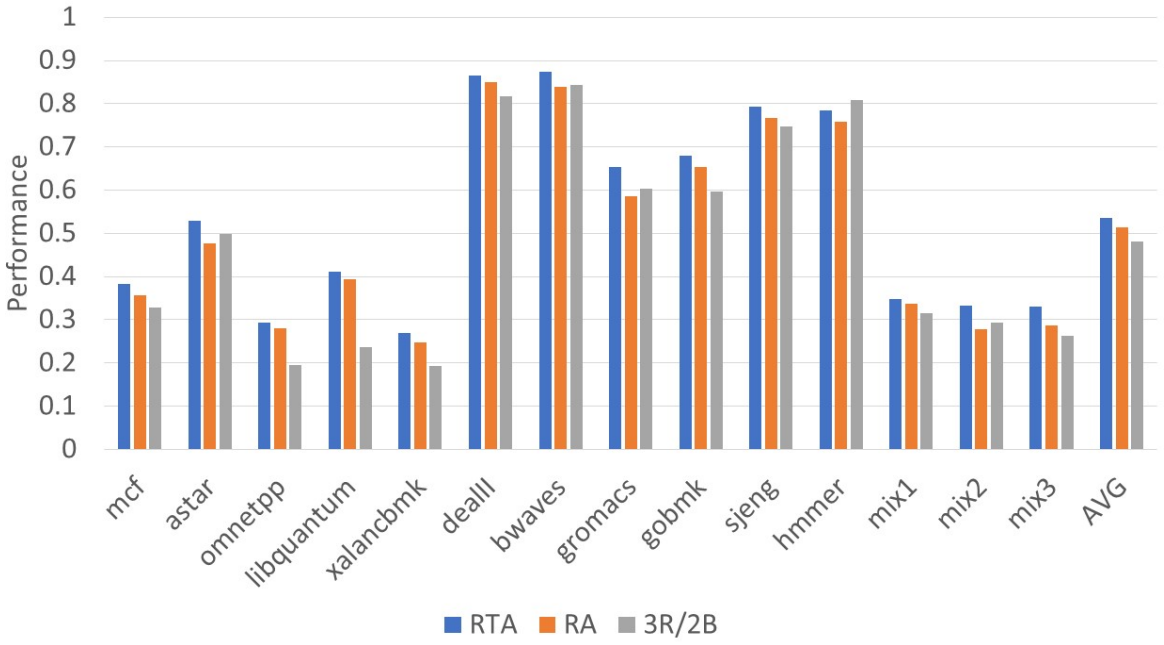

UTAH ARCH

Performance of Ranked Subsetting

- 2-way rank subsets
- Worse performance for 3R/2B
- 3% performance increase over RTA

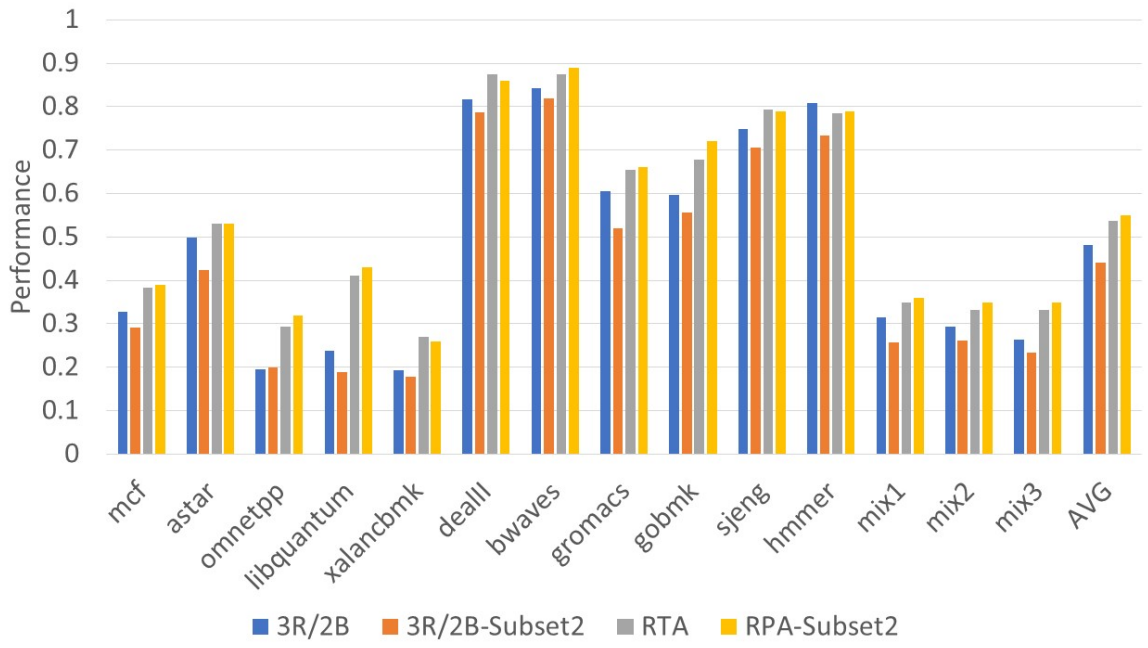

UTAH ARCH

of Ranked Subsets

- 4-way subsets – 18 cycle turns
- Sweet spot – 2-way rank subsets

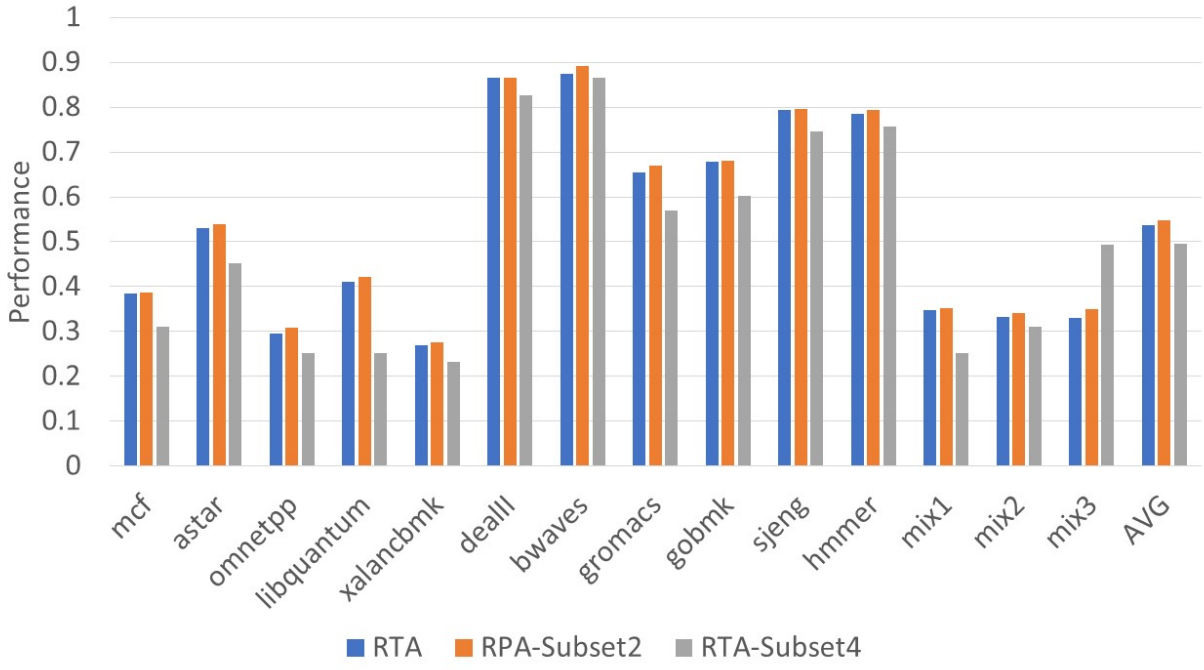

UTAH ARCH

Overall Results

- 11% performance increase over 3R/2B with RTA
- 14% with ranked subsetting (RPA-Subset2).

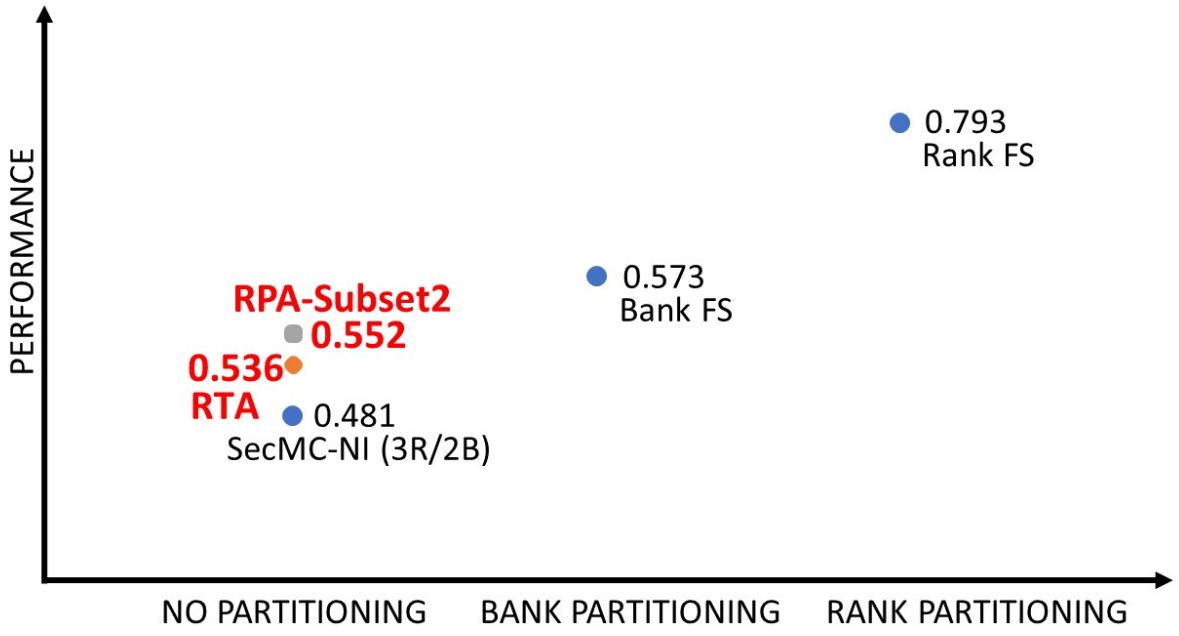

UTAH ARCH

Summary

- Shared memory controllers are vulnerable to timing channels
- RTA improves on the state of the art (SecMC-NI)
  - Reduced turn lengths
  - Additional constraints
  - Improved turn utilization tailored for low MLP
- Rank subsetting used to further improve performance
  - More parallelism
  - Fewer conflicts