ISCA 2021

Dvé: Improving DRAM Reliability and Performance On-Demand via Coherent Replication

Adarsh Patil,
Vijay Nagarajan, Rajeev Balasubramonian, Nicolai Oswald

THE UNIVERSITY *of* EDINBURGH
informatics

arm

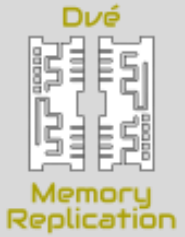

Outline

- Reliability benefits
- Performance gains
- On-demand Reliability

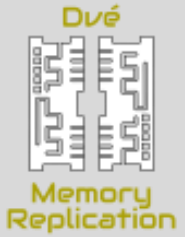

Outline

- Reliability benefits
- Performance gains
- On-demand Reliability

Dvé Memory Replication

R

Increasing DRAM Faults

Bloomberg

Markets

How One Piece of Hardware Took Down a $6 Trillion Stock Market

By Gearoid Reidy, Shoko Oda, Min Jeong Lee, and Toshiro Hasegawa
2 October 2020, 10:47 BST *Updated on 5 October 2020, 01:48 BST*

That all changed on Thursday, when a piece of hardware called the No. 1 shared disk device, one of two square-shaped data-storage boxes, detected a memory error. These devices store management data used across the servers, and distribute information such as commands and ID and password combinations for terminals that monitor trades.

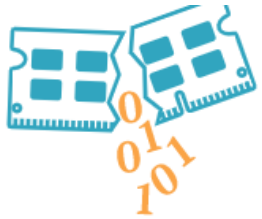

RAMBleed

Reading Bits in Memory Without Accessing Them

RAMBleed is a side-channel attack that enables an attacker to read out physical memory belonging to other processes. The implications of violating

ECCPLOIT: ECC MEMORY VULNERABLE TO ROWHAMMER ATTACKS AFTER ALL

Where many people thought that high-end servers were safe from the (unpatchable) Rowhammer bitflip

IEEE SPECTRUM

Hardware

DRAM's Damning Defects—and How They Cripple Computers

By **Ioan Stefanovici, Andy Hwang and Bianca Schroeder**
Posted 23 Nov 2015 | 16:00 GMT

COMPUTERWORLD UNITED KINGDOM

NEWS

Google: DRAM error rates vastly higher than previously thought

PCs will likely require error correction code in the future due to DRAM issues

By Lucas Mearian
Senior Reporter, Computerworld | 8 OCTOBER 2009 23:51 GMT

ZDNet MENU UK

DRAM error rates: Nightmare on DIMM street

A two-and-a-half year study of DRAM on 10s of thousands Google servers found DIMM error rates are hundreds to thousands of times higher than thought -- a mean of 3,751 correctable errors per DIMM per year.This is the world's first large-scale study of RAM errors in the field.

By Robin Harris for Storage Bits | October 4, 2009 -- 22:04 GMT (23:04 BST) | Topic: Hardware

ZDNet MENU UK

DRAM errors: from soft to hard

Every system uses dynamic random access memory (DRAM), but how good is it? Bad news: not nearly as good as vendors would like us to think. Good news: we're learning.

By Robin Harris for Storage Bits | October 24, 2012 -- 16:26 GMT (17:26 BST) | Topic: Storage

Dvé
Memory Replication

R

Progression of Reliability Mechanisms

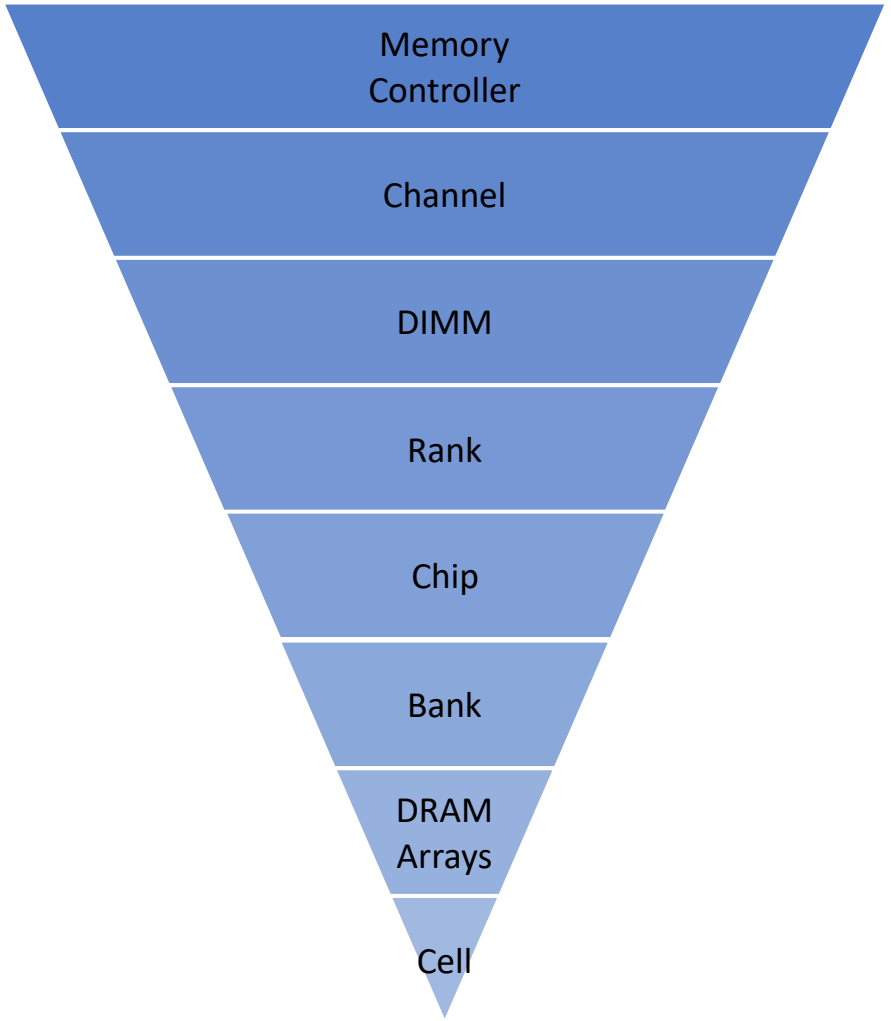

Dvé
Memory Replication

R

Progression of Reliability Mechanisms

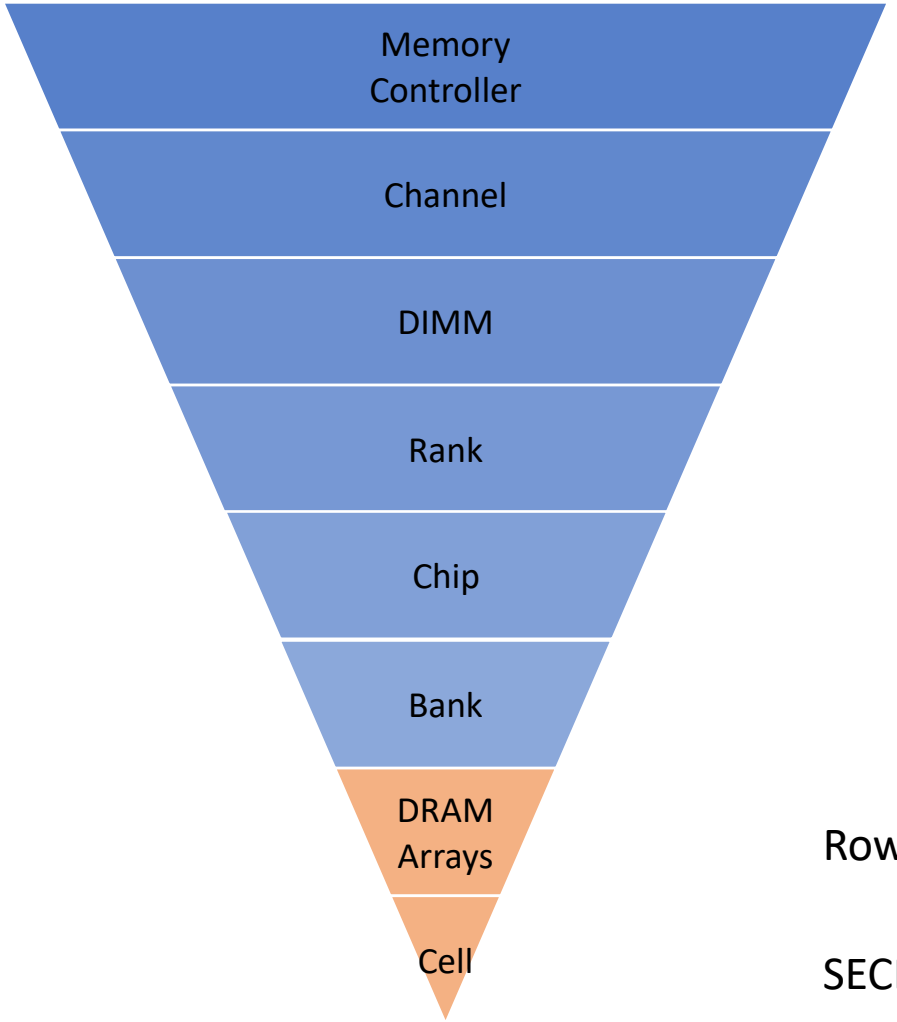

Row/column sparing

SECDED ECC, in-DRAM ECC

Dvé
Memory Replication

R

Progression of Reliability Mechanisms

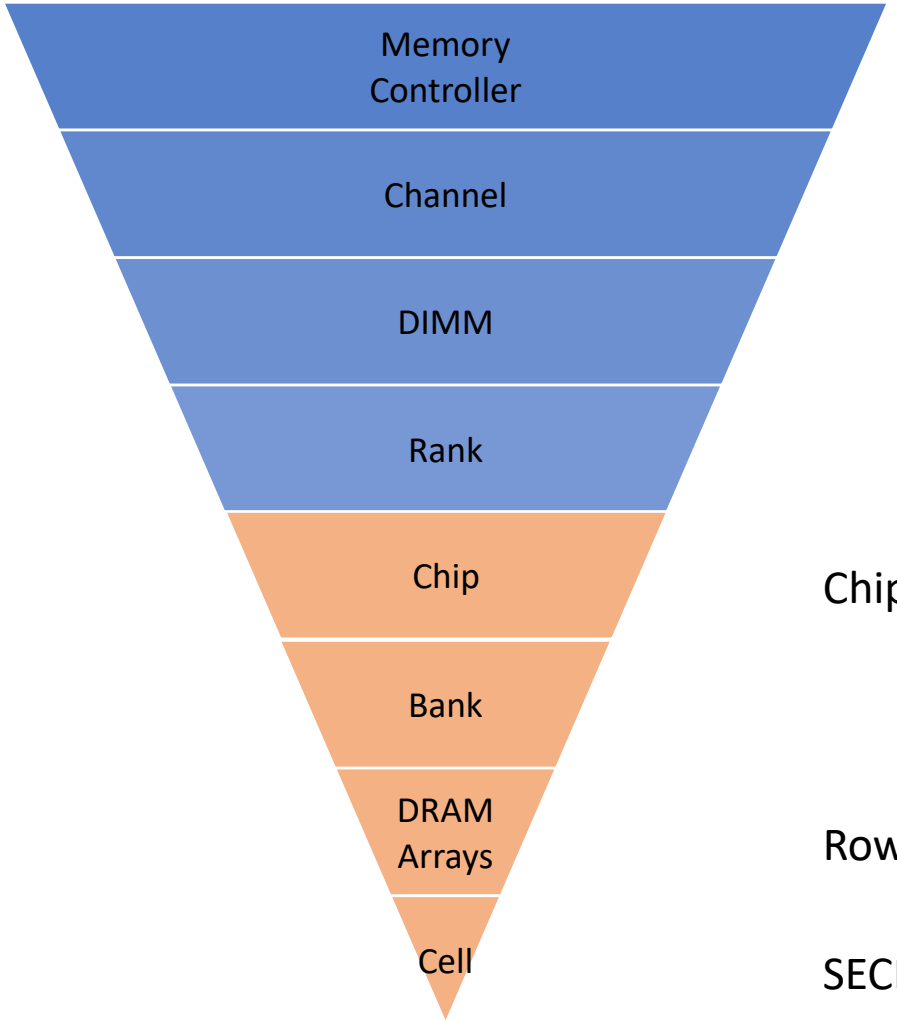

Chipkill ECC / Multi-tier ECC

Row/column sparing

SECDED ECC, in-DRAM ECC

Dvé
Memory Replication

R

Progression of Reliability Mechanisms

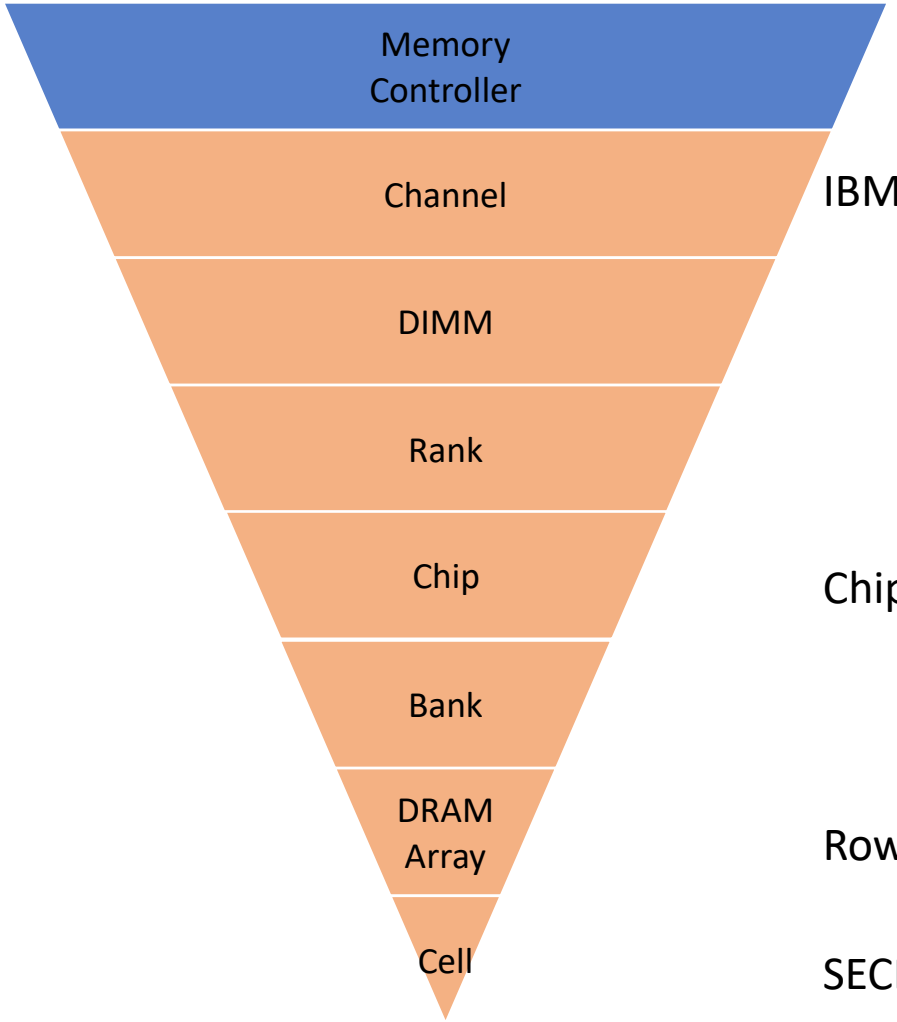

IBM RAIM, Intel Memory Mirroring

Chipkill ECC / Multi-tier ECC

Row/column sparing

SECDED ECC, in-DRAM ECC

Dvé
Memory Replication

R

Progression of Reliability Mechanisms

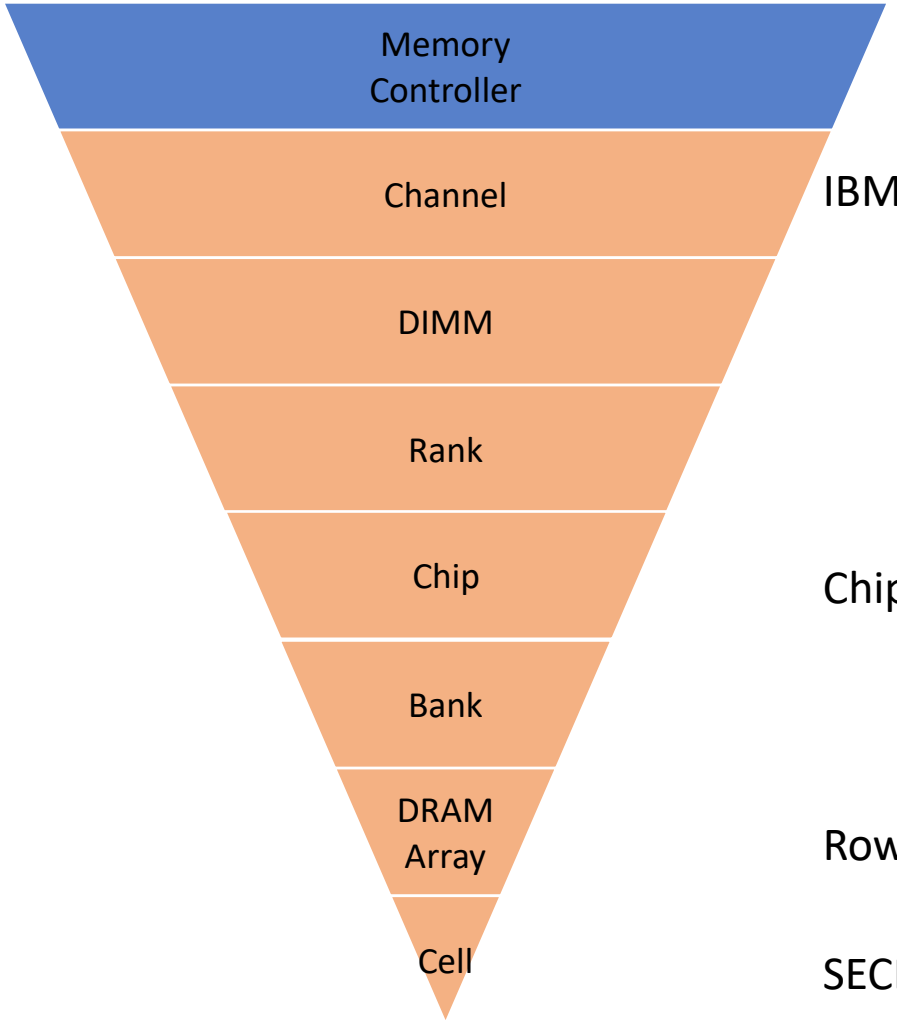

IBM RAIM, Intel Memory Mirroring

Chipkill ECC / Multi-tier ECC

Row/column sparing

SECDED ECC, in-DRAM ECC

Performance overheads

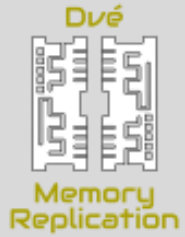

Replication for Reliability

Dvé insights

- ❑ Full data replica (not ECC code)
- ❑ Keep Replicas as far apart and disjoint as possible
- ❑ Tolerate errors arising from anywhere in the memory path

Dvé
Memory Replication

R

Replication for Reliability

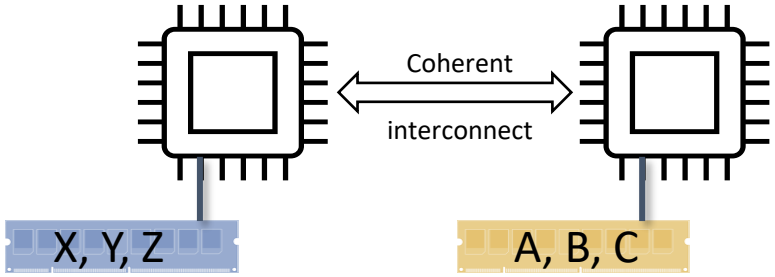

Dvé insights

- ❑ Full data replica (not ECC code)
- ❑ Keep Replicas as far apart and disjoint as possible
- ❑ Tolerate errors arising from anywhere in the memory path

Dvé
Memory Replication

R

Replication for Reliability

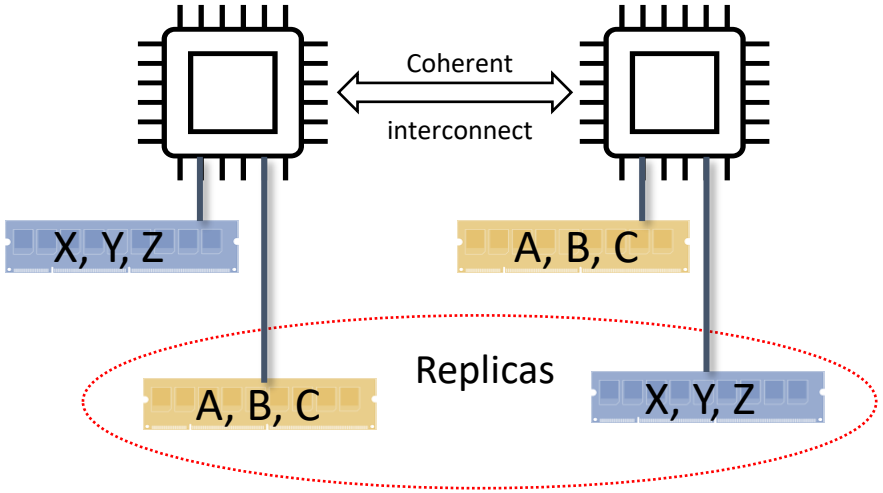

Dvé insights

- ❑ Full data replica (not ECC code)
- ❑ Keep Replicas as far apart and disjoint as possible
- ❑ Tolerate errors arising from anywhere in the memory path

Dvé
Memory Replication

R

Replication for Reliability

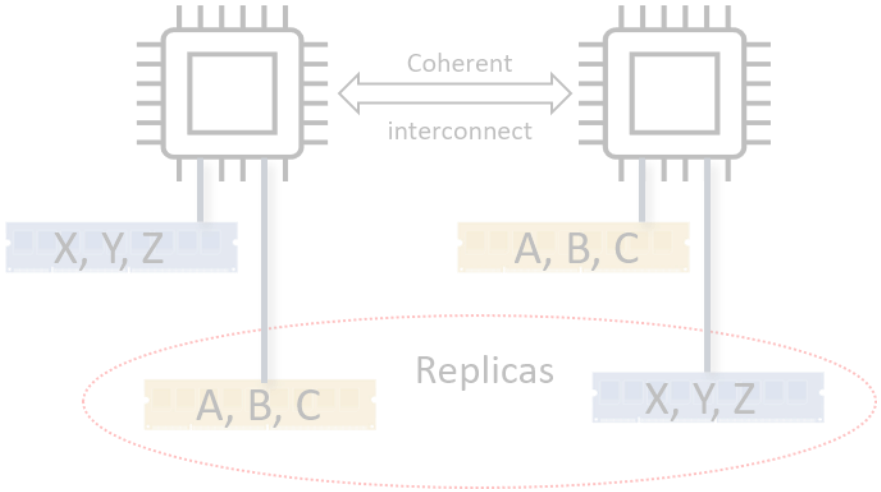

Dvé insights

- ❑ Full data replica (not ECC code)
- ❑ Keep Replicas as far apart and disjoint as possible
- ❑ Tolerate errors arising from anywhere in the memory path

For Detection

- ❑ Existing ECC, CRC, Parity
- ❑ Strong detection-only code
- ❑ Other diagnostic capabilities

For Correction

- ❑ Rely on replica

Dvé
Memory Replication

R

Replication for Reliability

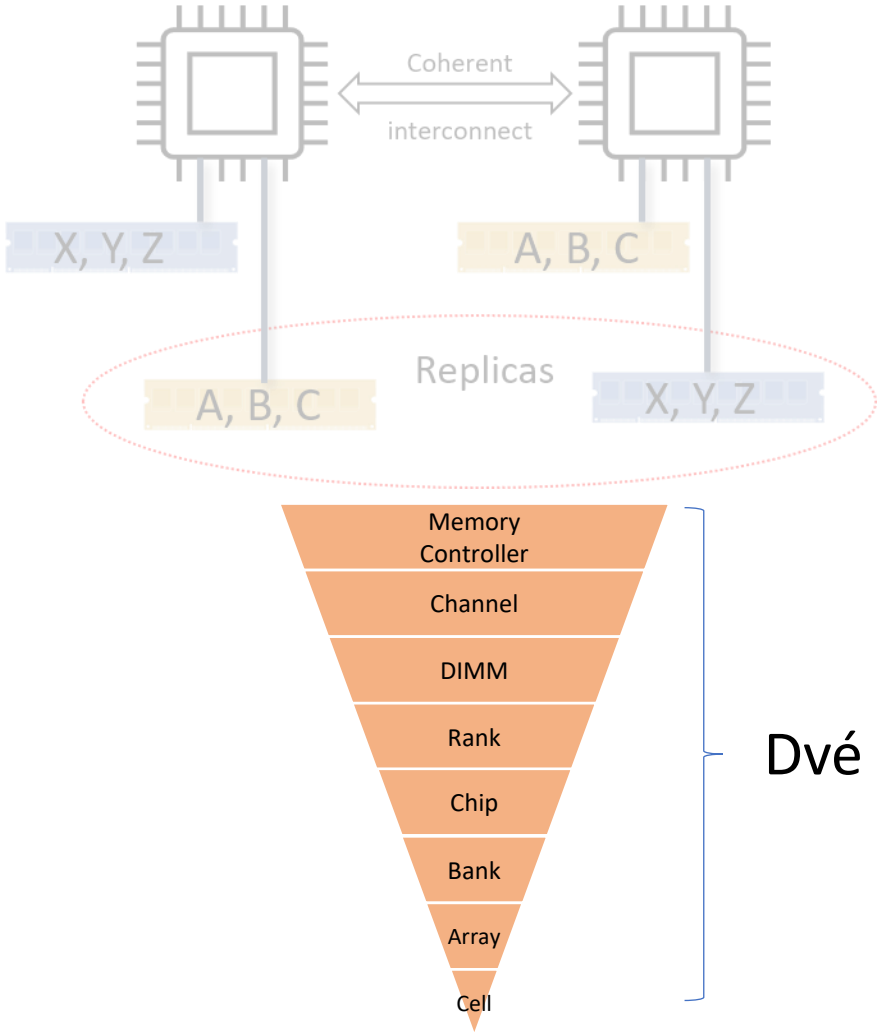

Dvé insights

- ❑ Full data replica (not ECC code)
- ❑ Keep Replicas as far apart and disjoint as possible
- ❑ Tolerate errors arising from anywhere in the memory path

For Detection

- ❑ Existing ECC, CRC, Parity
- ❑ Strong detection-only code
- ❑ Other diagnostic capabilities

For Correction

- ❑ Rely on replica

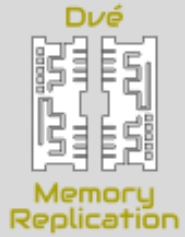

R Quantifying Reliability *(Onus Probandi)*

Analytical Modelling

- Device FIT rate: 66.1 [Sridharan et. al., SC '12]
- Error rates: DUE and SDC
- Equipped with same detection scheme

Comparison Points

Chipkill: (SSC-DSD ECC code)	guarantees recovery from 1 DRAM chip failure in a rank
IBM RAIM: (RAID-3 across 5 channels)	guarantees recovery from 1 channel failure
Intel Memory Mirroring: (channel-level replication)	guarantees recovery from 1 channel failure

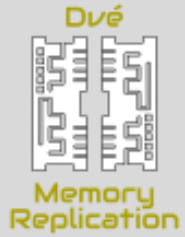

Quantifying Reliability *(Onus Probandi)*

Key Results

Comparison against	DUE Rate Improvement	SDC Rate Improvement
Chipkill (Dvé equipped with TSD)	4x	~10^6x
IBM RAIM (Dvé equipped with Chipkill)	172x	0.63x
Intel Mirroring (Dvé equipped with TSD + temperature scaled FIT rate)	11%	1x

Dvé Memory Replication

R

Quantifying Reliability *(Onus Probandi)*

Intuition

- "k-out-of-n" model systems vs "parallel n" model system
- Bottom-up vs Top-down design
- Lower bound analysis

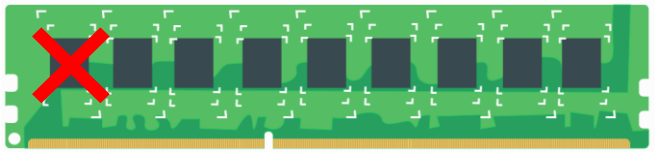

Chipkill ECC

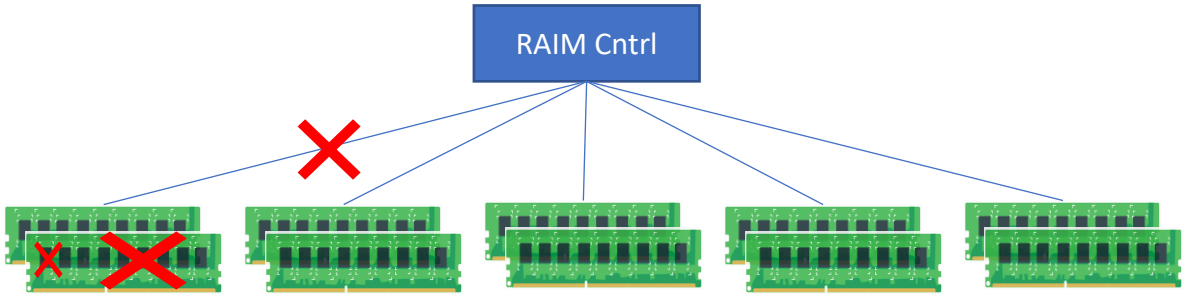

IBM RAIM

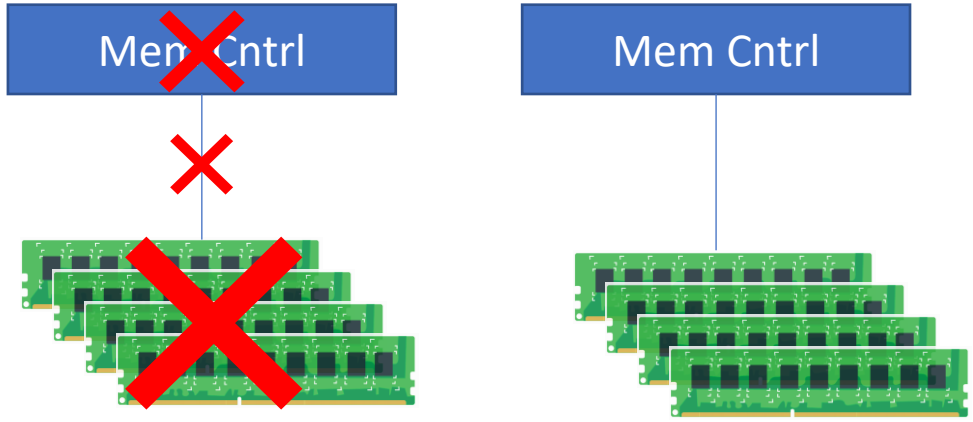

Dvé

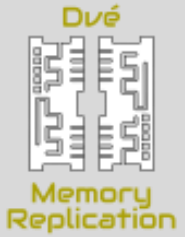

Outline

- Reliability benefits
- Performance gains
- On-demand Reliability

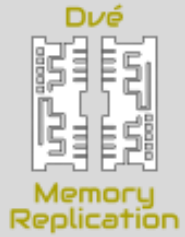

P

Coherent Replication

Dvé insights

- ❑ Use replica to improve performance

Dvé
Memory Replication

P

Coherent Replication

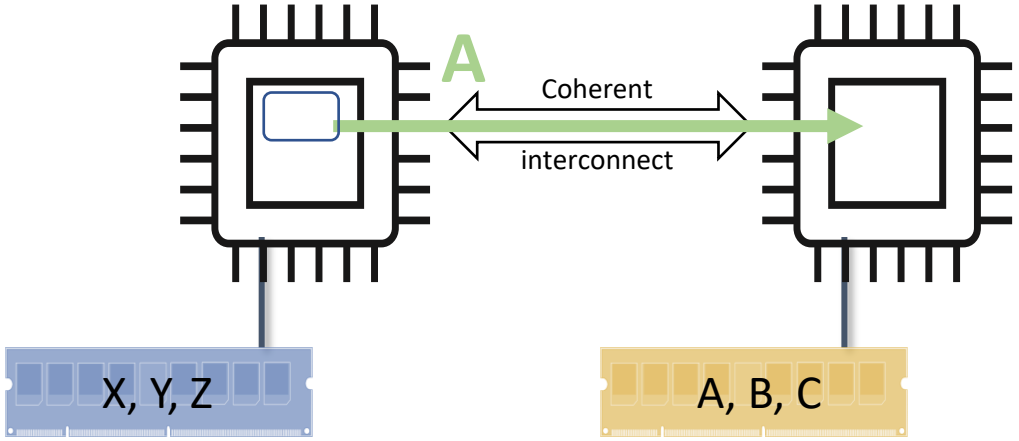

Dvé insights

- ❑ Use replica to improve performance

Dvé
Memory Replication

P

Coherent Replication

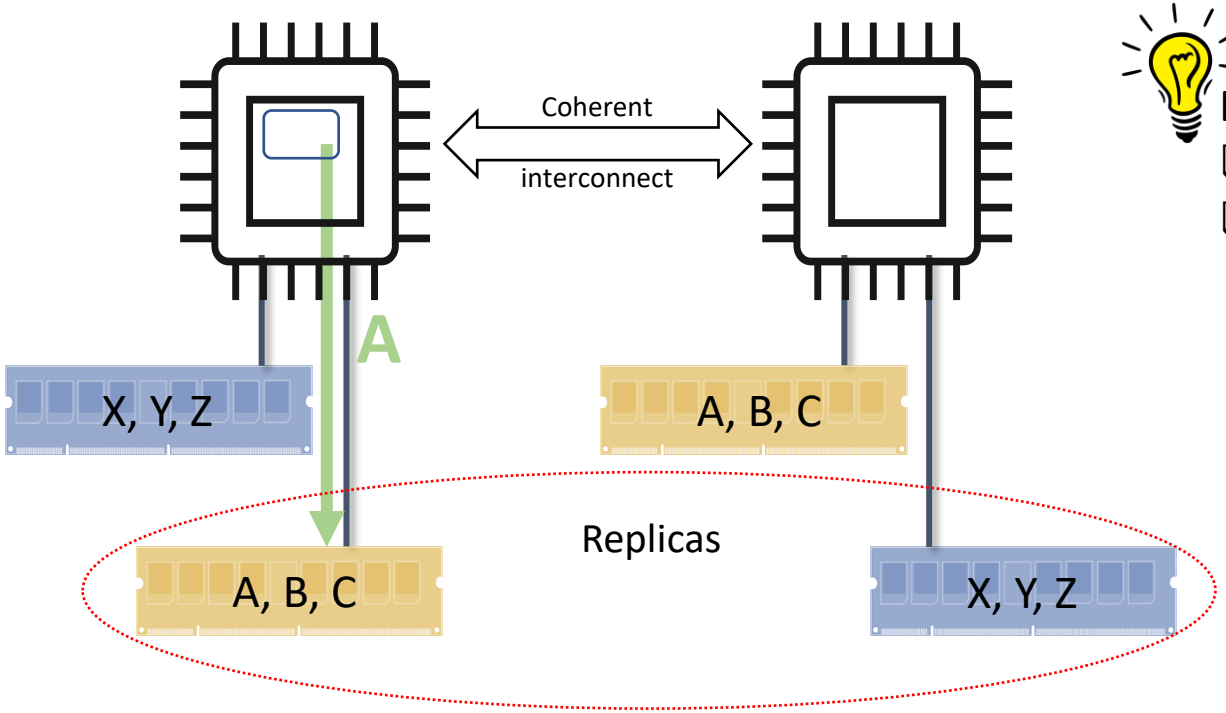

Dvé insights

- ❑ Use replica to improve performance
- ❑ Route memory requests to nearest replica

Dvé Memory Replication

P

Coherent Replication

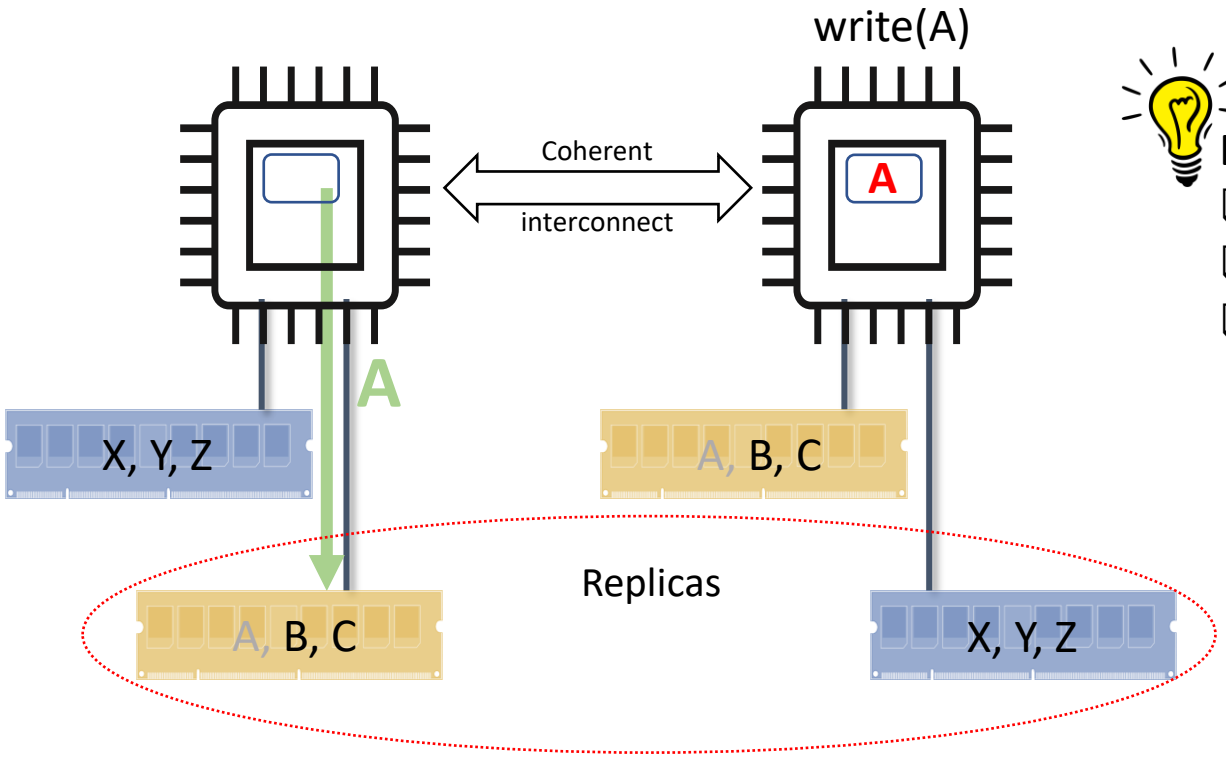

Dvé insights

- ❑ Use replica to improve performance
- ❑ Route memory requests to nearest replica
- ❑ Ensure safe access to replica

Dvé
Memory Replication

P

Coherent Replication

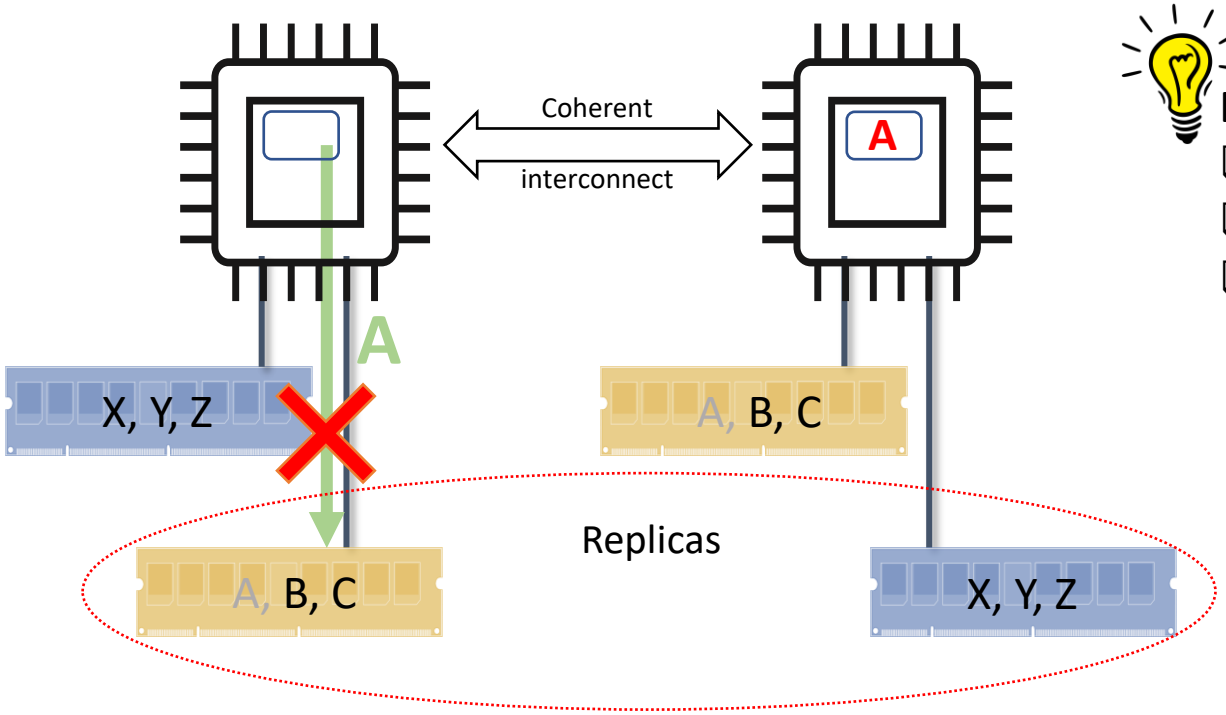

Dvé insights

- ❑ Use replica to improve performance
- ❑ Route memory requests to nearest replica
- ❑ Ensure safe access to replica

Dvé
Memory Replication

P

Coherent Replication

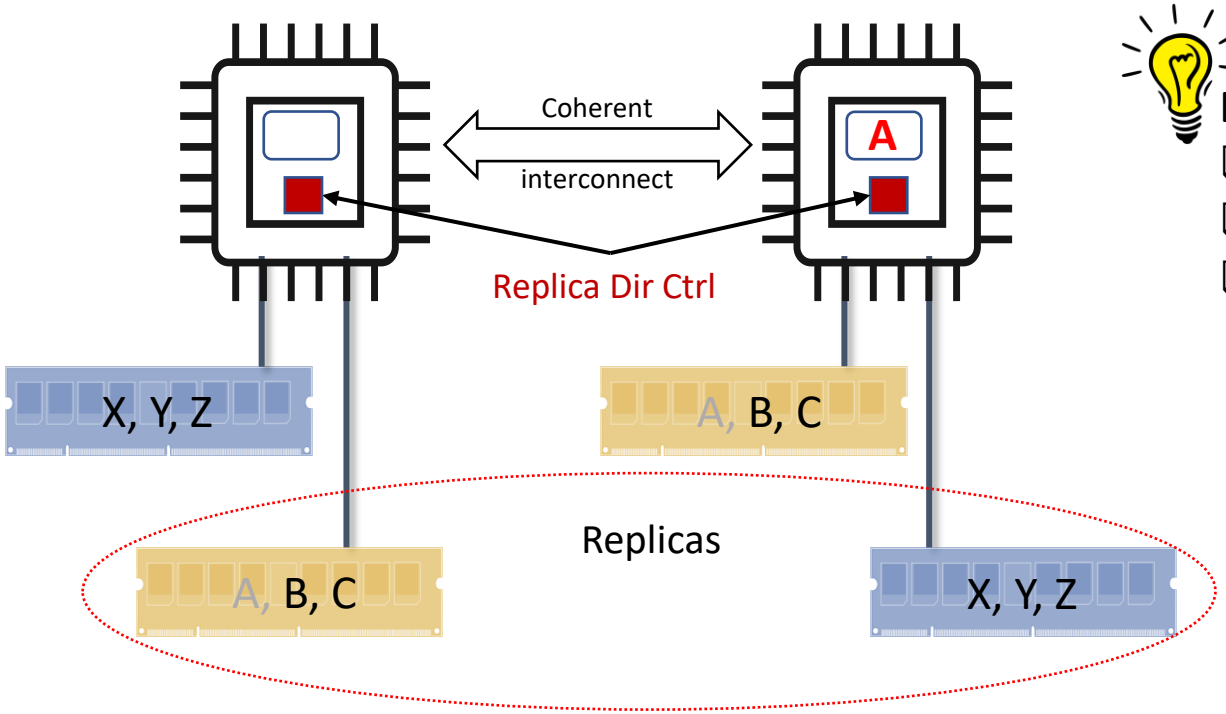

Dvé insights

- ❑ Use replica to improve performance
- ❑ Route memory requests to nearest replica
- ❑ Ensure safe access to replica

Dvé
Memory Replication

P

Coherent Replication

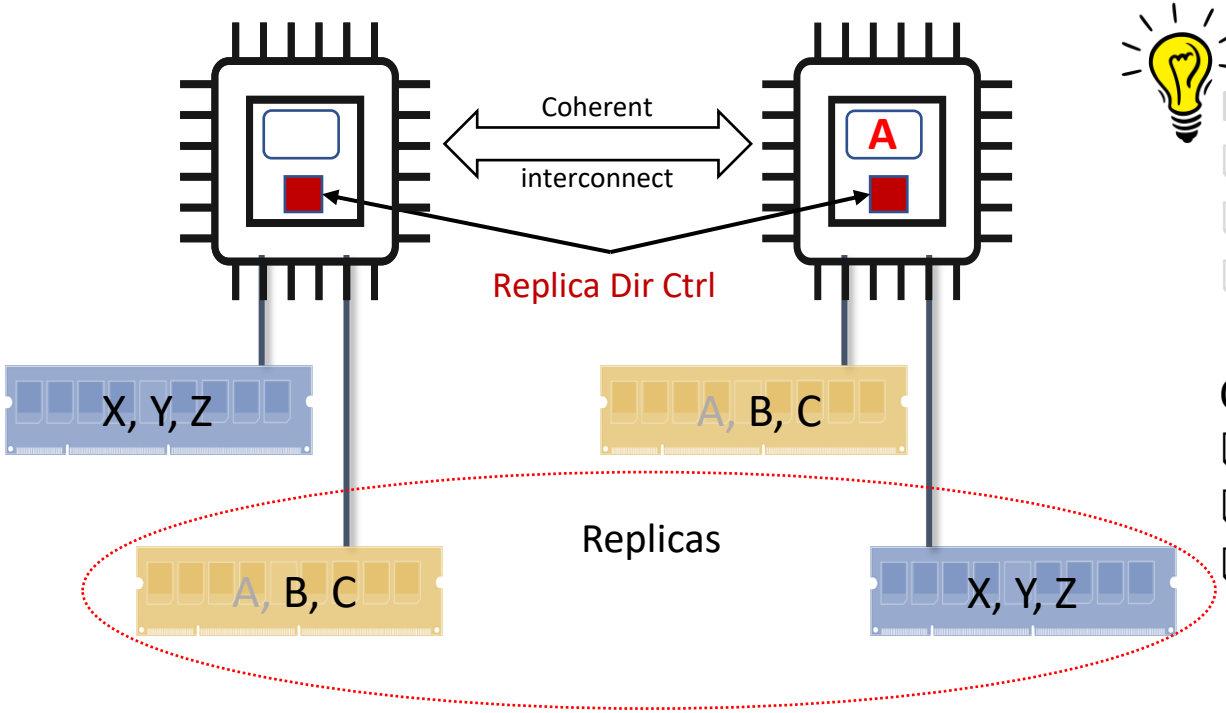

Dvé insights

- ❑ Use replica to improve performance
- ❑ Route memory requests to nearest replica
- ❑ Ensure safe access to replica

Coherent Replication

- ❑ Builds on existing cache coherence protocols
- ❑ maintain the replicas in sync (for reliability)
- ❑ provide coherent access to both replicas during fault-free operation (for performance)

Dvé
Memory Replication

P

Coherent Replication

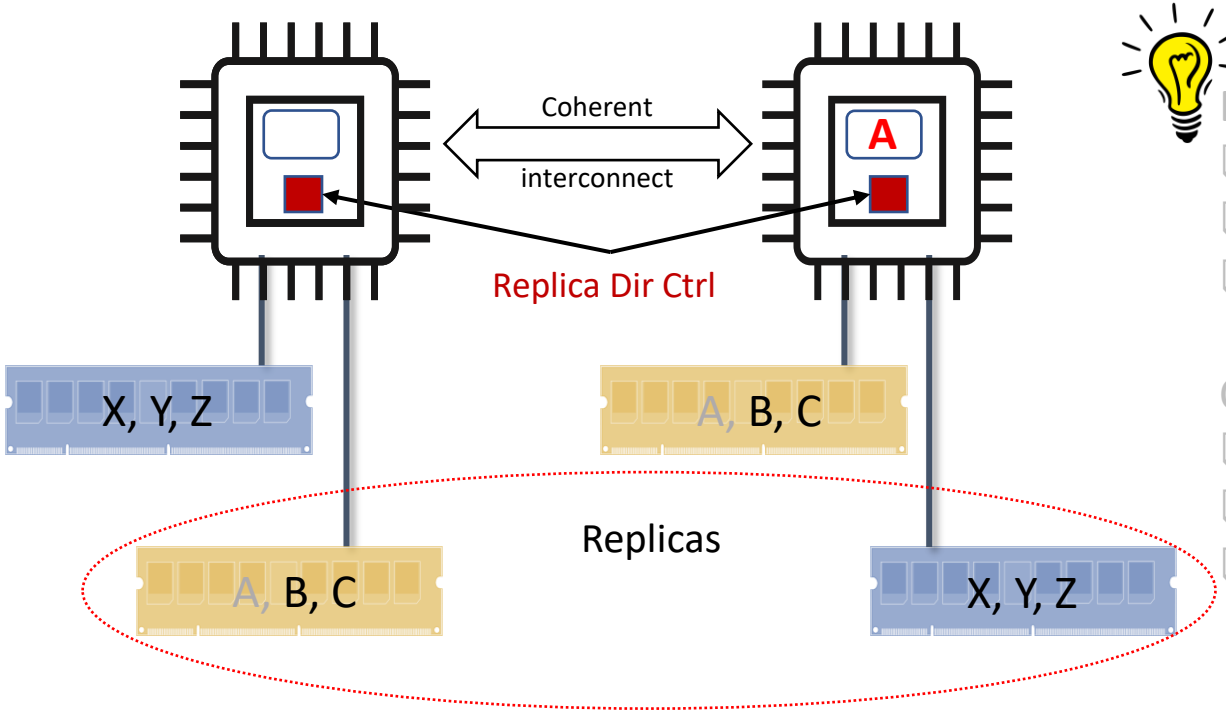

Dvé insights

- ❑ Use replica to improve performance
- ❑ Route memory requests to nearest replica
- ❑ Ensure safe access to replica

Coherent Replication

- ❑ Builds on existing cache coherence protocols
- ❑ maintain the replicas in sync (for reliability)
- ❑ provide coherent access to both replicas during fault-free operation (for performance)

Mechanisms

- ❑ Allow-based
- ❑ Deny-based

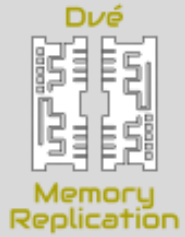

Performance Eval *(Onus Probandi)*

Simulation

- SynchroTrace driven gem5 [TACO 2018]
- Processor: 2-socket, 8 core/socket
- Caches: L1 (private per core, 64KB), L2 (shared per socket, 8MB)
- Memory: 2 × 8GB DDR4-2400Mhz
- Coherence Protocol: Hierarchical MOESI (intra-socket), MOSI (inter-socket)
- Interconnect: Inter-socket point-to-point (50ns), intra-socket mesh

Benchmarks

- OpenMP and Pthreads based multithreaded workloads
- 7 benchmark suites – NAS PB, Parboil, Rodinia, PARSEC, SPLASH-2x, SPEC 2017, HPC (assorted)

Comparison Points

- Baseline NUMA: requests routed to node where data is housed
- Intel Memory Mirroring++ (hypothetical): load balances reads between mirrored channels

Dvé
Memory Replication

P

Performance Eval *(Onus Probandi)*

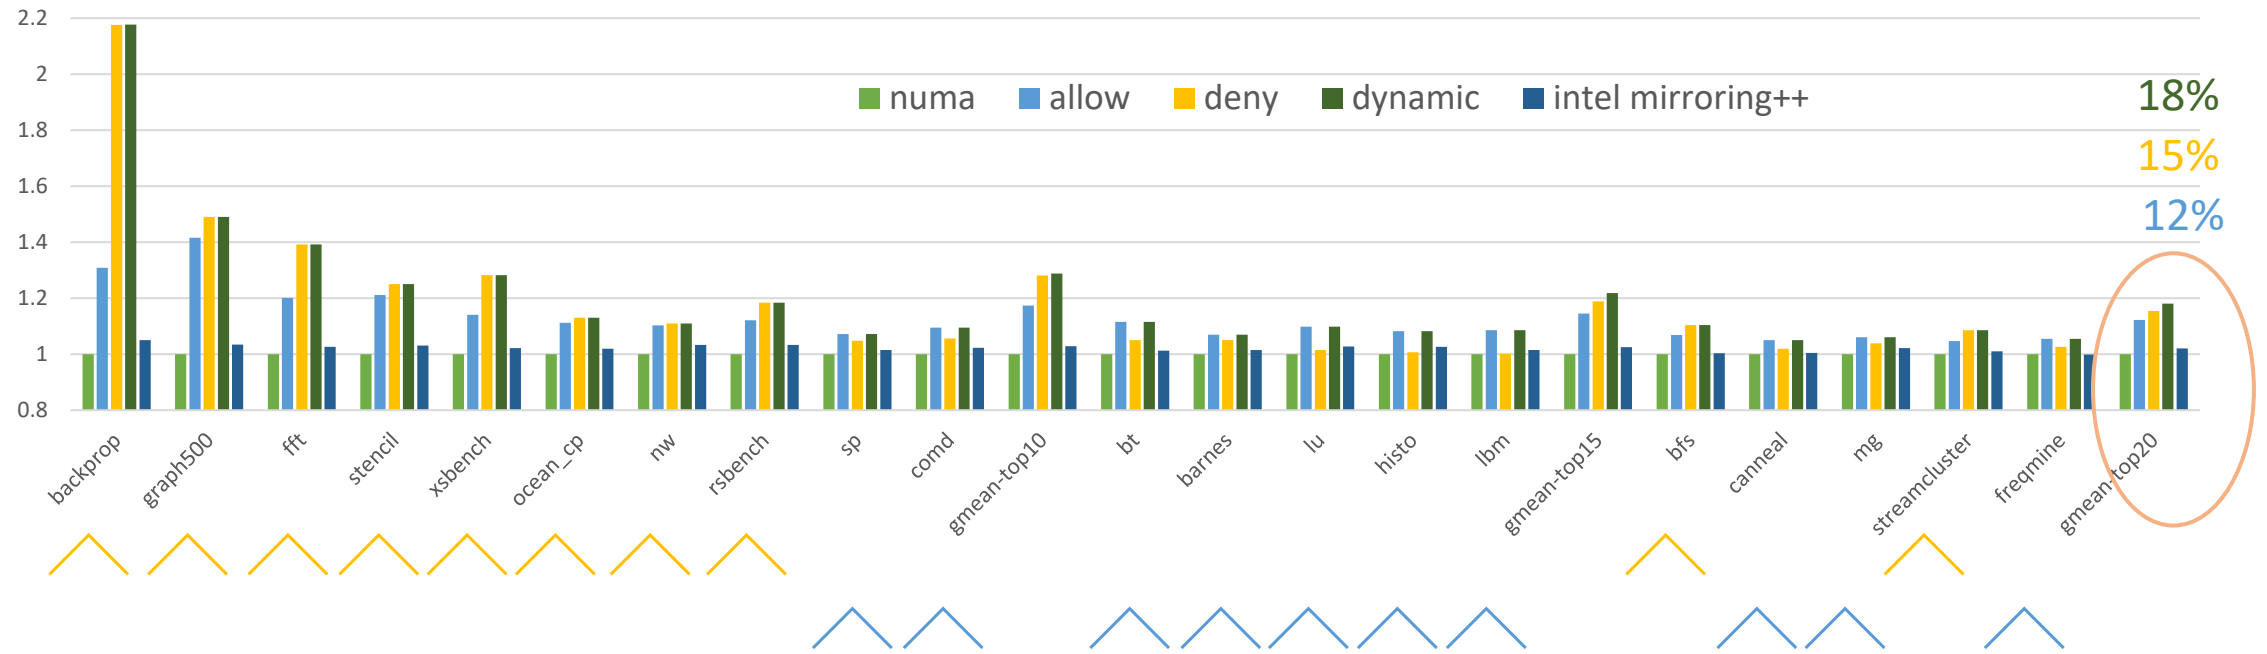

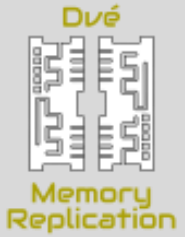

Outline

- Reliability benefits
- Performance gains
- On-demand Reliability

Dvé
Memory Replication

o

Capacity overheads?

Dvé Memory Replication

o

Capacity overheads?

Dvé insights

- ❑ Utilize idle memory

Skewed memory utilization

- ❑ 50% of the memory is idle in 90% of the servers
- ❑ Provisioning for peak

Dvé
Memory Replication

o

Capacity overheads?

Dvé insights

- ❑ Utilize idle memory
- ❑ Overheads applicable only as and when demanded by the application

Skewed memory utilization

- ❑ 50% of the memory is idle 90% of the time
- ❑ Provisioning for peak

Interface to allocate high-reliability memory

- ❑ Hardware-software co-design
- ❑ OS support

Dvé
Memory Replication

o

Capacity overheads?

Dvé insights
- ❑ Utilize idle memory
- ❑ Overheads applicable only as and when demanded by the application

Skewed memory utilization
- ❑ 50% of the memory is idle 90% of the time
- ❑ Provisioning for peak

Interface to allocate high-reliability memory
- ❑ Hardware-software co-design
- ❑ OS support

Flexible trade-off between capacity and reliability

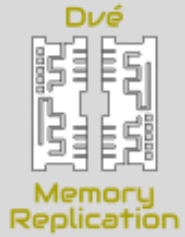

o On-demand Replication *(Onus Probandi)*

Mapping physical address ⟷ replica physical address

- Mapping replica page pairs
  - OS creates page pairs in replica map table (RMT)
  - Single system-wide RMT to create/destroy replica page pairs
  - Hardware-walked RMT at directory controller

- Carving/managing space required for replication
  - Estimate maximum DRAM resident set size
  - Steal memory using balloon drivers
  - Monitor page fault rate for thresholds
  - Modular design allows fallback to baseline reliability

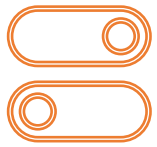

- When should replication enabled or disabled?
  - Notification from Control Plane (managed as a soft-setting)
  - Several configurations possible: per-VM, per-container, kernel-only
  - explicitly specified by application at malloc

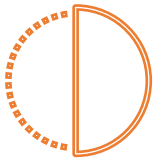

System wide Replication
- Entire memory space replica
- Fixed function mapping

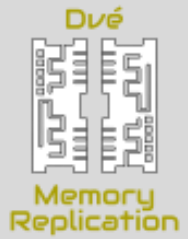

Summary

Replication for reliability

Lowers DUE by

4x over Chipkill
172x over IBM RAIM
11% over Intel Memory Mirroring

Coherent Replication for Performance

Improves performance by

5% - 117% over baseline NUMA
3% - 107% over an improved
Intel mirroring scheme

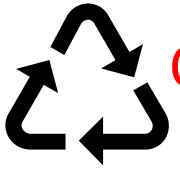

On-demand Replication

hardware-software co-design
using OS/compiler support

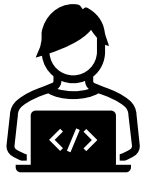

Artifacts available

https://github.com/adarshpatil/dve
https://adar.sh/dve