Virtualized and Flexible ECC for Main Memory

Doe Hyun Yoon and Mattan Erez

Dept. Electrical and Computer Engineering
The University of Texas at Austin

Memory Error Protection

- Applying ECC uniformly – ECC DIMMs
  - Simple and transparent to programmers
- Error protection level
  - Fixed, design-time decision
- Chipkill-correct used in high-end servers
  - Constrain memory module design space
    - Allow only x4 DRAMs
    - Lower energy efficiency than x8 DRAMs
- ***Virtualized ECC* – objectives**
  - To provide *flexible* memory error protection
  - To *relax* design constraints of chipkill

Virtualized ECC

- Two-tiered error protection
- Tier-1 Error Code (T1EC)
  - Simple error code for detection or light-weight correction
- Tier-2 Error Code (T2EC)
  - Strong error correcting code
- Store T2EC within the memory namespace itself
  - OS manages T2EC
- Flexible memory error protection
  - Different T2EC for different data pages
  - Stronger protection for more important data

Virtualized ECC – Example

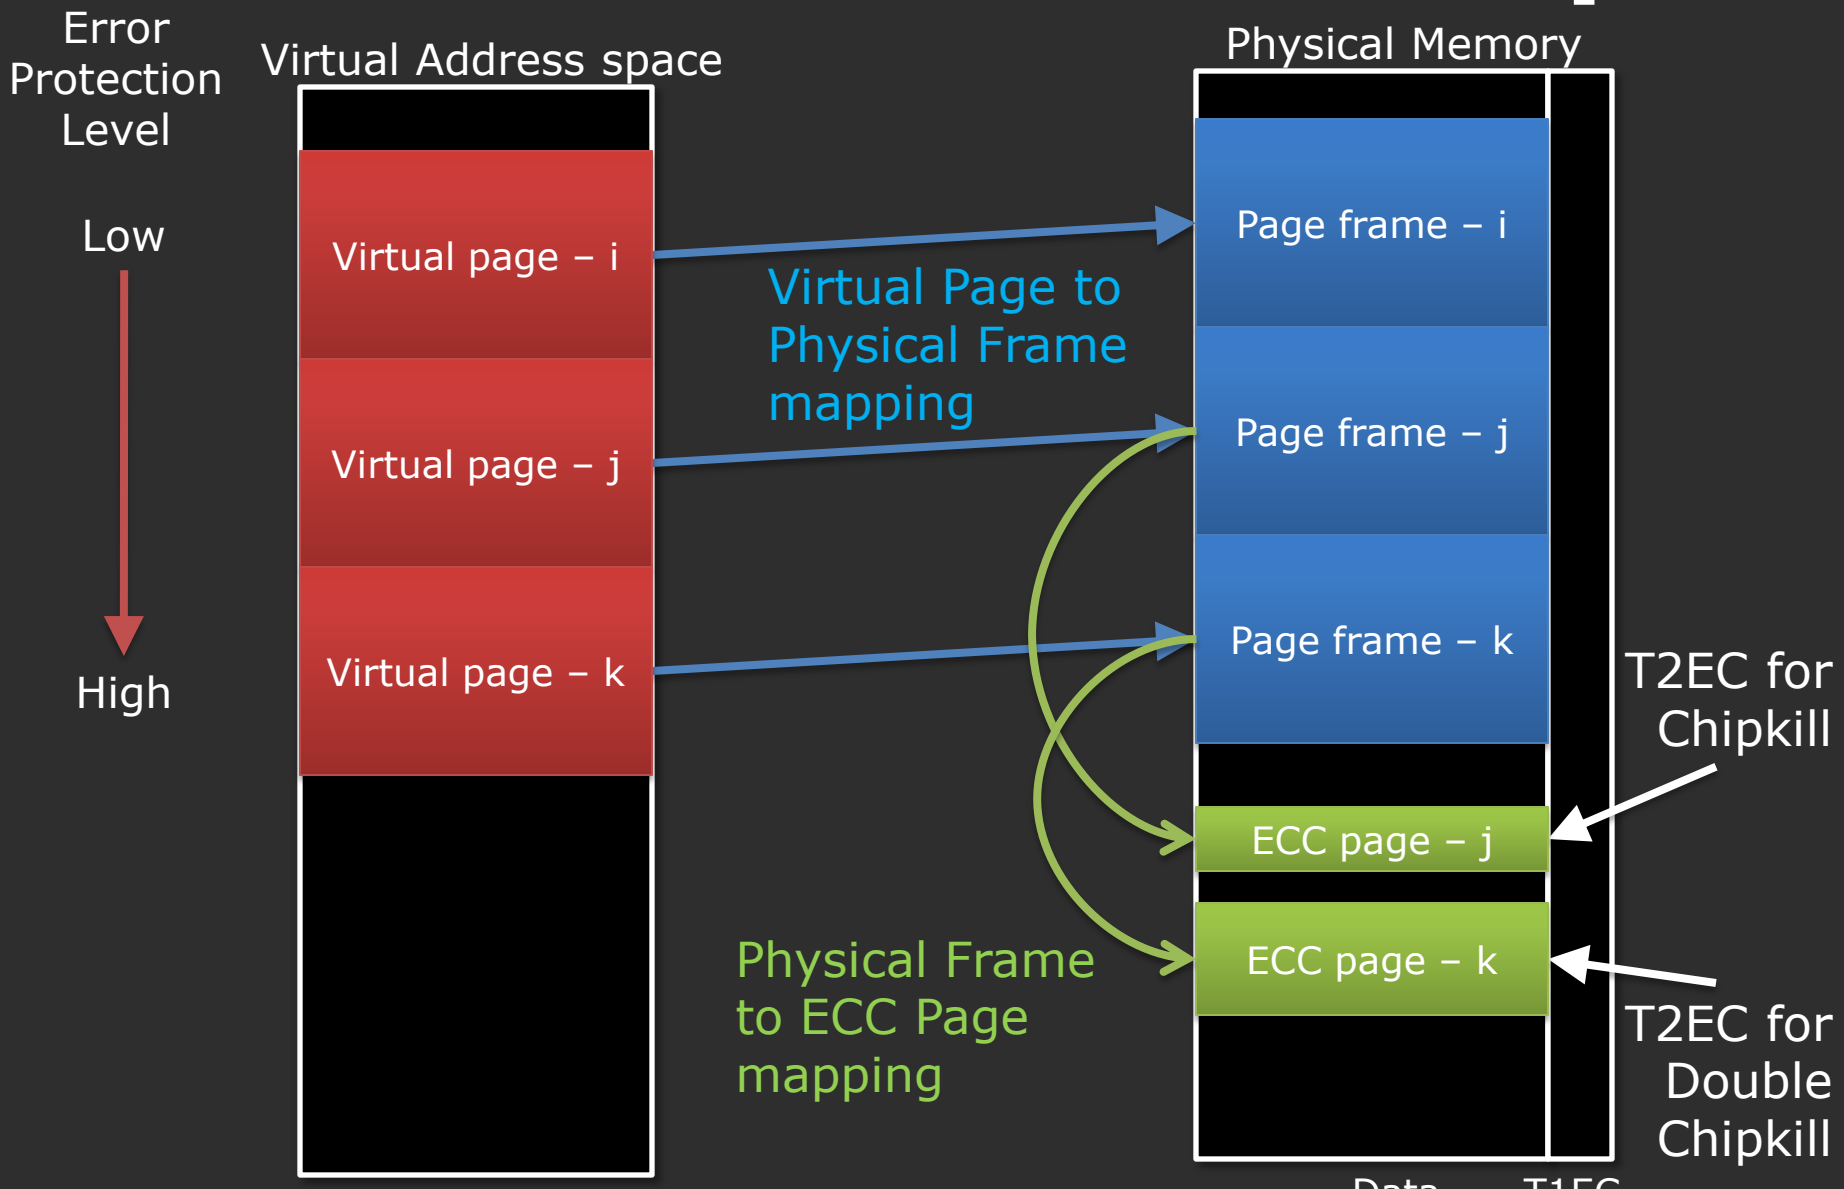

VIRTUALIZED ECC

Observations on Memory Errors

- Per-system error rate is still low
  - Most of time, we try to detect errors finding no error
- To detect errors is a common case operation
  - Need a low latency, low complexity error detection mechanism

  ➔ T1EC
- To correct errors is an uncommon case operation
  - Correction can be complex, take a long time
  - But, still need to manage error correction info somewhere

  ➔ Virtualized T2EC

Uniform ECC

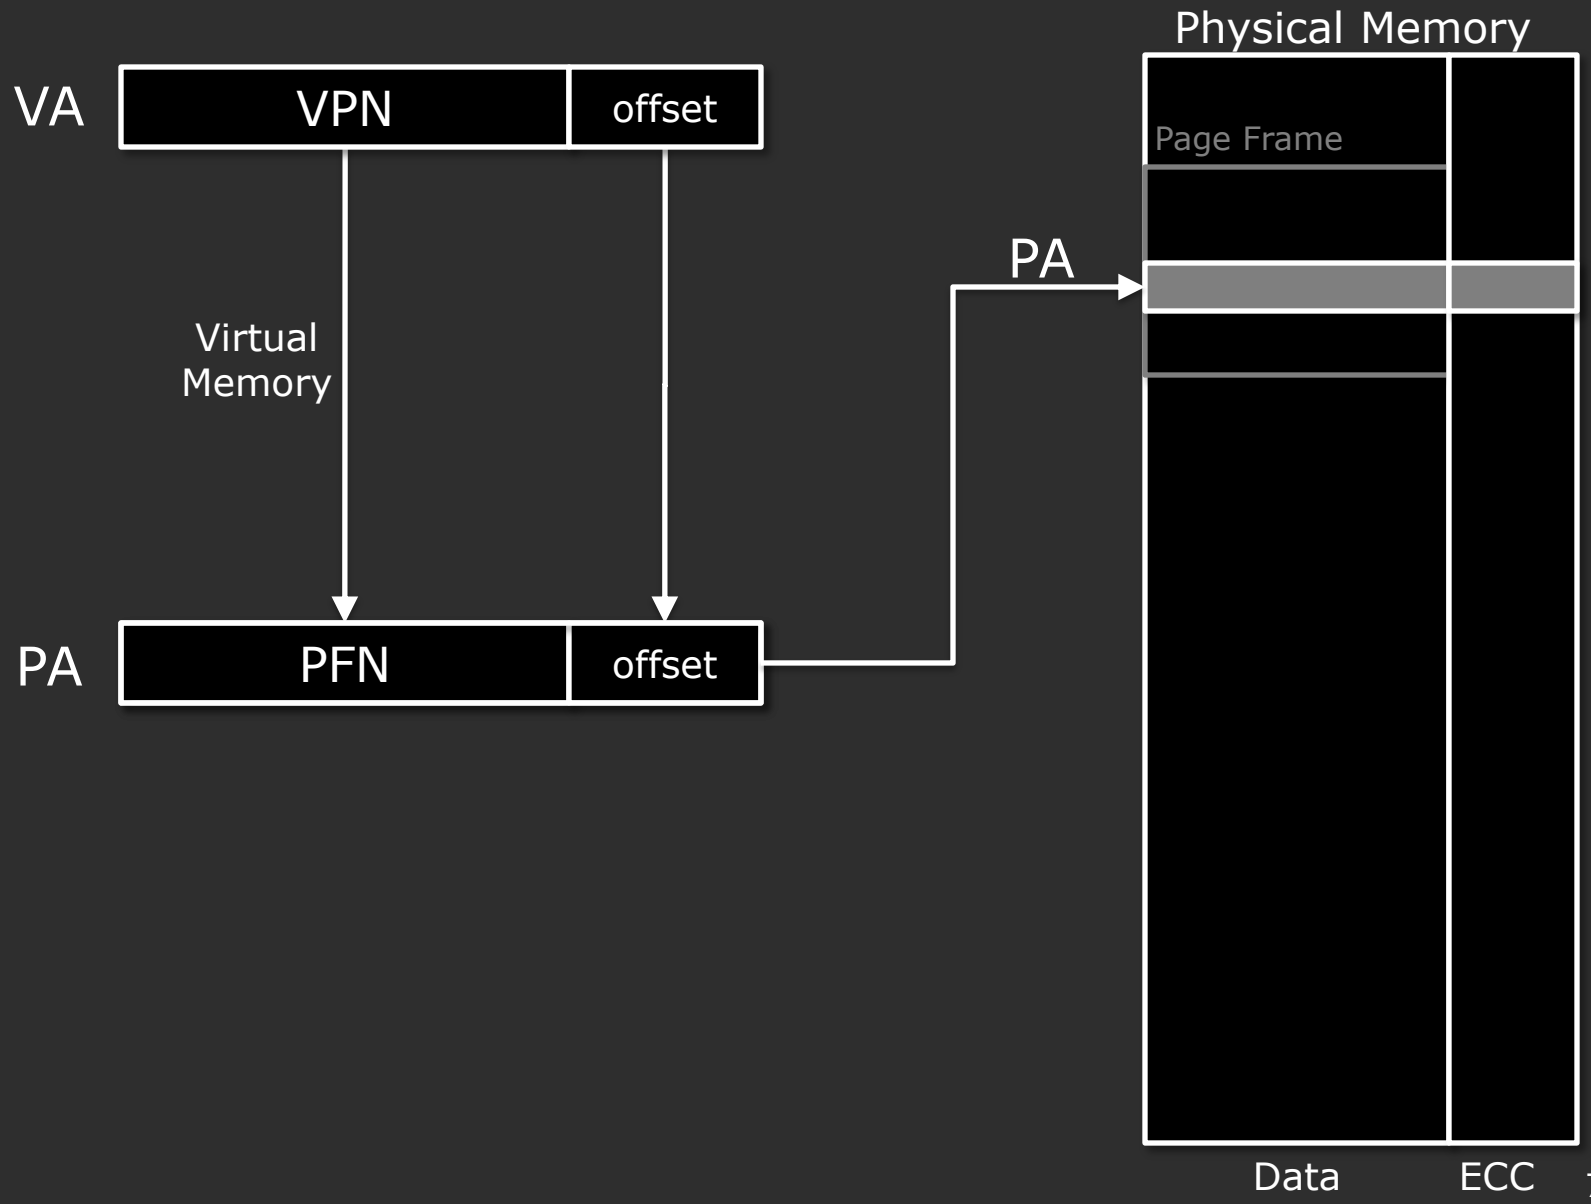
Physical Memory
VA
VPN
offset
Page Frame
PA
Virtual
Memory
PA
PFN
offset
Data
ECC

Virtualized ECC

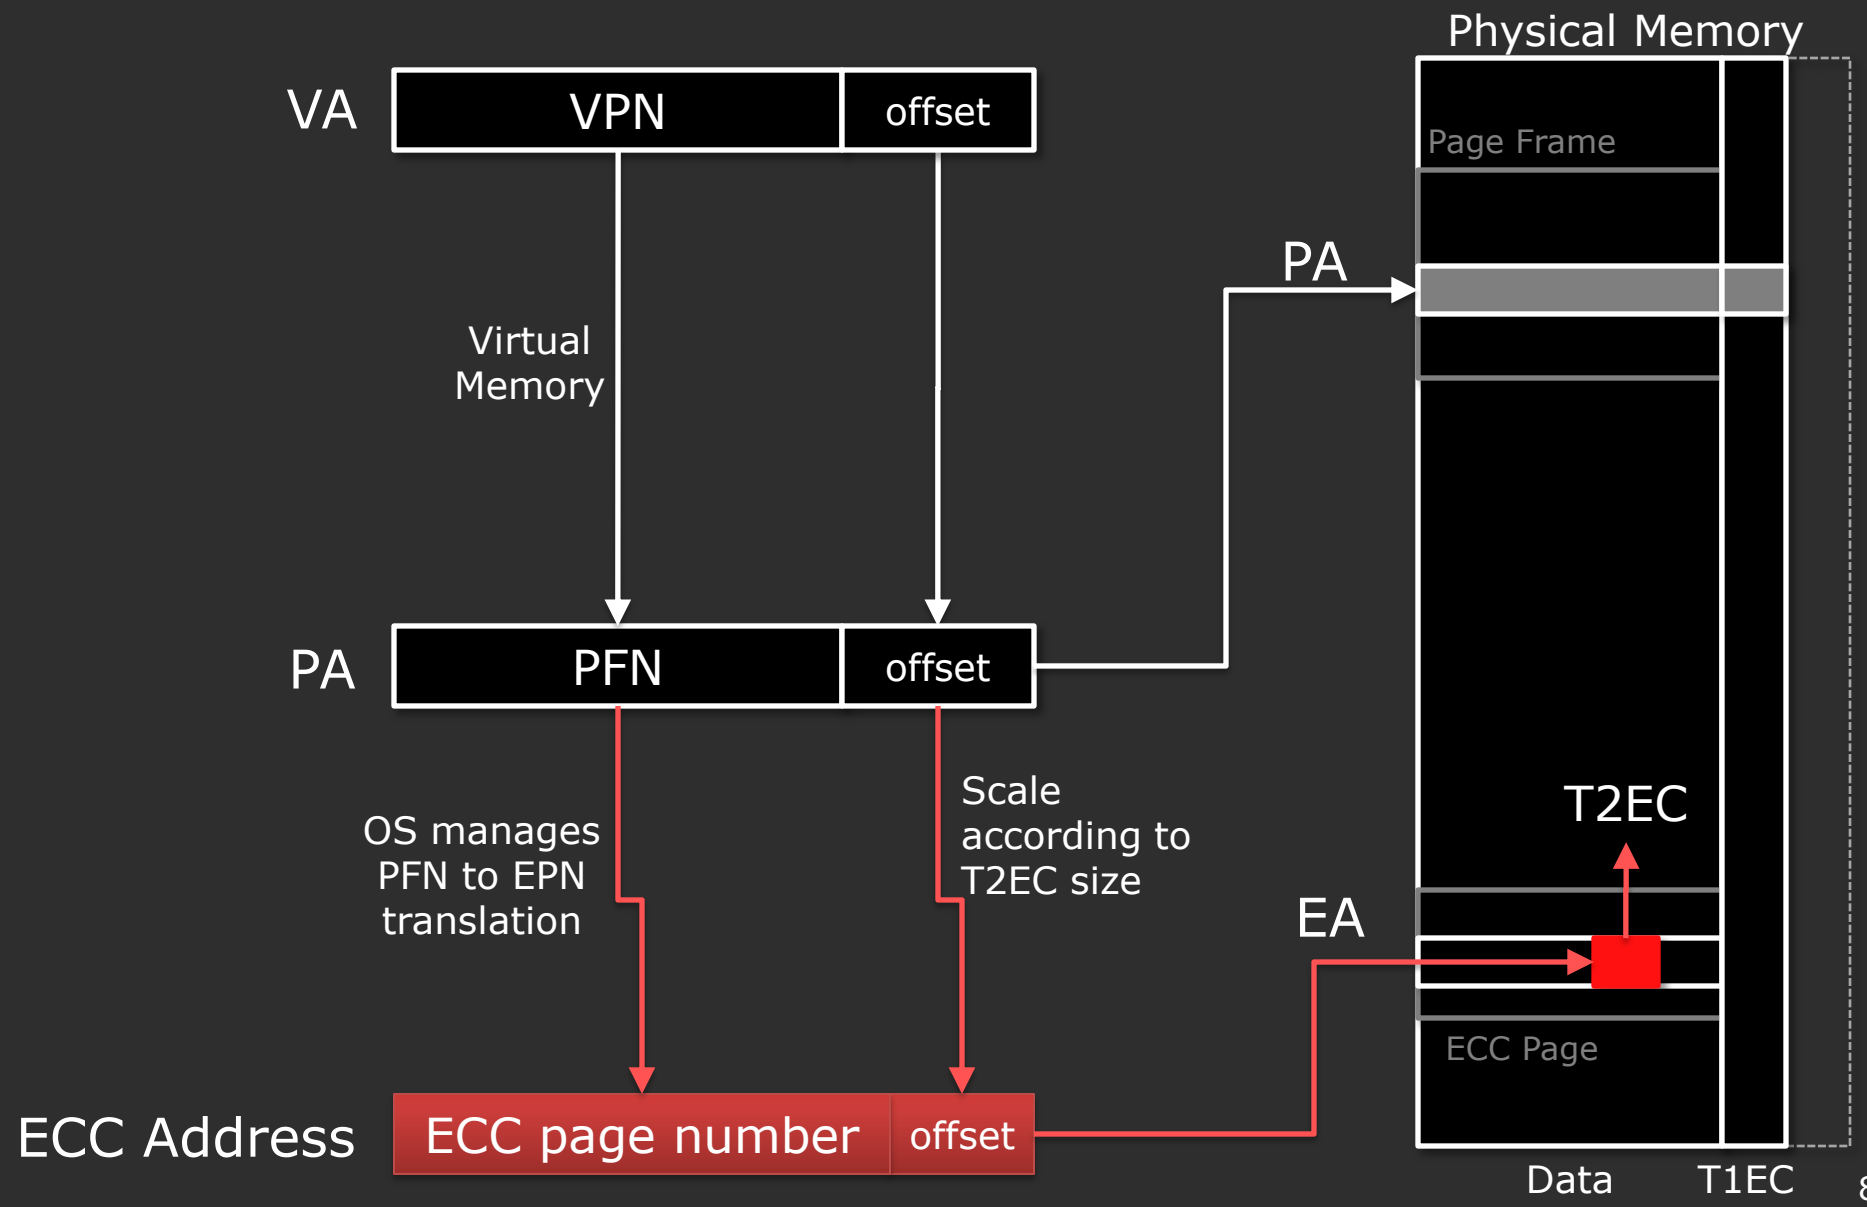

Update only valid T2EC to DRAM

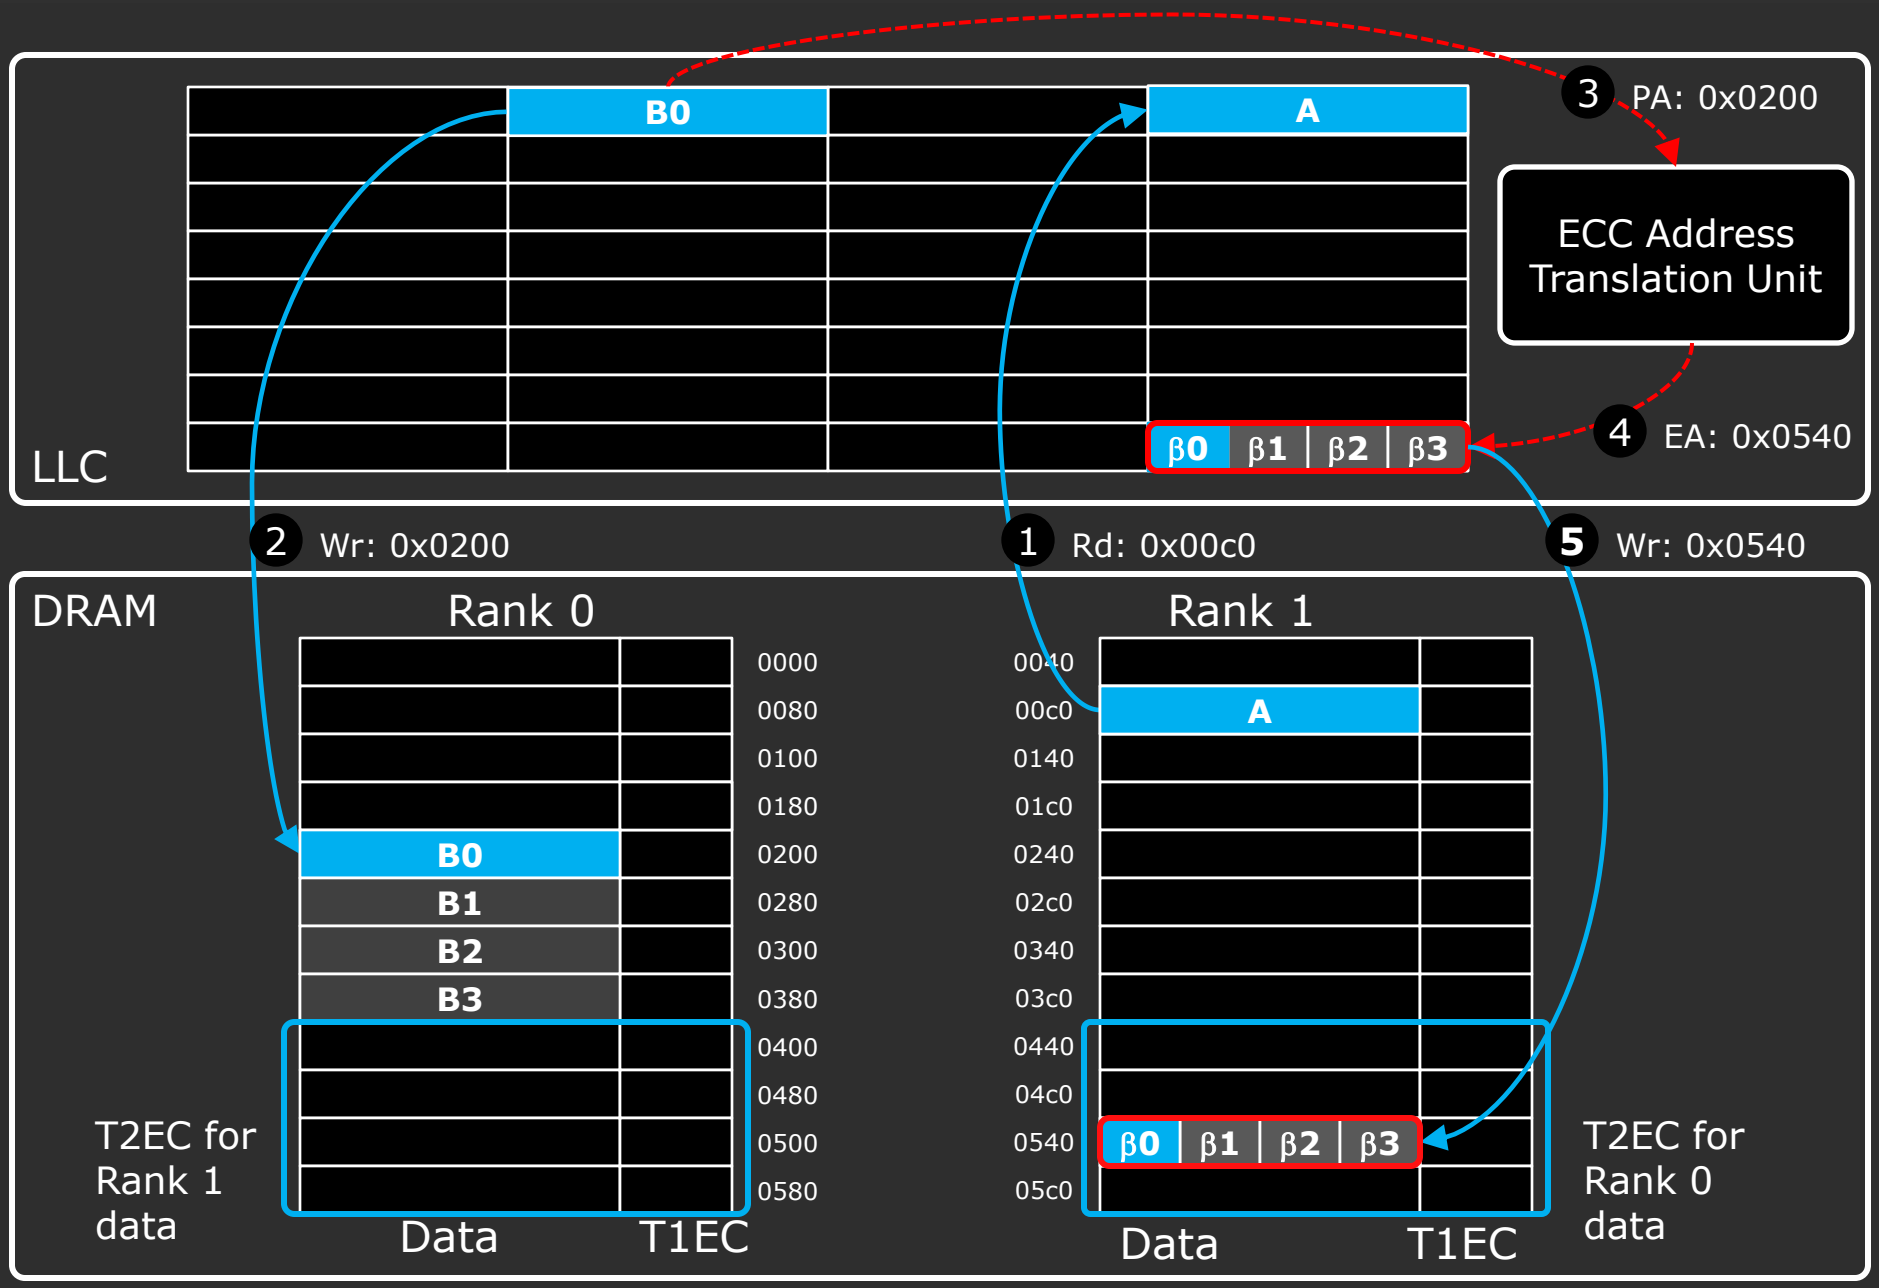

Penalty with V-ECC

- Increased data miss rate
  - T2EC lines in LLC reduce effective LLC size
- Increased traffic due to T2EC write-back
  - One-way write-back traffic
    - Not in a critical-path

CHIPKILL-CORRECT

Chipkill-correct

- Single Device-error Correct
  Double Device-error Detect
  - Can tolerate a DRAM failure
  - Can detect a second DRAM failure

- Chipkill requires x4 DRAMs
- x8 chipkill is impractical
  - But, x8 DRAM is more energy efficient

Baseline x4 Chipkill

- Two x4 ECC DIMMs
  - 128bit data + 16bit ECC (redundancy overhead: 12.5%)
  - 4 check symbol error code using 4-bit symbol
- Access granularity
  - 64B in DDR2 (min. burst 4 x 128 bit)
  - 128B in DDR3 (min. burst 8 x 128 bit)

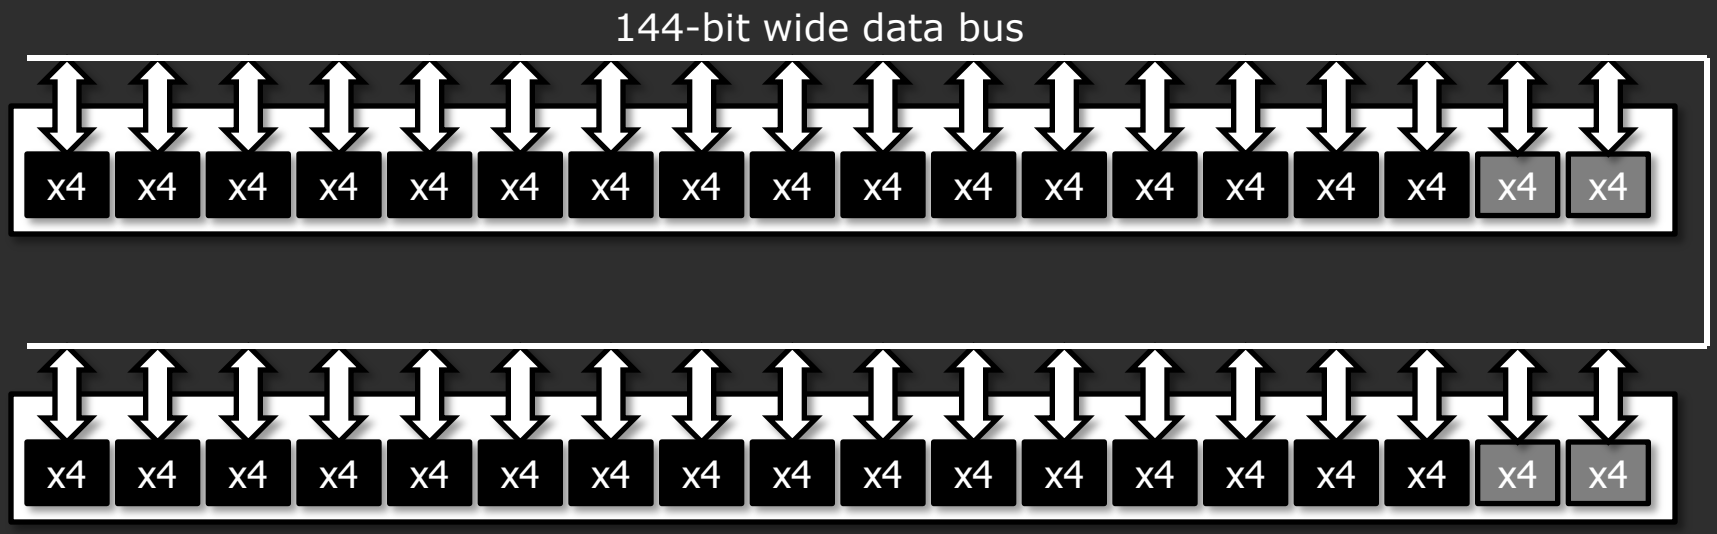

x8 Chipkill

- x8 chipkill with the same access granularity
  - 152-bit wide data path
    - 128-bit data + 24-bit ECC
    - Redundancy overhead: 18.75%
  - Need a custom-designed DIMM
    - Increase the system cost a lot

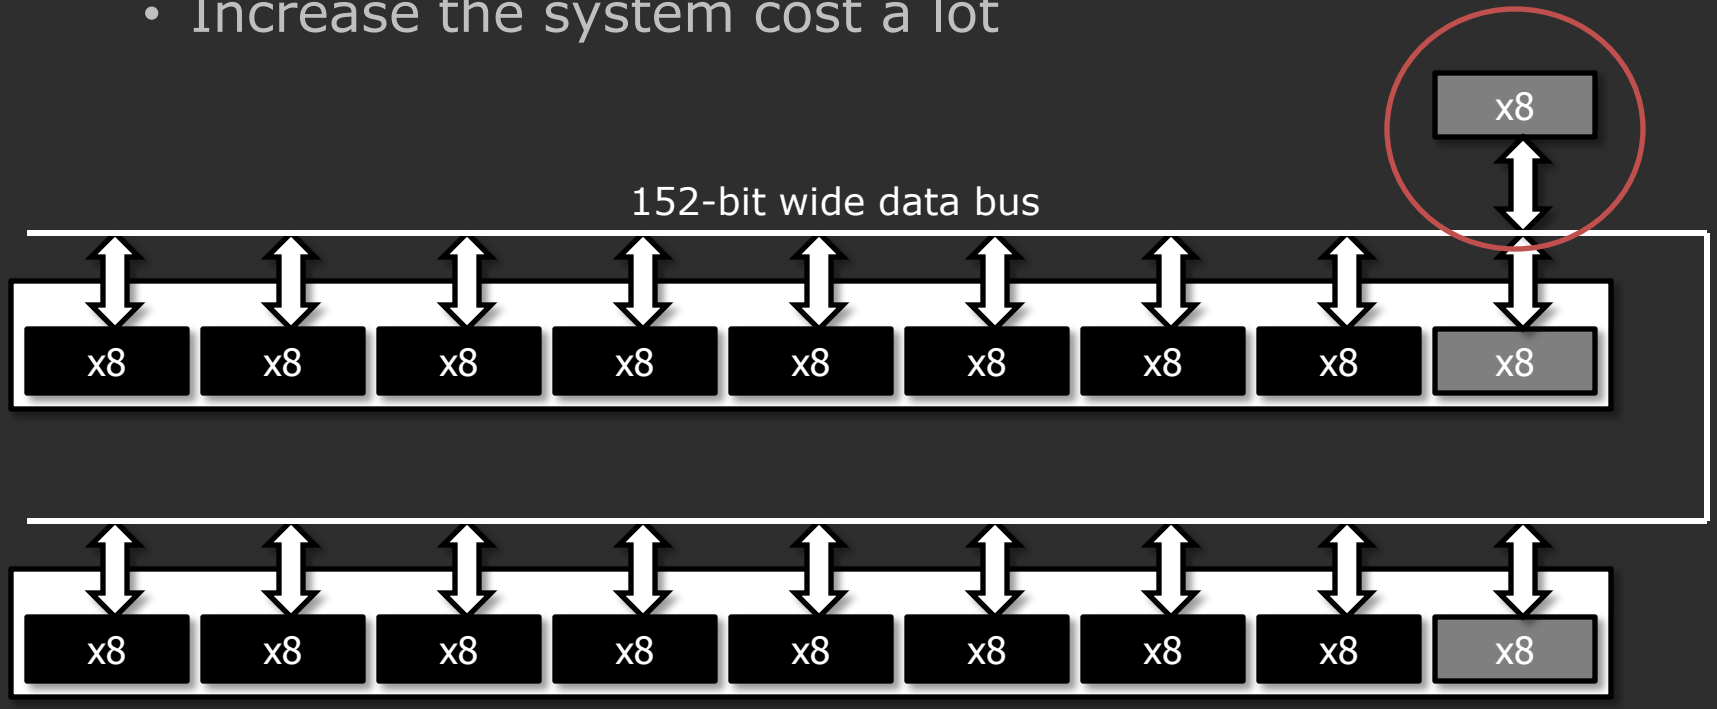

x8 Chipkill /w Standard DIMMs

- Increase access granularity
  - 128B in DDR2 (min. burst 4 x 256 bit)
  - 256B in DDR3 (min. burst 8 x 256 bit)

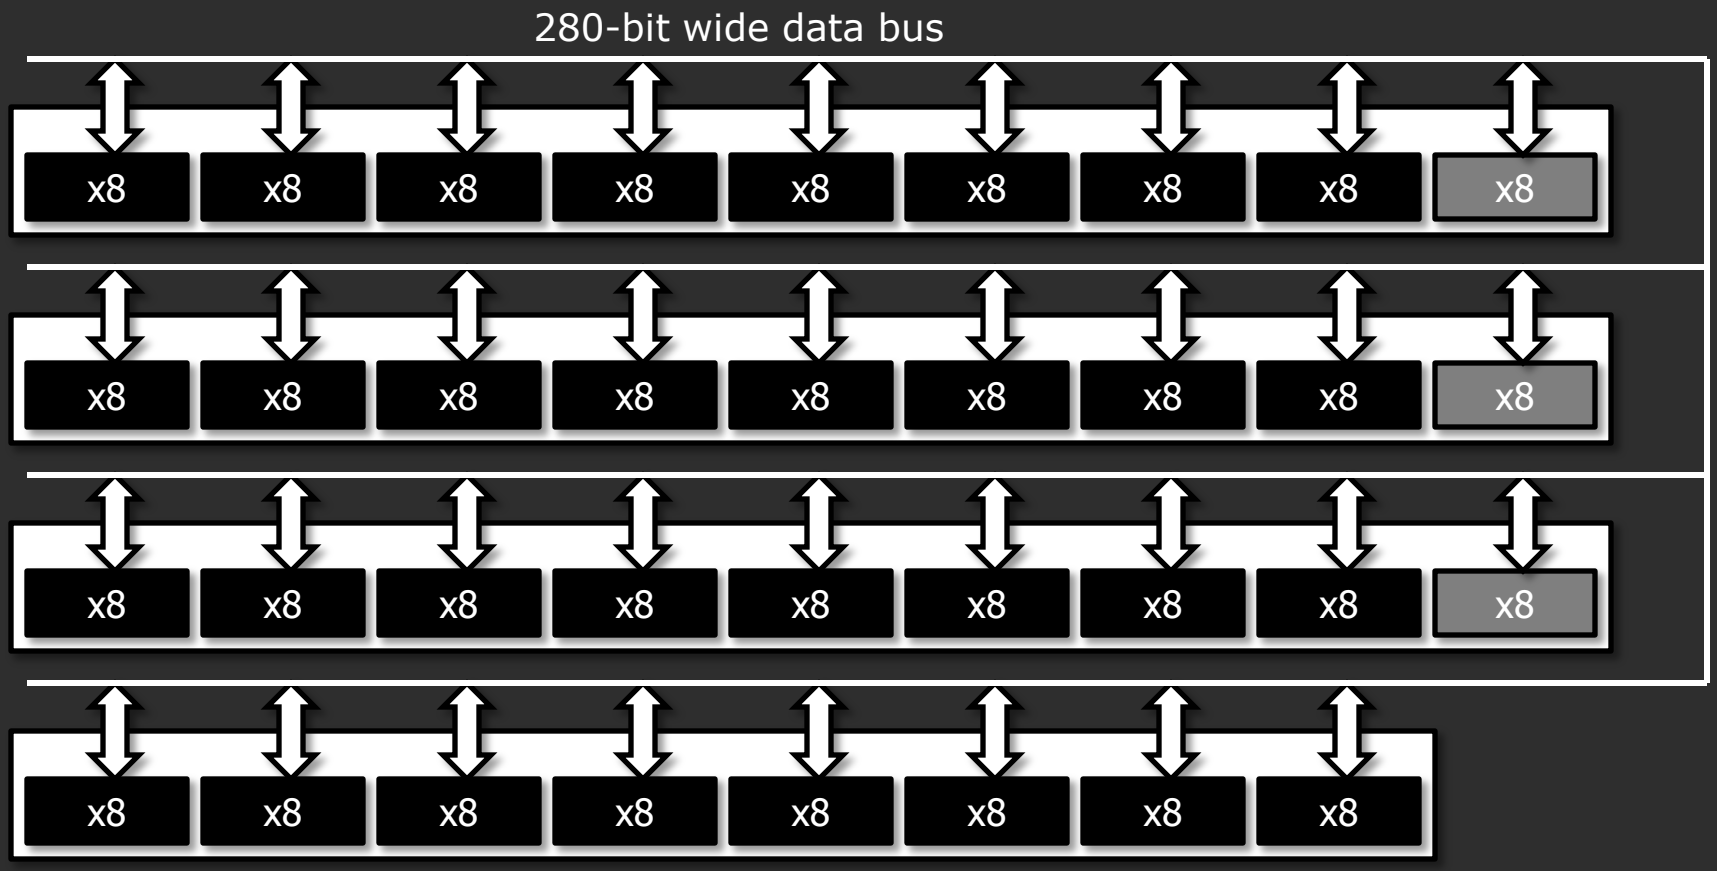

V-ECC for Chipkill

- Use 3 check symbol error codes
  - Single Symbol-error Correct and Double Symbol-error Detect
- T1EC
  - 2 check symbols
  - Detect up to 2 symbol error
- T2EC
  - 3rd check symbol
  - Combined T1EC/T2EC provides Chipkill

V-ECC: ECC x4 configuration

- Use 8-bit symbol error code
  - 2 bursts out of a x4 DRAM form an 8bit-symbol
    - Modern DRAMs have minimum burst of 4 or 8
- 1 x4 ECC DIMM + 1 x4 Non-ECC DIMM
- Each DRAM access in DDR2 (burst 4)
  - 64B data, 4B T1EC
  - 2B T2EC is virtualized within memory namespace
    - 32 T2ECs per 64B cache line

Virtualized within memory

T2EC

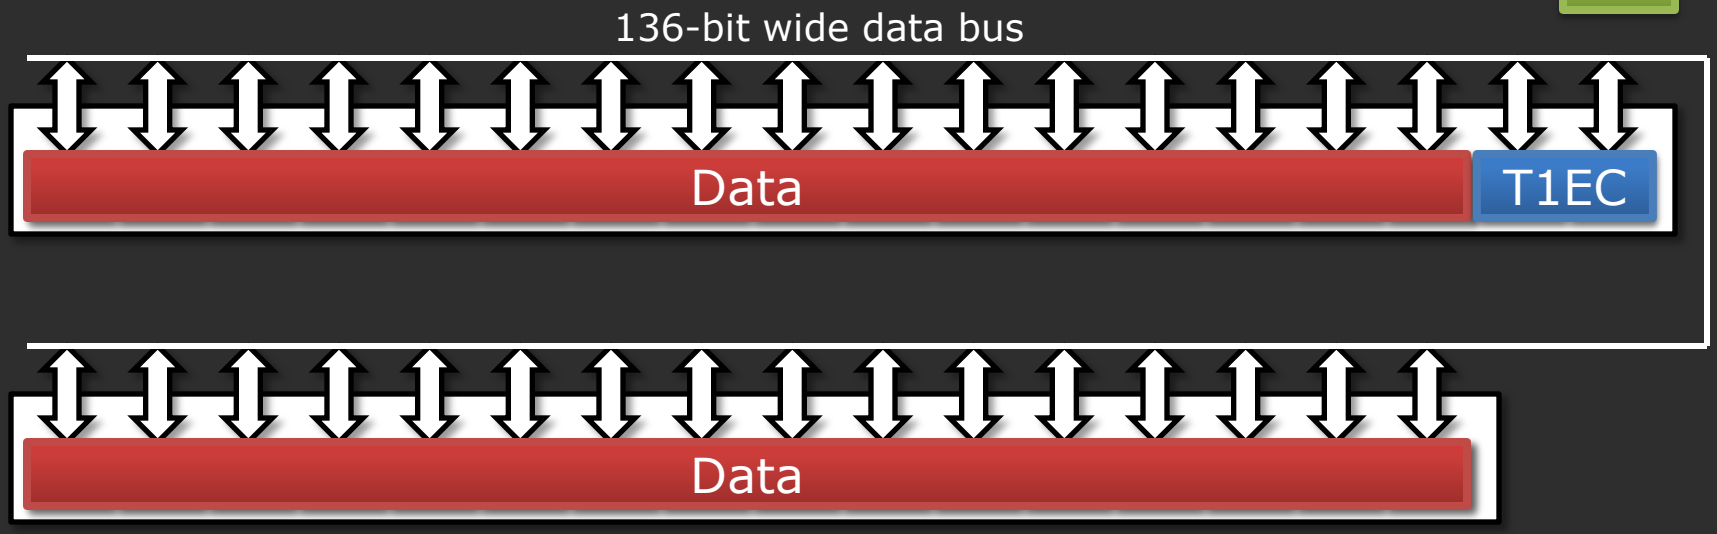

V-ECC: ECC x8 configuration

- Use 8-bit symbol error code
- 2 x8 ECC DIMMs
- Each DRAM access in DDR2 (burst 4)
  - 64B data, 8B T1EC
  - 4B T2EC is virtualized
    - 16 T2ECs per 64B cache line

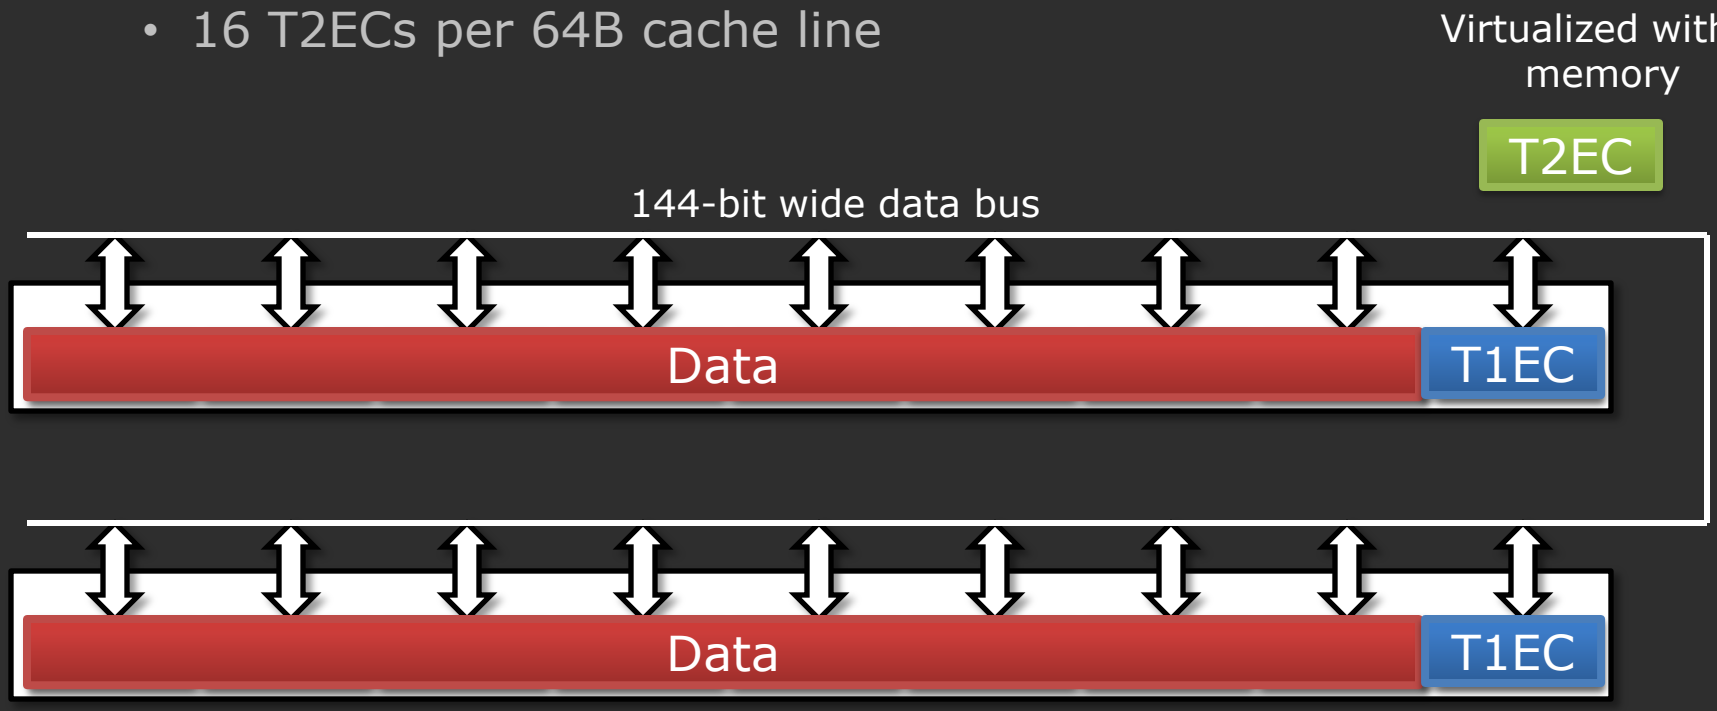

Flexible Error Protection

- Single HW with V-ECC can provide
  - Chipkill-detect, Chipkill-correct, and Double chipkill-correct
  - Use different T2EC for different pages

	Chipkill-Detect	**Chipkill-Correct**	**Double Chipkill-Correct**
ECC x4	0B	2B	4B
ECC x8	0B	4B	8B

- Maximize performance/power efficiency with Chipkill-Detect
- Stronger protection at the cost of additional T2EC access

EVALUATION

Simulator/Workload

- GEMS + DRAMsim
  - An out-of-order SPARC V9 core
  - Exclusive two-level cache hierarchy
  - DDR2 800MHz – 12.8GB/s (128-bit wide data path)
    - 1 channel 4 ranks
- Power model
  - WATTCH for processor power – scaled to 45nm
  - CACTI for cache power – cacti 45nm
  - Micron model for DRAM power – commodity DRAMs
- Workloads
  - 12 data intensive applications from SPEC CPU 2006 and PARSEC
  - Microbenchmarks: STREAM and GUPS

Normalized Execution Time

- Less than 1% penalty on average
- Performance penalty ↑
  - Spatial locality ↓
  - Write-back traffic ↑

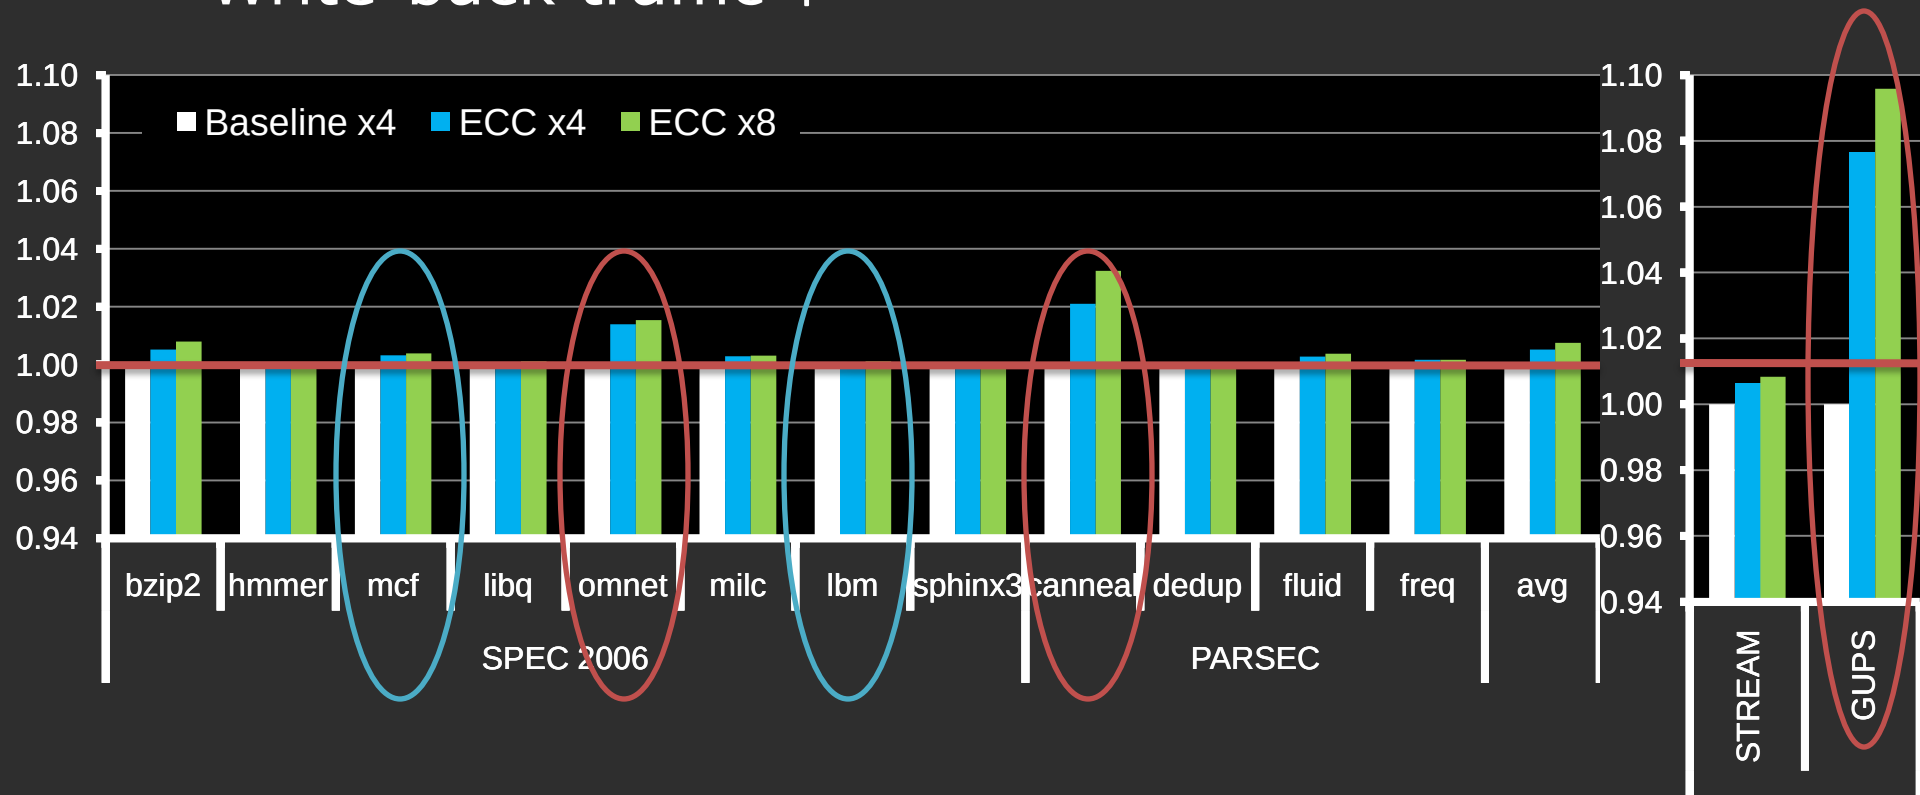

System Energy Efficiency

- Energy Delay Product (EDP) gain
  - ECC x4: 1.1% on average
  - ECC x8: 12.0% on average

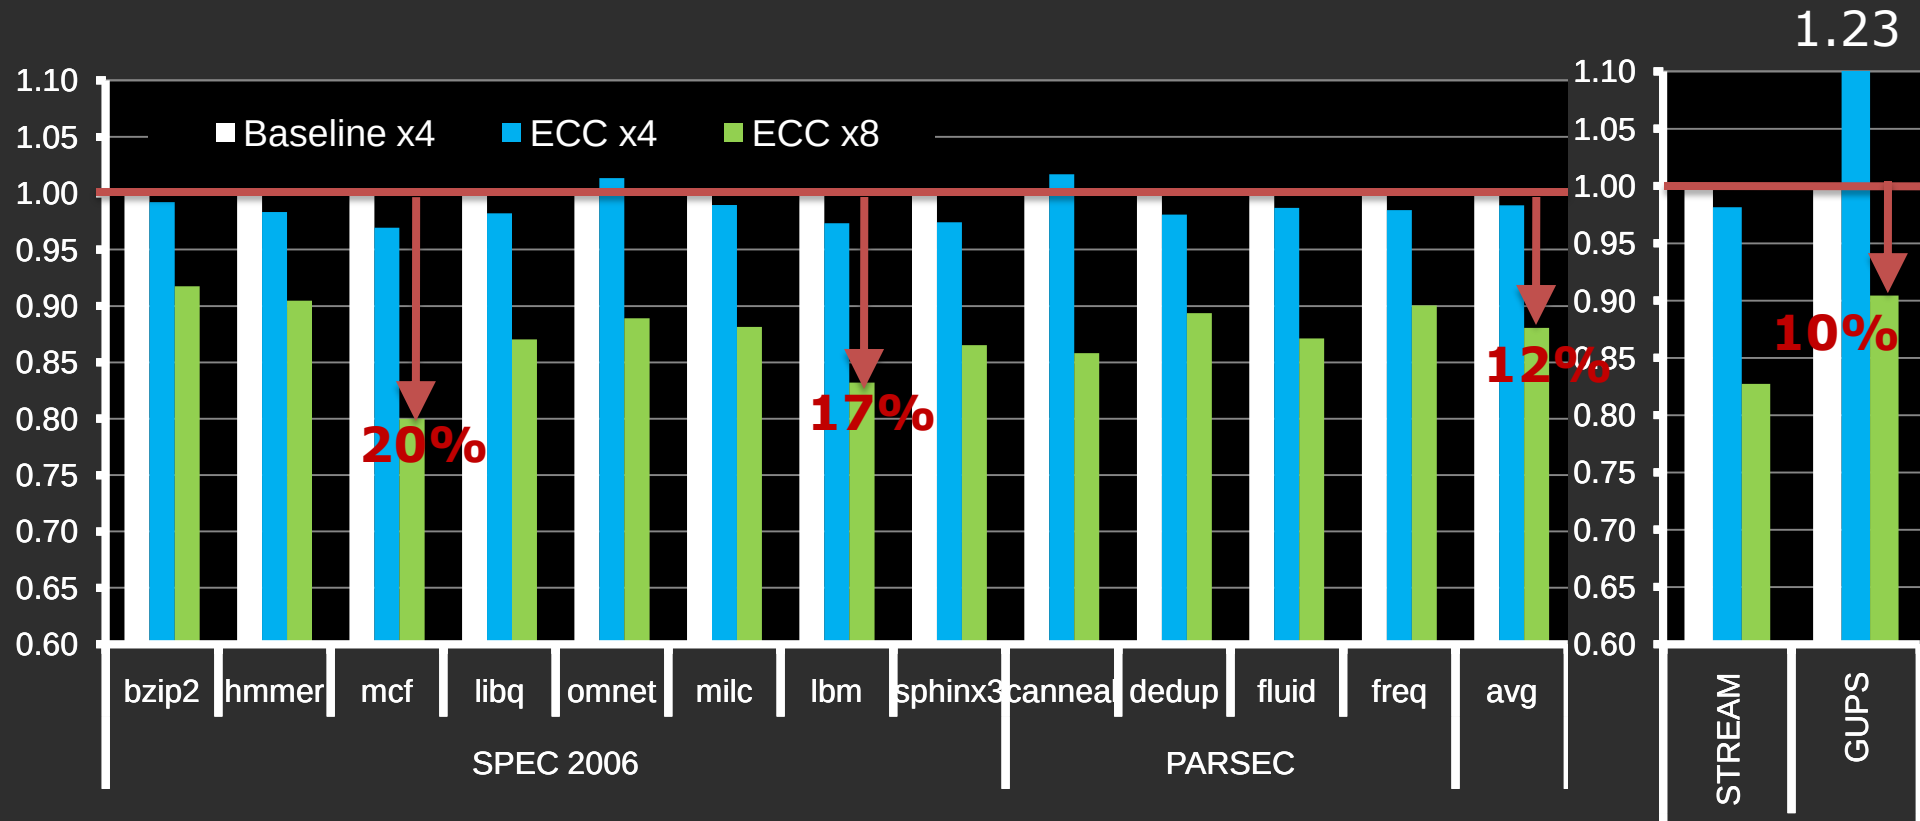

Flexible Error Protection

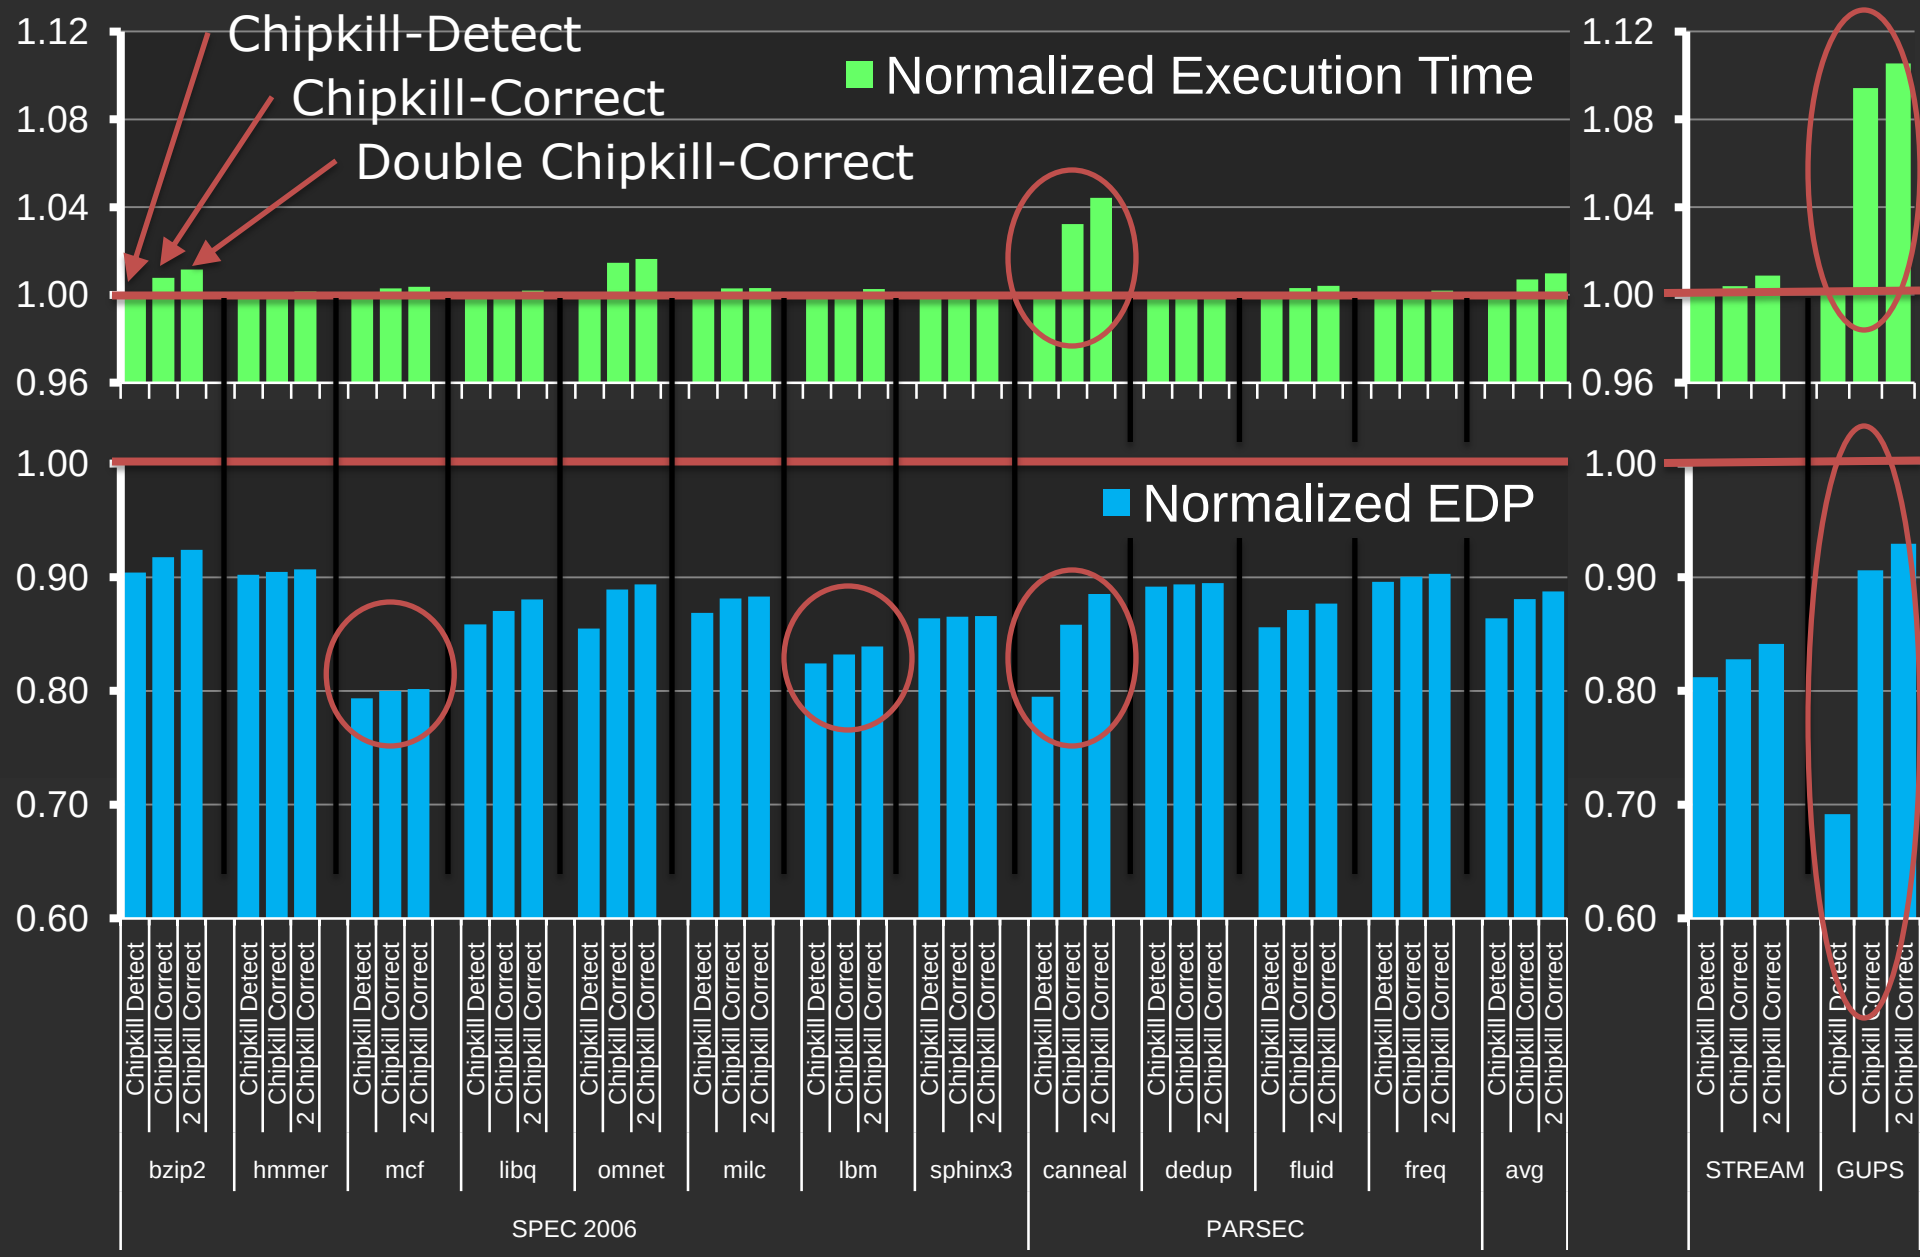

Conclusion

- Virtualized ECC
  - Two-tiered error protection, virtualized T2EC
- Improved system energy efficiency with chipkill
  - Reduce DRAM power consumption by 27%
  - Improve system EDP by 12%
- Performance penalty – 1% on average
- Error protection even for Non-ECC DIMMs
  - Can be used for GPU memory error protection
- Flexibility in error protection
  - Adaptive error protection level by user/system demand
  - Cost of error protection is proportional to protection level

Virtualized and Flexible ECC for Main Memory

Doe Hyun Yoon and Mattan Erez

Dept. Electrical and Computer Engineering
The University of Texas at Austin