

# Virtualized ECC: Flexible Reliability in Memory Systems

---

Doe Hyun Yoon  
Advisor: Mattan Erez

Electrical and Computer Engineering  
The University of Texas at Austin

- Reliability concerns are growing
  - Memory failure rate is increasing
    - Shrinking semiconductor design rules & lower V<sub>dd</sub>
  - More and more memory cells
    - Higher capacity & larger number of nodes in a system
  - Industry consistently requires high reliability levels
    - Chipkill and even higher reliability levels
- Traditional solutions
  - Apply ECC uniformly across memory locations
    - Not all apps/data need same error tolerance levels
  - Cost increases with tolerance levels
    - Hard to meet the required level of reliability at low cost

# Virtualized ECC


Uniform ECC



Data block

ECC

Virtualized ECC



Virtualized ECC

Stronger ECC

Finer access granularity

- Store redundant info within memory hierarchy
- Dynamic mapping between data and ECC
  - Flexible protection level/access granularity
- Two-tiered protection

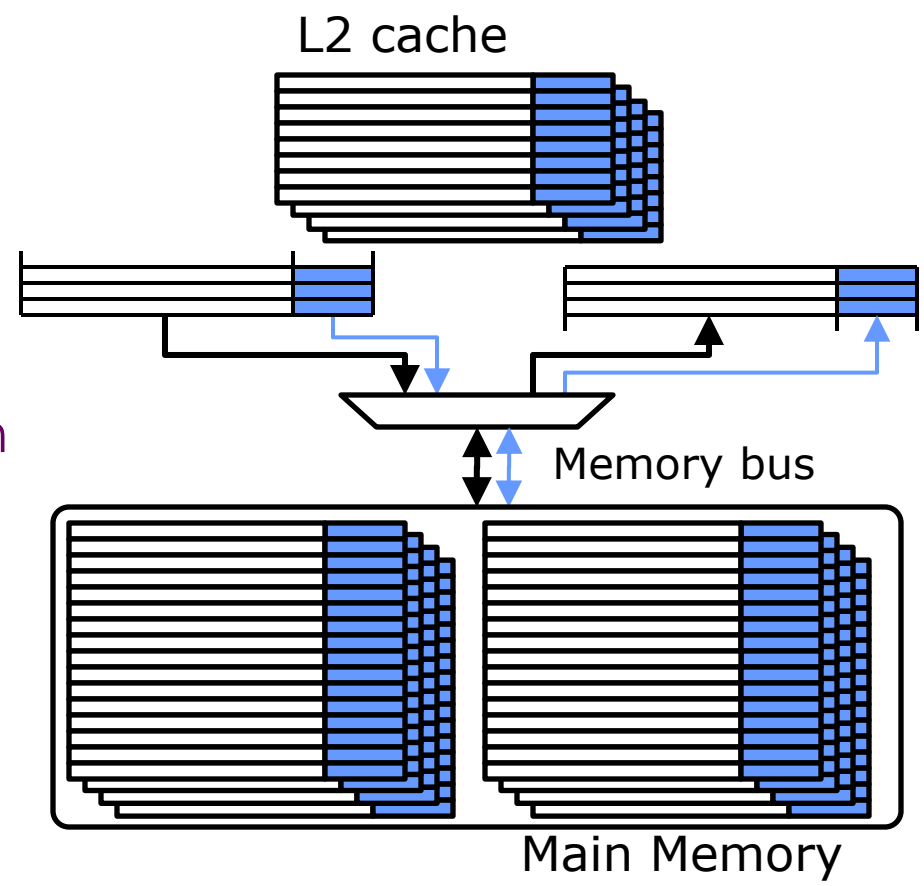
# Outline

---

- Motivation
- Traditional Memory Protection
- Virtualizing Redundant Information
- Memory Mapped ECC (MME)
- Virtualized and Flexible ECC (V-ECC)
- Adaptive Granularity Memory System (AGMS)
- Conclusions and Future Work

# Traditional Memory Protection

- Uniform ECC
  - Simple and transparent to the programmer
- Dedicated and aligned HW resources
  - Waste storage and bandwidth
  - Better reliability at higher cost
- Design-time decision
  - Error tolerance level
  - Access granularity



**Users *always* pay for the cost of memory protection**

# Virtualizing Redundant Information

---

- Store all or part of redundant information within the memory hierarchy
  - Minimize dedicated resources to redundant information
  - Two-tiered protection
    - Tier-1 error code: Low-cost, low-complexity for the common case
    - Tier-2 error code: Strong error correcting code, but rarely accessed
  - Decouple data and ECC storage
- Flexibility in memory protection
  - Adaptive error tolerance levels / access granularities
- **Users pay for the error tolerance they need, rather than overpay for what they might need**

# Examples of Virtualizing ECC

---

- Last-level cache protection with memory mapped ECC
  - **Memory Mapped ECC (MME) [ISCA'09]**
  - ECC FIFO [SC'09]
- Chipkill level flexible main memory protection
  - **Virtualized and Flexible ECC (V-ECC) [ASPLOS'10]**
- Adaptive memory access granularity with ECC
  - **Adaptive Granularity Memory System (AGMS) [On-going]**


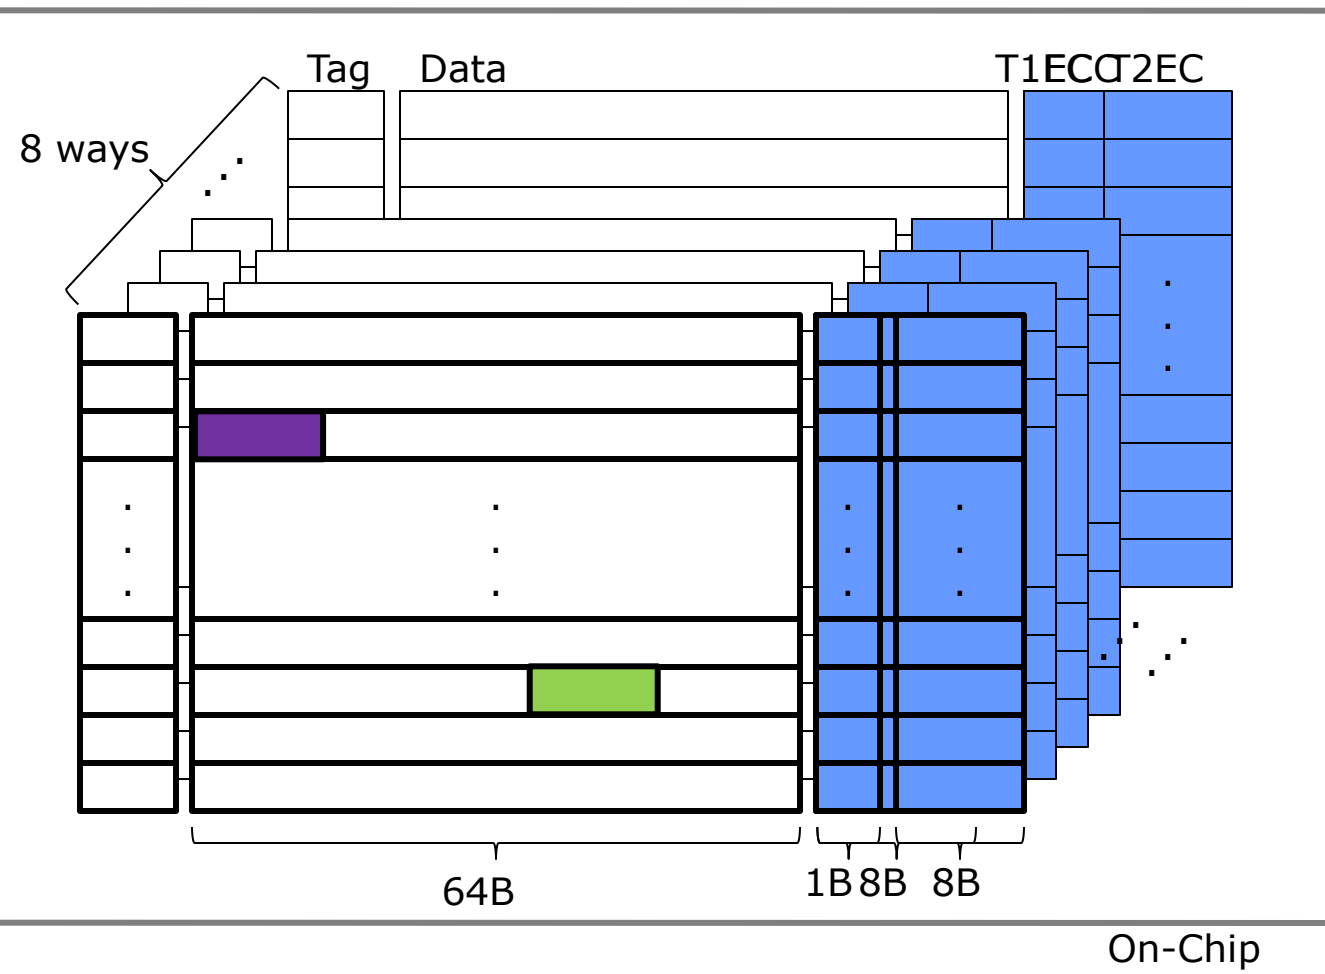


---

# Memory Mapped ECC [ISCA'09]



# Memory Mapped ECC [ISCA'09]



T2EC is memory mapped and cacheable

# Memory Mapped ECC (Cont'd)

---

- Last-level cache protection mechanism
- Two-tiered protection
  - Tier-1 error code (T1EC)
    - Low-overhead on-chip error code
  - Tier-2 error code (T2EC)
    - Strong memory mapped error correcting code
  - LLC is dynamically and transparently partitioned into data and T2EC
- Fixed, one-on-one mapping between physical cache lines and memory mapped T2EC
- Area saving: 15%, power saving: 9%



---

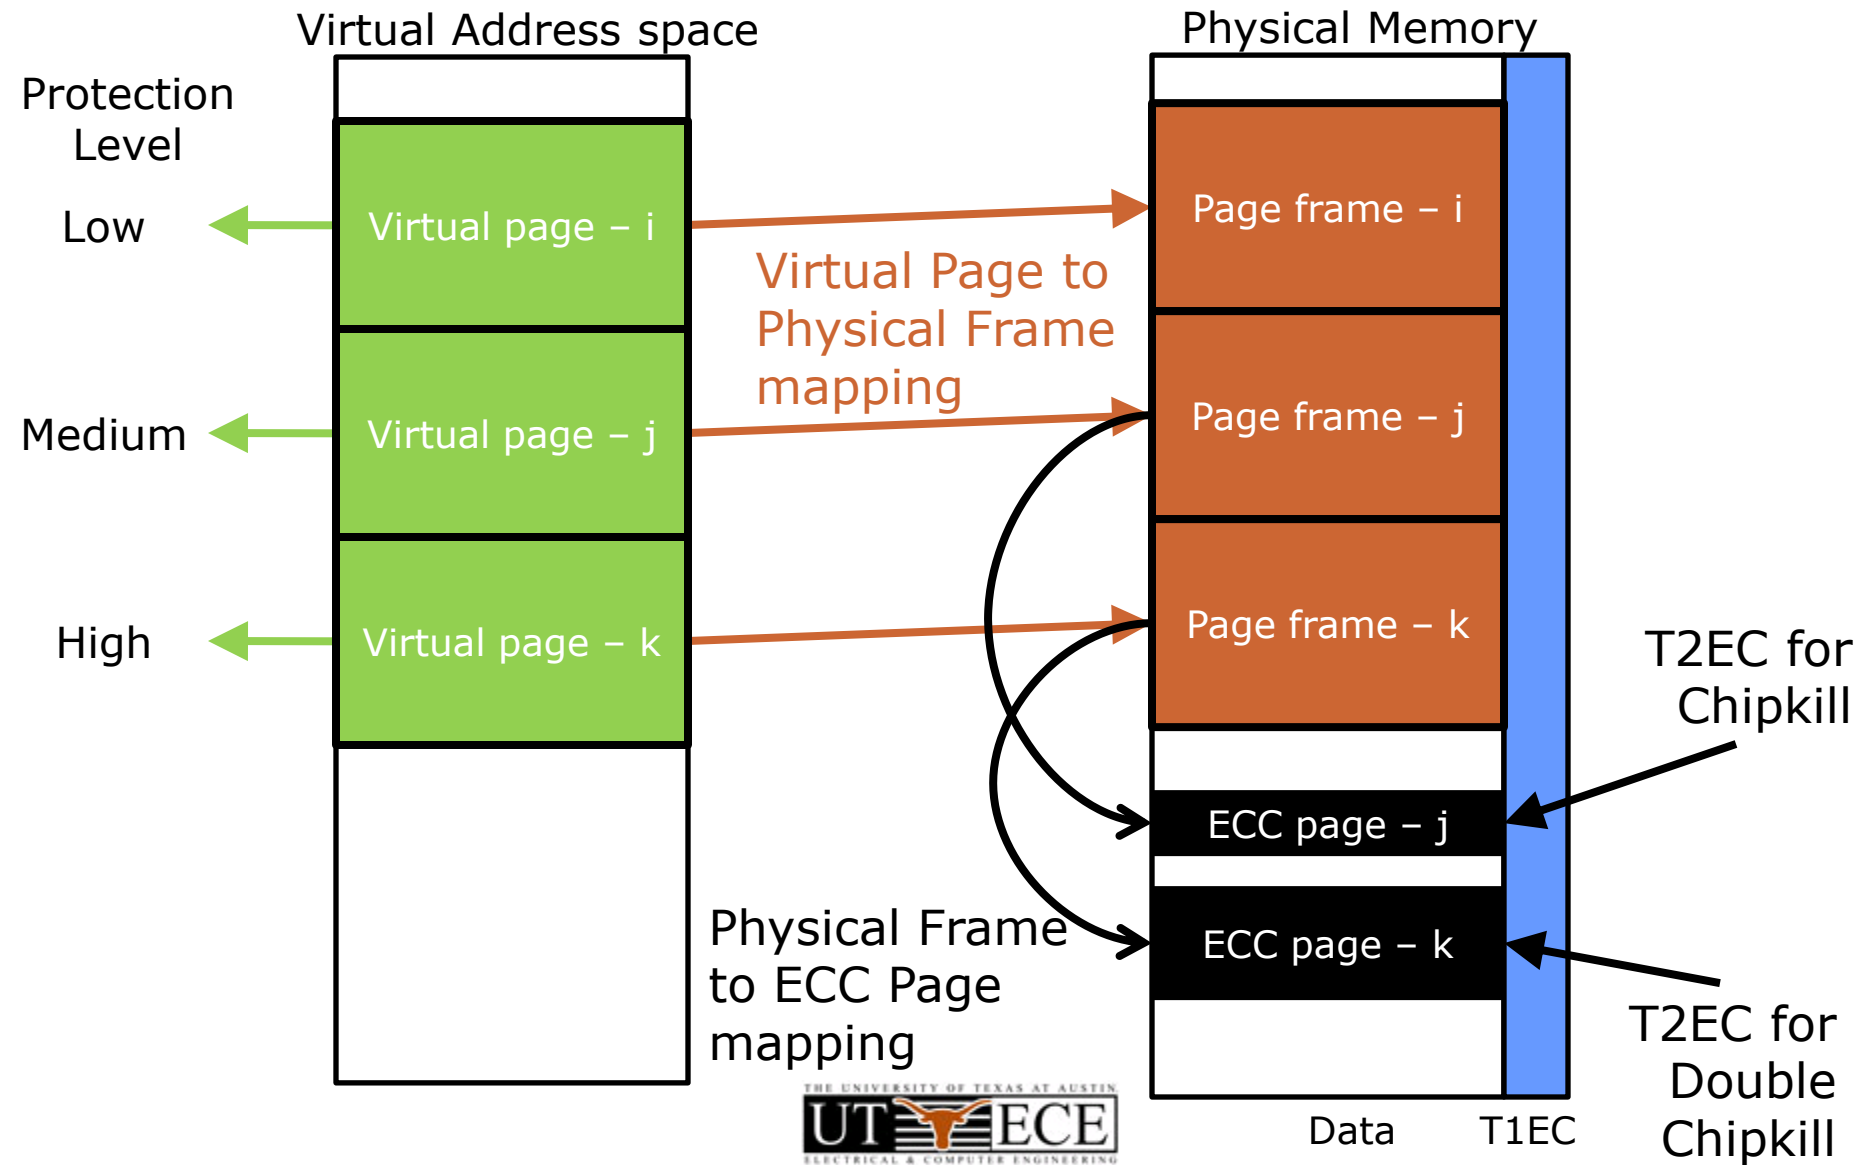
# Virtualized and Flexible ECC [ASPLOS'10]

# Virtualized and Flexible ECC [ASPLOS'10]

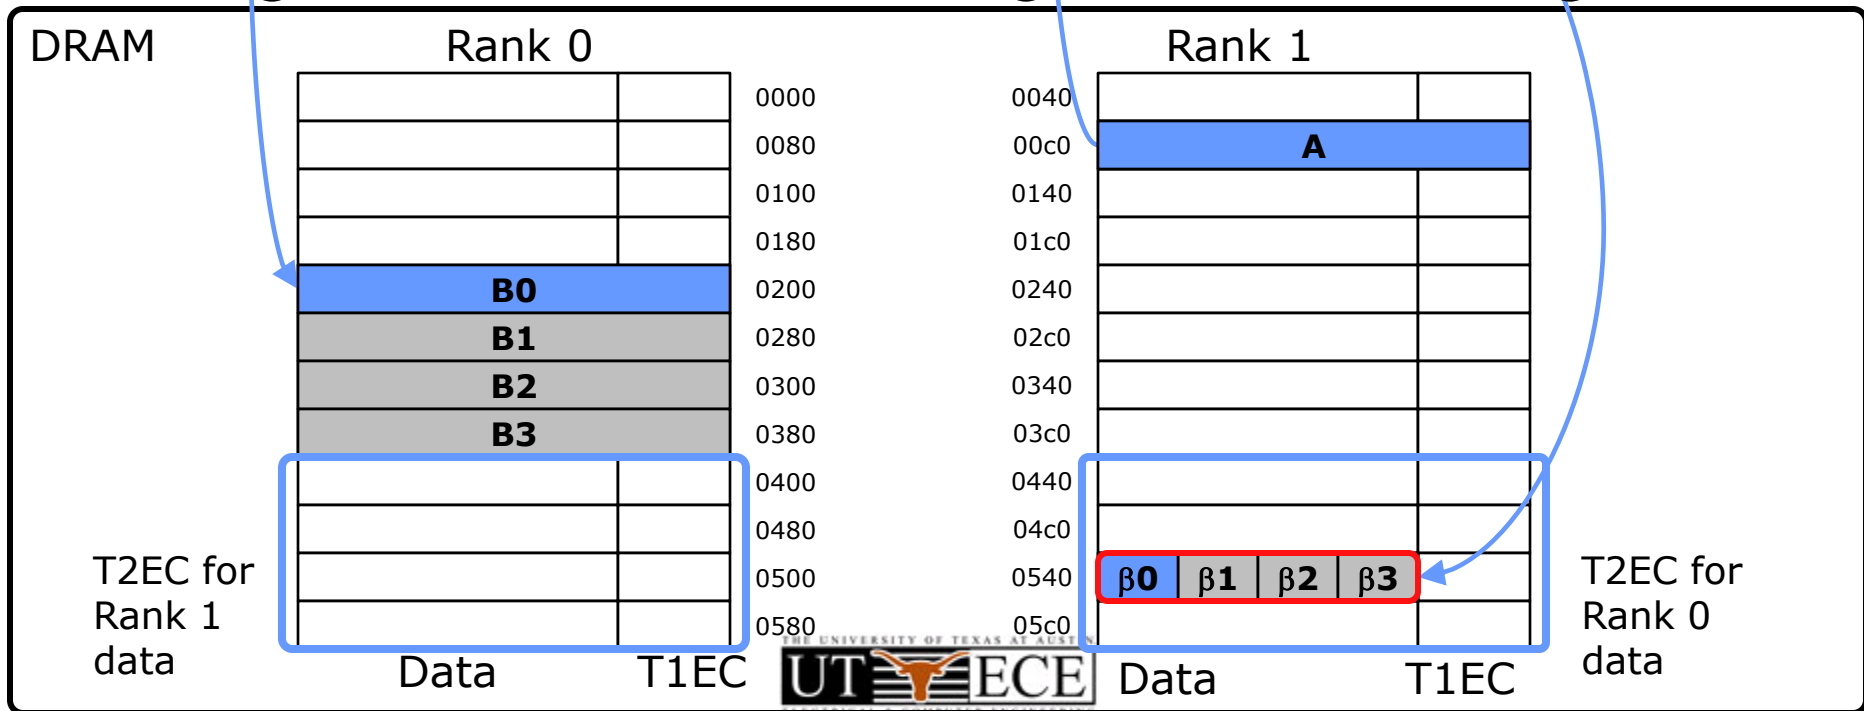

---

- Main memory protection mechanism
- Virtualize redundant information within the memory hierarchy
  - Augment Virtual Memory (VM)
  - Dynamic mapping between data and ECC
- Flexible memory protection
  - Single hardware can provide different tolerance levels
  - Allow adaptive tuning of reliability levels
- Enable protection even for Non-ECC DIMMs

# Virtualized and Flexible ECC (Cont'd)



# Write-back a T2EC line when evicted



# Performance Impact of 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 on the critical path

# Chipkill correct

---

- Single Device-error Correct and Double Device-error Detect
  - Can tolerate a DRAM failure
  - Can detect a second DRAM failure
- Traditional chipkill requires x4 DRAMs
- V-ECC x8
  - Two-tiered error code
  - Simpler T1EC relaxes module design constraints
    - Enable more energy efficient x8 configurations
  - T2EC for error correction is virtualized



# Flexible Error Protection

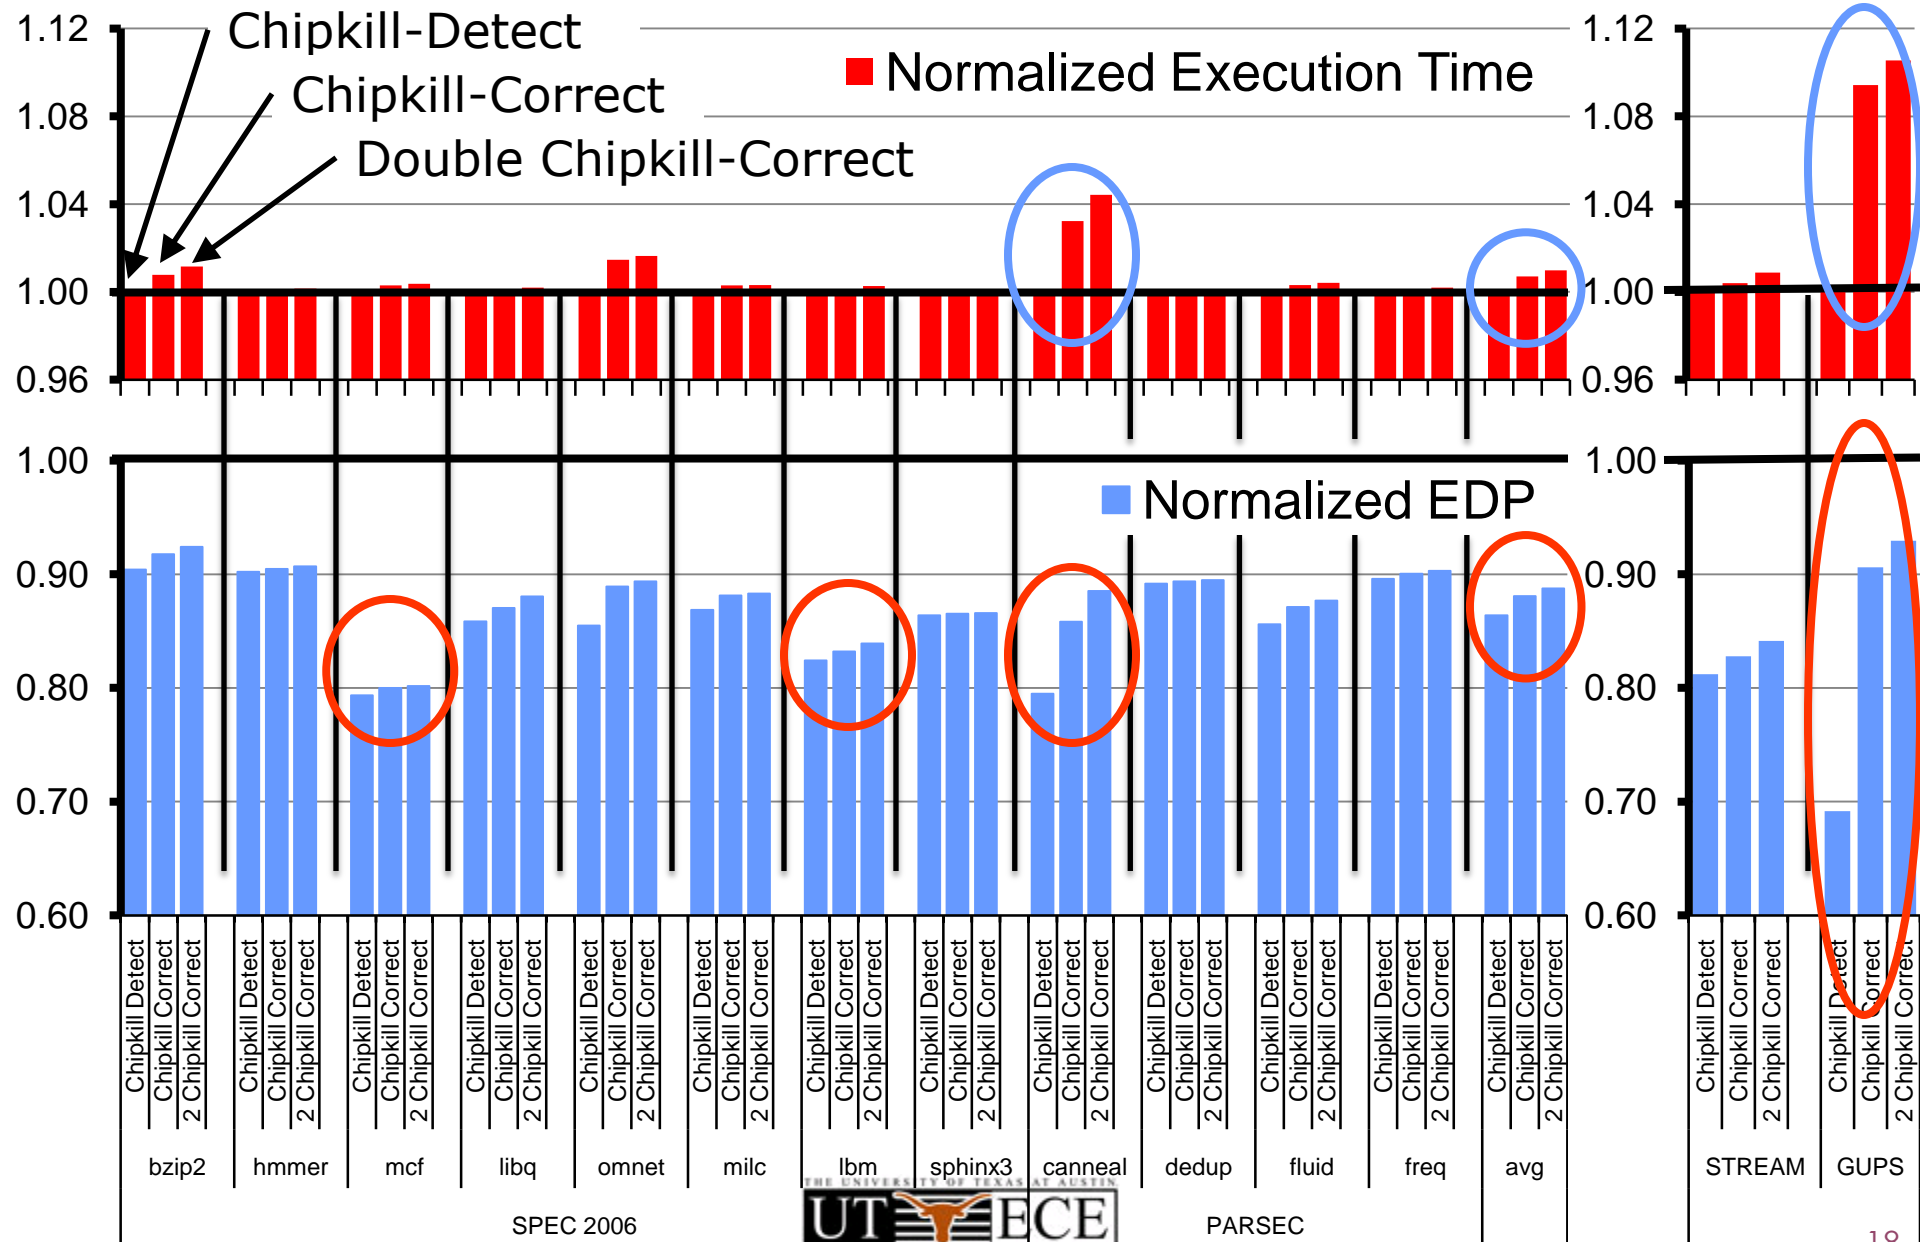
- Single HW with V-ECC can provide
  - Chipkill-detect, Chipkill-correct, and Double chipkill-correct
  - Use different T2EC for different pages

	Chipkill-Detect	Chipkill-Correct	Double Chipkill-Correct
T2EC per 64B	0B	4B	8B

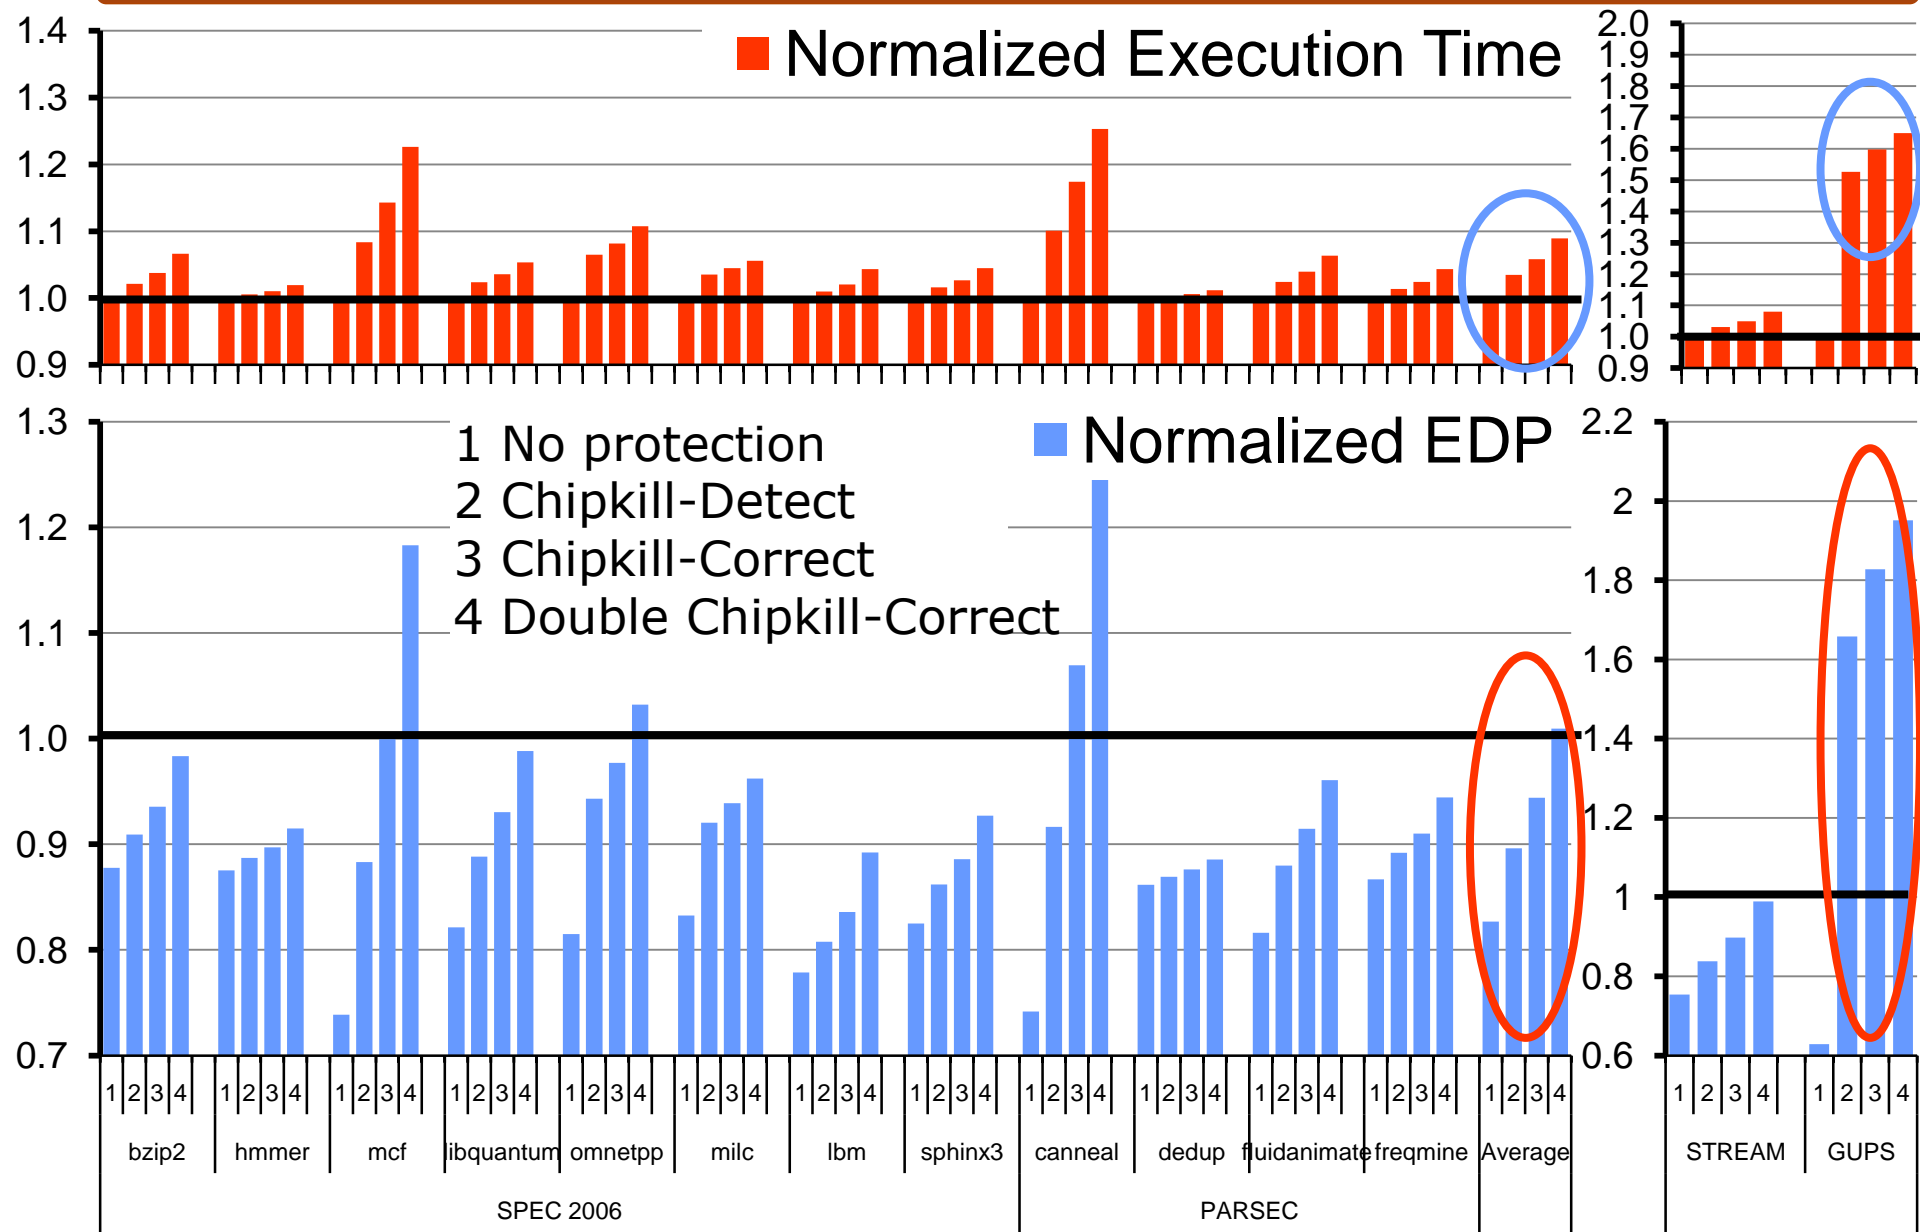
- Maximize performance/power efficiency with Chipkill-Detect
- Stronger protection at the cost of additional T2EC accesses



# VECC x8 (normalized to baseline x4 chipkill)



# V-ECC with Non-ECC DIMM (normalized to baseline x4 chipkill)





---

# Adaptive Granularity Memory System

# Adaptive Granularity Memory System

---


- Bandwidth wall
  - More and more cores/threads per chip
  - Off-chip bandwidth scaling is limited
    - due to pins and power
- Fixed, coarse access granularity
  - 64B or 128B data block with 12.5% redundancy overhead
  - Works well with most applications with spatial locality
- **Inefficient** for apps with low spatial locality
  - Waste off-chip BW for unnecessary data with a block
  - Can utilize off-chip BW better with fine-grained blocks
- Cache hierarchy, DRAM and uniform ECC prevents fine-grained accesses

# Coarse- and Fine-Grained Memory Access

- Coarse-grained access
  - Lower ECC/control overhead
  - Generally works well
  - But, poor throughput without spatial locality




- Fine-grained access
  - Higher ECC/control overhead
  - Better throughput, if apps have poor spatial locality



# Design examples


Wide channel for  
Coarse-grained access



Data

ECC

Narrow multi-channel for  
Fine-grained access



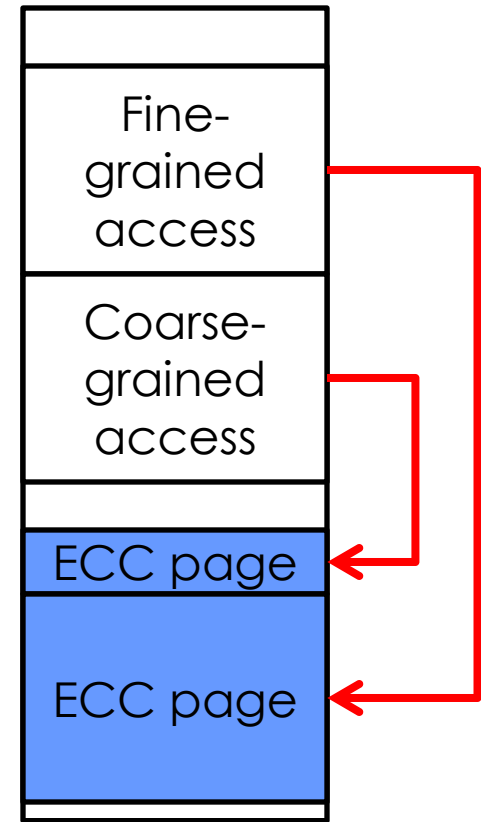
Data

ECC

Data

ECC

AGMS with  
virtualized ECC



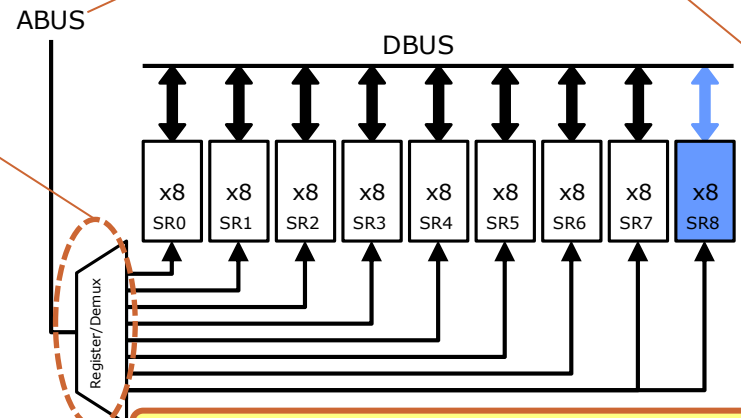
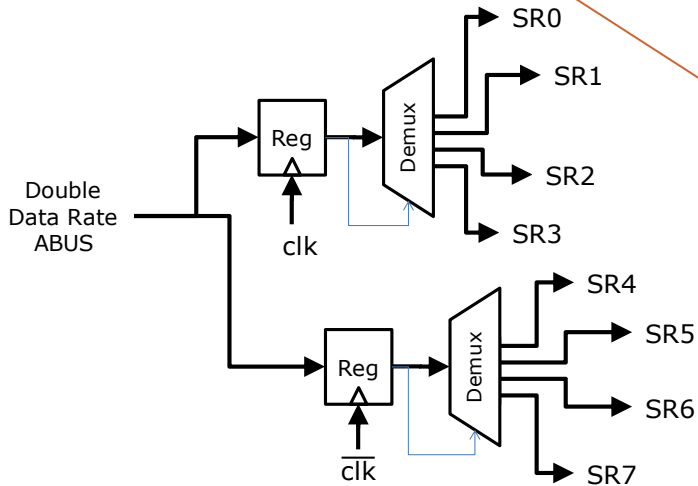
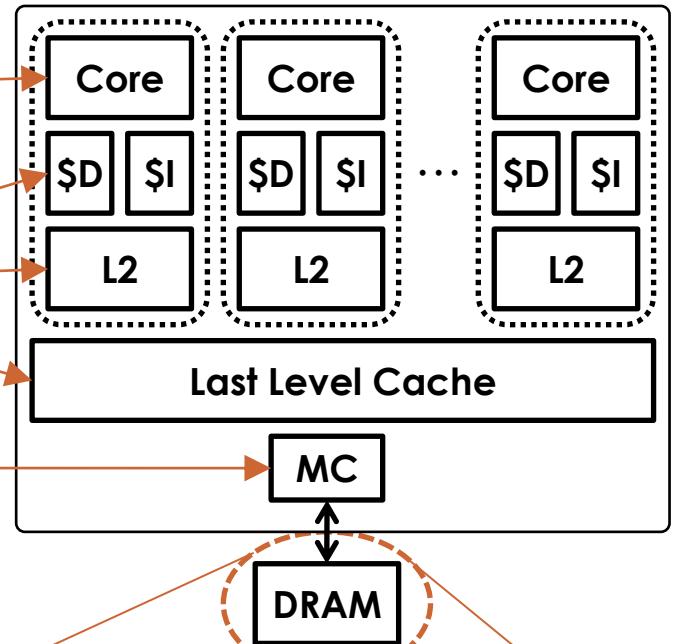
Data

# AGMS Design

Augment Virtual Memory (VM) to identify fine-grained pages  
-- 1 bit per PTE

Sectored caches to manage 8B data blocks

Memory controller can handle multiple access granularities




Higher Address/Command BW

Subranked DRAM to enable an 8B access granularity

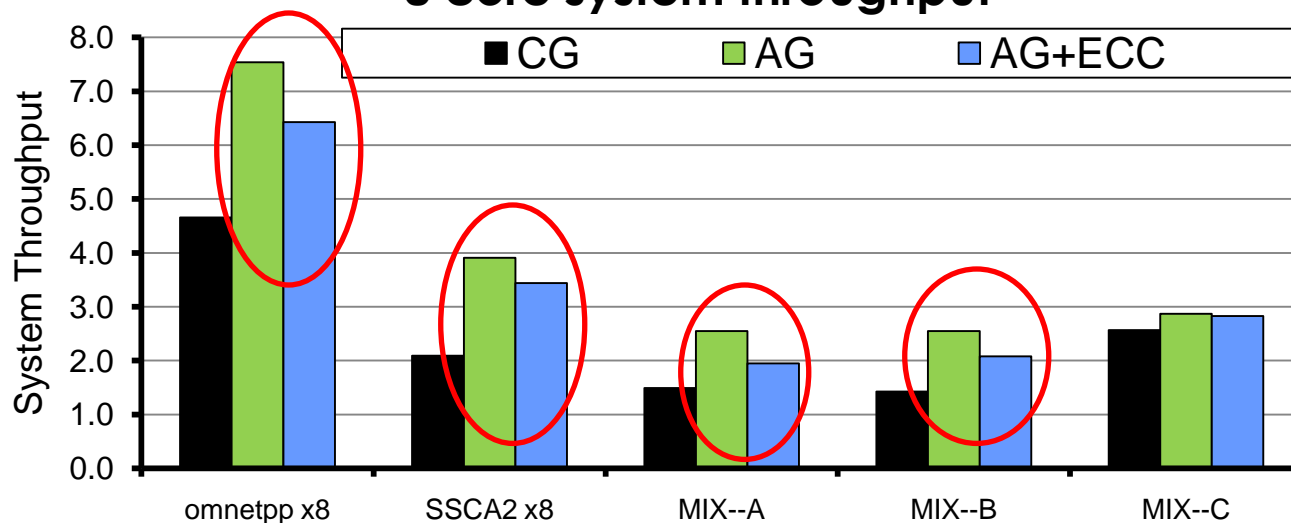


# Preliminary Results

## 4 core system throughput



## 8 core system throughput



# Conclusions and Future Work

---

- Virtualized ECC
  - Store all or part of redundant information within the memory hierarchy
  - Two-tiered protection
  - Minimize dedicated resources to redundancy
  - Flexibility in error tolerance level/access granularity
- Future work
  - GPU memory protection
  - Non-Volatile memory protection
  - Adaptive/tunable reliability
  - Generic meta-data storage