XED: EXPOSING ON-DIE ERROR DETECTION INFORMATION FOR STRONG MEMORY RELIABILITY

Prashant Nair, Georgia Tech Vilas Sridharan*, AMD Inc.* Moinuddin Qureshi, Georgia Tech

ISCA-43, June 20th 2016 Seoul, Republic of Korea





INTRODUCTION

DRAM Scaling → High Capacity Memories Two types of DRAM faults

INTRODUCTION

DRAM Scaling → High Capacity Memories Two types of DRAM faults



INTRODUCTION

DRAM Scaling → High Capacity Memories Two types of DRAM faults



Runtime Faults					
Fault Mode	Transient Fault Rate (FIT)	Permanent Fault Rate (FIT)			
Bit	14.2	18.6			
Word	1.4	0.3			
Column	1.4	5.6			
Row	0.2	8.2			
Bank	0.8	10			
*Total	18	42.7			
Sridharan et. al. SC13					

DRAM vendors plan to use "On-Die ECC"

- Mitigates scaling faults transparently
- Enables good DIMM with bad chips (yield)
- Part of: LPDDR4, DDR4, DDR5 (proposed)















On-Die ECC: Single Error Correction, Double Error Detection Code (SECDED)





On-Die ECC fixes scaling faults invisibly

Runtime faults

Fault	Transient	Permanent		
Mode	Fault Rate (FIT)	Fault Rate (FIT)		
Bit	14.2	18.6		

ECC-DIMM (9-Chips)

CHIP	ECC Chip								
					•				

Runtime faults

- Chip faults common
- Need strong ECC

Fault Mode	Transient Fault Rate (FIT)	Permanent Fault Rate (FIT)
Bit	14.2	18.6
Word	1.4	0.3
Column	1.4	5.6
Row	0.2	8.2
Bank	0.8	10
*Total	18	42.7



Runtime chip faults → Chipkill (strong ECC)



Runtime chip faults → Chipkill (strong ECC)

18 DRAM Chips



Runtime chip faults \rightarrow Chipkill (strong ECC)



Runtime chip faults \rightarrow Chipkill (strong ECC)

18 DRAM Chips



GOAL AND CHALLENGE

<u>GOAL</u>: Use On-Die ECC to mitigate runtime faults "Chipkill-level reliability using x8 ECC-DIMM"

GOAL AND CHALLENGE

<u>GOAL</u>: Use On-Die ECC to mitigate runtime faults "Chipkill-level reliability using x8 ECC-DIMM"

<u>CHALLENGE</u>: On-Die ECC is invisible, expose it without changing the memory interface

OUTLINE

- BACKGROUND
- XED
- CASE STUDIES
- EVALUATION
- SUMMARY

What if the chip can inform that it failed?



What if the chip can inform that it failed?



What if the chip can inform that it failed?



Parity + Location \rightarrow Reconstruct Data for Faulty Chip

What if the chip can inform that it failed?



Parity + Location \rightarrow Reconstruct Data for Faulty Chip

Fix chip-faults using only 9 Chips

XED: EXPOSED ON-DIE ERROR DETECTION

XED consists of three components

- Strong detection in addition to SEC
- Parity-based correction
- Transparently identifying faulty chip

XED: ON-DIE ECC AS DETECTION CODE



XED: ON-DIE ECC AS DETECTION CODE

















EXPOSE ON-DIE ERROR INFO

OPTION 1: Use additional wires



EXPOSE ON-DIE ERROR INFO

OPTION 1: Use additional wires


OPTION 1: Use additional wires







Incompatible with DDR memory standards

Needs a new protocol

Worse for pin-constrained future systems!

Memory Controller

OPTION 2: Use additional burst/transaction



OPTION 2: Use additional burst/transaction



OPTION 2: Use additional burst/transaction



OPTION 2: Use additional burst/transaction



Expose On-Die error detection with minor changes

XED: ON-DIE ERROR INFO FOR FREE

On detecting an error, the DRAM chip sends a 64bit "Catch-Word" (CW) instead of data



Memory Controller

XED: ON-DIE ERROR INFO FOR FREE

On detecting an error, the DRAM chip sends a 64bit "Catch-Word" (CW) instead of data



XED: MUX TO SEND CATCH-WORDS



Simple MUX to chose between Data and Catch-Word

XED: ON-DIE ERROR INFO FOR FREE

On detecting an error, the DRAM chip sends a 64bit "Catch-Word" (CW) instead of data

Chips provisioned with a unique Catch-Word

No additional wires/bandwidth overheads

Compatible with existing memory protocols

Memory Controller

64-bit Catch-Words identify the faulty chip

Catch Word (CW) ≠ Valid Data (D2)



Catch Word (CW) \neq Valid Data (D2) Then \rightarrow PA \neq D0 \bigoplus D1 \bigoplus CW \bigoplus ... \bigoplus D7



Catch Word (CW) \neq Valid Data (D2) Then \rightarrow PA \neq D0 \bigoplus D1 \bigoplus CW \bigoplus ... \bigoplus D7



Catch Word (CW) = Valid Data (D2)



Catch Word (CW) = Valid Data (D2) [*Collision*] Then \rightarrow PA = D0 \bigoplus D1 \bigoplus CW \bigoplus ... \bigoplus D7



Catch Word (CW) = Valid Data (D2) [*Collision*] Then \rightarrow PA = D0 \bigoplus D1 \bigoplus CW \bigoplus ... \bigoplus D7



Catch-Word collision: Doesn't affect correctness

• A chip stores 64 bits/cache-line \rightarrow 2⁶⁴ combinations

- A chip stores 64 bits/cache-line \rightarrow 2⁶⁴ combinations
- However even a 16Gb chip has only 2²⁸ cachelines

- A chip stores 64 bits/cache-line \rightarrow 2⁶⁴ combinations
- However even a 16Gb chip has only 2²⁸ cachelines
- Even if this entire chip contained different data there are nearly 2^{63.99} data combinations free!

- A chip stores 64 bits/cache-line \rightarrow 2⁶⁴ combinations
- However even a 16Gb chip has only 2²⁸ cachelines
- Even if this entire chip contained different data there are nearly 2^{63.99} data combinations free!



- A chip stores 64 bits/cache-line \rightarrow 2⁶⁴ combinations
- However even a 16Gb chip has only 2²⁸ cachelines
- Even if this entire chip contained different data there are nearly 2^{63.99} data combinations free!



The catch-word will most likely not collide

OUTLINE

- BACKGROUND
- XED
- CASE STUDIES
- EVALUATION
- SUMMARY

XED FOR SCALING ERRORS

On-Die ECC

- Single Error Correction
- Always detects scaling errors (single-bit)

CASE STUDY 1: SINGLE SCALING FAULT

Scaling fault within a single chip



Parity reconstructs data from chip with scaling error

CASE STUDY 2: MULTIPLE SCALING FAULTS

Scaling faults within multiple chips



Disable XED + Retry

CASE STUDY 3: CHIP FAULT

Catch-Word identifies the faulty chip



Parity reconstructs data from failed chip

CASE STUDY 4: CHIP + SCALING FAULT

Parity detects error even after retry \rightarrow Chip Failure



Disable XED + Diagnosis to locate chip failure

OUTLINE

- BACKGROUND
- XED
- CASE STUDIES
- EVALUATION



• SUMMARY

USIMM : 8 Cores, 4 Channels, 2 Ranks, 8 Banks

FaultSim*: Memory Reliability Simulator

- Real World Fault Data
- 7 year system lifetime,
- Billion Monte-Carlo Trails
- Metric: Probability of System Failure
- Scaling Fault-Rate: 10⁻⁴

RESULTS: RELIABILITY

XED vs Commercial ECC schemes



RESULTS: RELIABILITY

XED vs Commercial ECC schemes



XED provides strong reliability while using fewer chips



Lower the better



Lower the better



Execution time: 21% J, EDP : 34% J



Execution time: 21% J, EDP : 34% J

OUTLINE

- BACKGROUND
- XED
- CASE STUDIES
- EVALUATION
- SUMMARY


- DRAM Scaling introduces errors \rightarrow On-Die ECC
- On-Die ECC is invisible to the memory system

- DRAM Scaling introduces errors \rightarrow On-Die ECC
- On-Die ECC is invisible to the memory system
- Exposing On-Die ECC: Efficient Runtime ECC

- DRAM Scaling introduces errors \rightarrow On-Die ECC
- On-Die ECC is invisible to the memory system
- Exposing On-Die ECC: Efficient Runtime ECC
- XED
 - Exposes On-Die Error Detection using Catch-Words

- DRAM Scaling introduces errors \rightarrow On-Die ECC
- On-Die ECC is invisible to the memory system
- Exposing On-Die ECC: Efficient Runtime ECC
- XED
 - Exposes On-Die Error Detection using Catch-Words
 - 2X fewer chips as compared to Chipkill

- DRAM Scaling introduces errors \rightarrow On-Die ECC
- On-Die ECC is invisible to the memory system
- Exposing On-Die ECC: Efficient Runtime ECC
- XED
 - Exposes On-Die Error Detection using Catch-Words
 - 2X fewer chips as compared to Chipkill
 - 4X higher reliability as compared to Chipkill

- DRAM Scaling introduces errors \rightarrow On-Die ECC
- On-Die ECC is invisible to the memory system
- Exposing On-Die ECC: Efficient Runtime ECC
- XED
 - Exposes On-Die Error Detection using Catch-Words
 - 2X fewer chips as compared to Chipkill
 - 4X higher reliability as compared to Chipkill
 - 21% lower execution time as compared to Chipkill

- DRAM Scaling introduces errors \rightarrow On-Die ECC
- On-Die ECC is invisible to the memory system
- Exposing On-Die ECC: Efficient Runtime ECC
- XED
 - Exposes On-Die Error Detection using Catch-Words
 - 2X fewer chips as compared to Chipkill
 - 4X higher reliability as compared to Chipkill
 - 21% lower execution time as compared to Chipkill
- XED \rightarrow No change in memory protocols

THANK YOU



"You are in a pitiable condition, if you have to conceal what you wish to tell" - Publilius Syrus

THANK YOU



"You are in a pitiable condition, if you have to conceal what you wish to tell" - Publilius Syrus

BACKUP

2-Chip Failures



Chipkill (18-chips)







2-Chip Failures



Chipkill







2-Chip Failures



Chipkill







2-Chip Failures \rightarrow Extend to Multi-Chip Failures



Chipkill







SDC AND DUE

SDC AND DUE RATE OF XED

Source of Vulnerability	Rate over 7 years
XED: Scaling-Related Faults	No SDC or DUE
XED: Row/ Column/ Bank Failure	1.4×10^{-13} (SDC)
XED: Word Failure	6.1×10^{-6} (DUE)
Data Loss from Multi-Chip Failures	5.8×10^{-4}

ADDITIONAL BURST/TRANSACTION



XED VS LOT-ECC



SPEC PARSEC BIOBENCH COMM GMEAN