

HPCA 18

Reliability-aware Data Placement for Heterogeneous memory Architecture

Manish Gupta^Ψ, Vilas Sridharan*, David Roberts*, Andreas Prodromou^Ψ,
Ashish Venkat^Ψ, Dean Tullsen^Ψ, Rajesh Gupta^Ψ



Have you read everything on your car insurance?

Assicurato ha diritto all'indennizzo di invalidità permanente a condizione che la stessa si manifesti entro due anni dall'Infortunio.

La valutazione dell'invalidità permanente sarà effettuata in base alla tabella che segue nella pagina successiva. Se la lesione comporta una diminuzione anziché la perdita totale anatomica o funzionale di un organo o arti, le percentuali della indennità vengono ridotte in proporzione alla funzionalità residua.

La perdita totale anatomica o funzionale di più organi od arti comporta l'applicazione di una tabella di invalidità pari alla somma delle singole percentuali attribuite per ciascuna lesione con un massimo del 100%. Per i casi non previsti dalla tabella, il grado di

ART. C.3 - ESCLUSIONI

Sono esclusi dall'assicurazione i sinistri determinati da:

- a) partecipazione a corse o gare e relative prove ufficiali e verifiche preliminari e finali previste nel regolamento particolare di gara;
- b) tumulti popolari, atti di terrorismo, vandalismo, attentati ai quali l'Assicurato abbia partecipato attivamente;
- c) guerra, insurrezioni, terremoti, eruzioni vulcaniche;
- d) trasmutazione del nucleo dell'atomo come pure dovuti ad esposizione a radiazioni ionizzanti;

ART. C.4 - LIQUIDAZIONE

con un'eccezione: l'arresto di un arto superiore del piede un arto inferiore all'altezza di sotto al ginocchio un rene un occhio ambedue gli occhi un rene la milza sordità completa di un orecchio sordità completa di ambedue gli orecchi perdita totale dell'udito voce postumi di trauma distorsione cervicale con contrattura muscolare limitazione dei movimenti del capo del collo

Liquidazione incaricato dalla Società giustificati oppur interv Sanità quota prede anticip sanita presen incaric danno

The insurance does not cover those accidents caused by:

[...]

exposure to ionizing radiation*

Reliability Matters



Autonomous vehicles

- Safety is important

High performance computing

- Long running scientific jobs

Why Focus on Memory?

Cielo at Los Alamos National Lab



8-core AMD Opteron™ CPUs

8,944 nodes : 1,144,832 DRAM

DDR-3 DRAM, Chipkill-correct ECC

Hopper at NERSC / Lawrence Berkeley National Lab



12-core AMD Opteron™ CPUs

6,384 nodes : 817,152 DRAM

DDR-3 DRAM, Chipkill-detect ECC

- Most of your computer is in fact memory
- The probability of a bit upset is proportional to silicon surface area

Large-scale Systems Magnifies Failures

Cielo at Los Alamos National Lab



8-core AMD Opteron™ CPUs
8,944 nodes : 1,144,832 DRAM
DDR-3 DRAM, Chipkill-correct ECC


Hopper at NERSC / Lawrence Berkeley National Lab



12-core AMD Opteron™ CPUs
6,384 nodes : 817,152 DRAM
DDR-3 DRAM, Chipkill-detect ECC

Even though failure rate for each device seems low, the systems have millions of devices and failure rates are additive

Heterogeneous Memory Architecture



- Heterogeneous Memory Architectures (HMA) consist of multiple memory modules.
 - For example: An HMA system with HBM + DDRx
- Most research on HMAs present only performance trade-offs of placing data in one memory over the other
- Heterogeneity in two axes: 1) Reliability and 2) Performance
- We present techniques to balance both axes

Outline

- Motivation
- Background
- Estimating Data Vulnerability using AVF
- Data AVF vs. Hotness
- Evaluation Methodology
- Results
- Summary

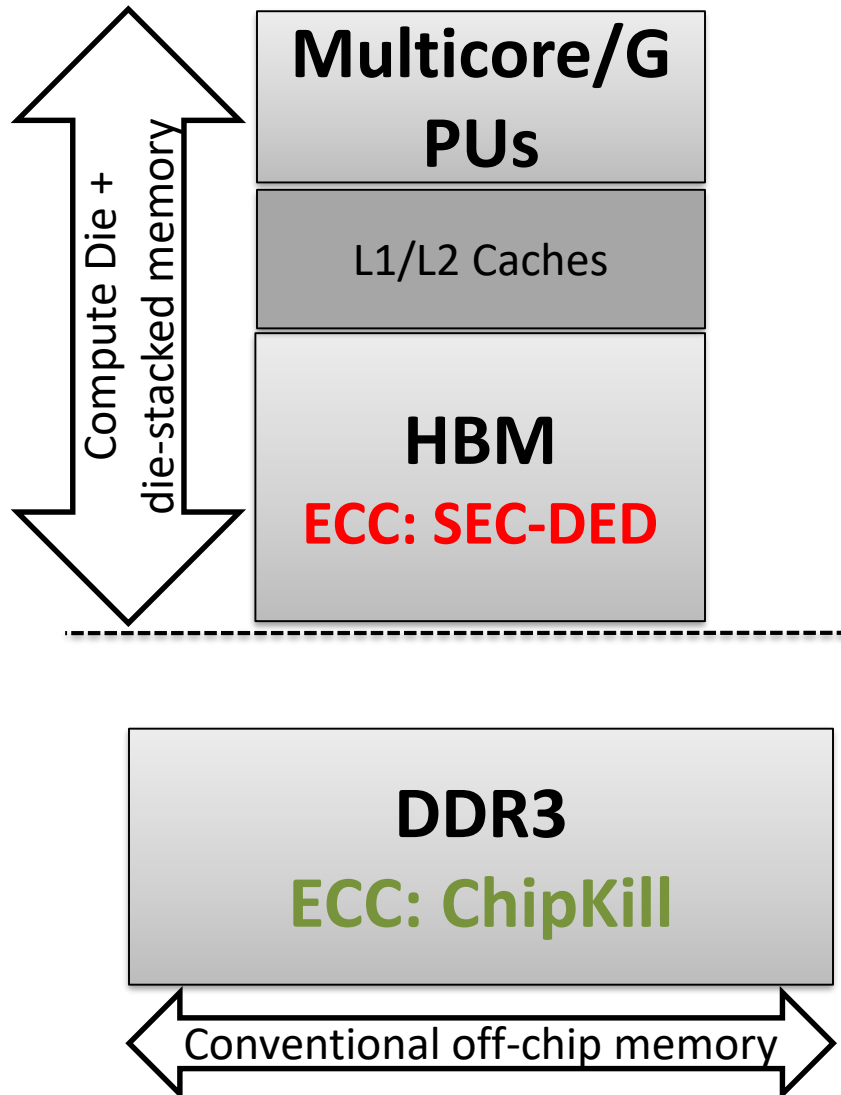
Background: Faults vs. Errors

- **Faults** are underlying cause of a hardware failure
 - **Permanent Faults** For example: consistently wrong value returned from memory due to hardware fault (stuck-at bit)
 - **Transient Faults** For example: soft errors due to single-event upsets or voltage droop
- **Errors** are manifestation of faults
 - Errors can be detected and/or corrected. For example using error correcting codes (ECC)

FIT (Failure in Time)

- **Failure In Time (FIT) is a measure to quantify system reliability**
- **1 FIT for large-scale system such as “Cielo”**
 - 1 FIT per node with 8,944 nodes = Failure every **12.8 years**
 - 1 FIT per DIMM for 71,552 DIMMs = Failure every **1.6 years**
 - 1 FIT per DRAM 1,144,832 DRAM chips: Failure every **36 days**
- **Real FIT rates (FIT rates for components on Cielo)**
 - Target **socket FIT rate of 1000**: failure every **2.3 days**
 - Target **DRAM chip FIT rate of 35**: failure every **1 days**

Heterogeneous Memory Architecture (HMA an Example System)



HMA system shows heterogeneity in not only performance but also **reliability**

SEC-DED (ECC)

- Single-bit Error Correct Double-bit Error Detect
- Easy to implement
- Loses efficacy with aging [1]
- 4x-8x higher bandwidth than DDR3

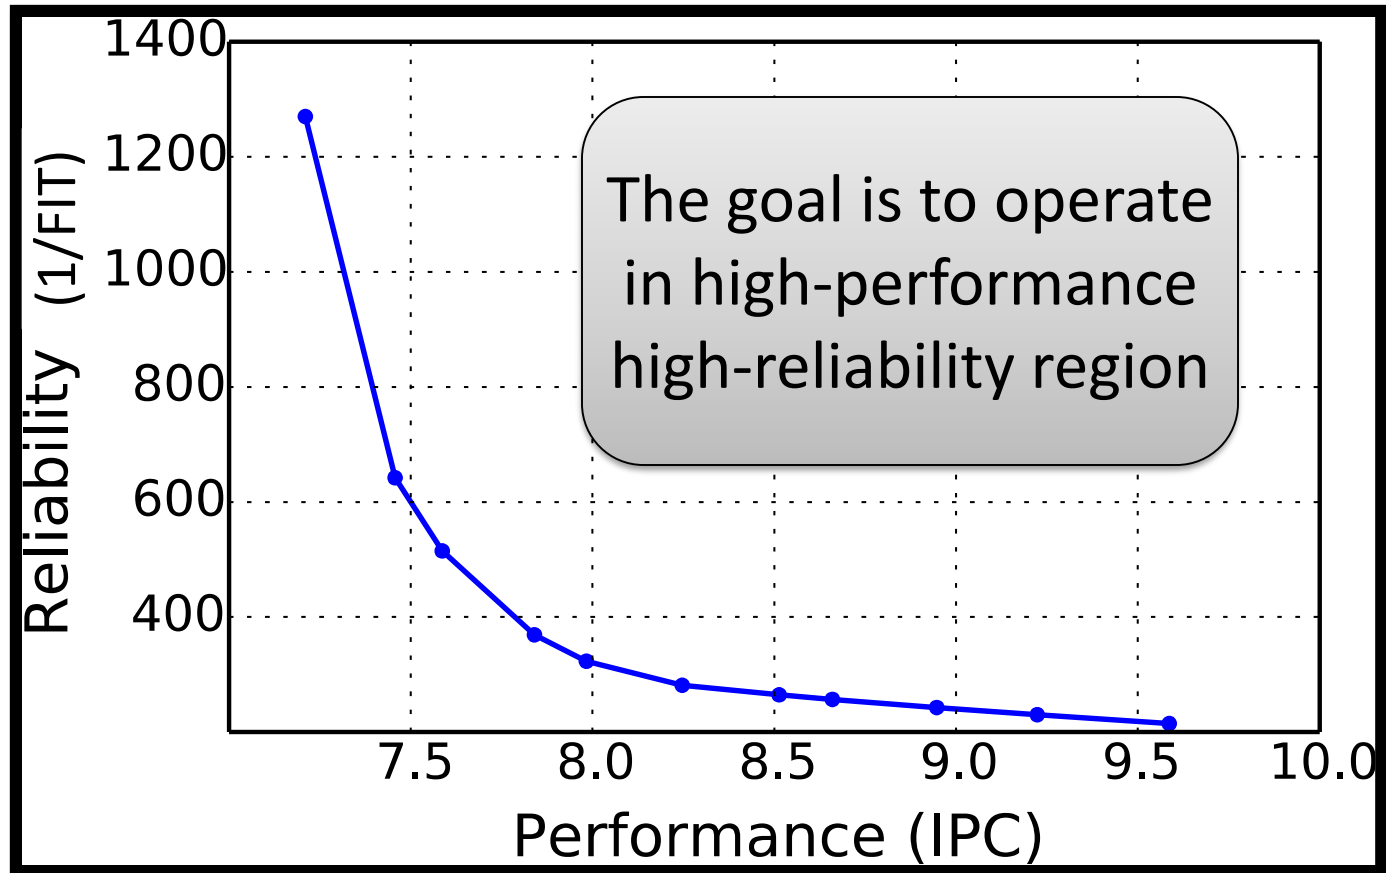
ChipKill (ECC)

- Symbol-based correcting code
- Requires distributing data to multiple devices
- ChipKill is 42x more effective than SEC-DED [2]
- Low bandwidth

[1] M. Gupta et al. Reliability vs. Performance Trade-off Study of Heterogeneous Memory Architectures in MEMESYS16

[2] V. Sridharan et al. A Study of DRAM Failures in the Field in SC12

Reliability vs. Performance




Reliability-aware Data Placement

Move **hot** pages to HBM memory

Data hotness

Estimate using access counters (WRs+RDs)



Data vulnerability


How to estimate page vulnerability?

Move **vulnerable** ("risky") pages to DDRx memory

Outline

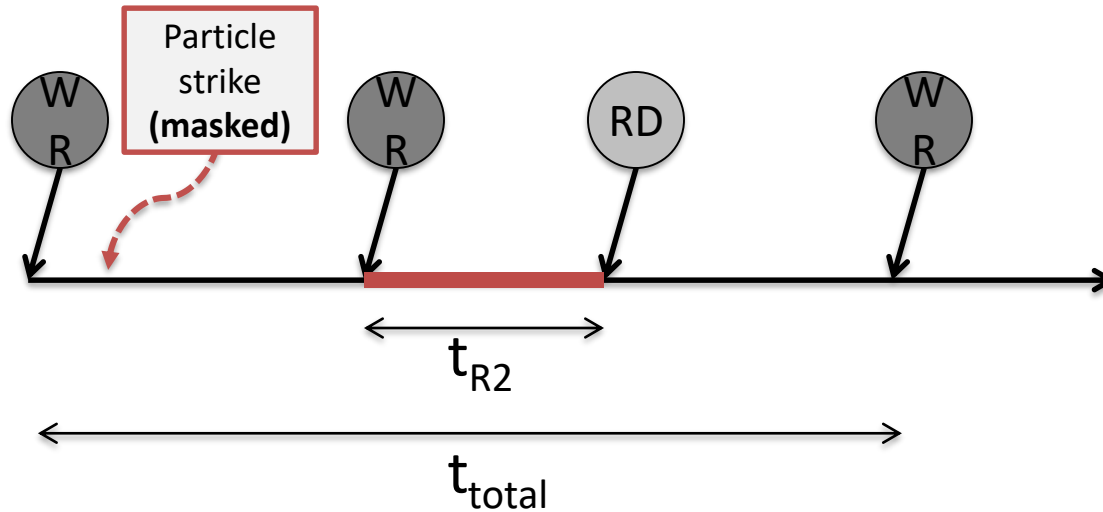
- Motivation
- Background
- Estimating Data Vulnerability using AVF
- Data AVF vs. Hotness
- Evaluation Methodology
- Results
- Summary

Data Vulnerability through Architectural Vulnerability Factor (AVF) [1]



$$\text{Vulnerability of a bit} = \frac{t_{R1} + t_{R2}}{t_{total}}$$

Data Vulnerability through Architectural Vulnerability Factor (AVF) [1]



$$\text{Vulnerability of a bit} = \frac{t_{R2}}{t_{total}}$$

Definitions: AVF and SER

$$AVF_{Mi} = \frac{\Sigma(\text{Vulnerability of a bit})}{(\text{Number bits in the structure } Mi)}$$

AVF
Vulnerability
of the
structure

Soft Error Rate (SER)


Probability of
uncorrectable software
visible error

$$SER_{Mi} = \text{Failure Probability}_{Mi} \times (AVF_{Mi})$$

Probability of uncorrectable
hardware fault
(Device FIT rate)

Scaling it with
vulnerability
factor
(risk factor)

The Goal



<u>Memory 1 (HBM)</u> High Bandwidth Low Reliability	<u>Memory 2 (DDRx)</u> Low Bandwidth High Reliability
Hot & low-risk pages	Cold & high-risk pages

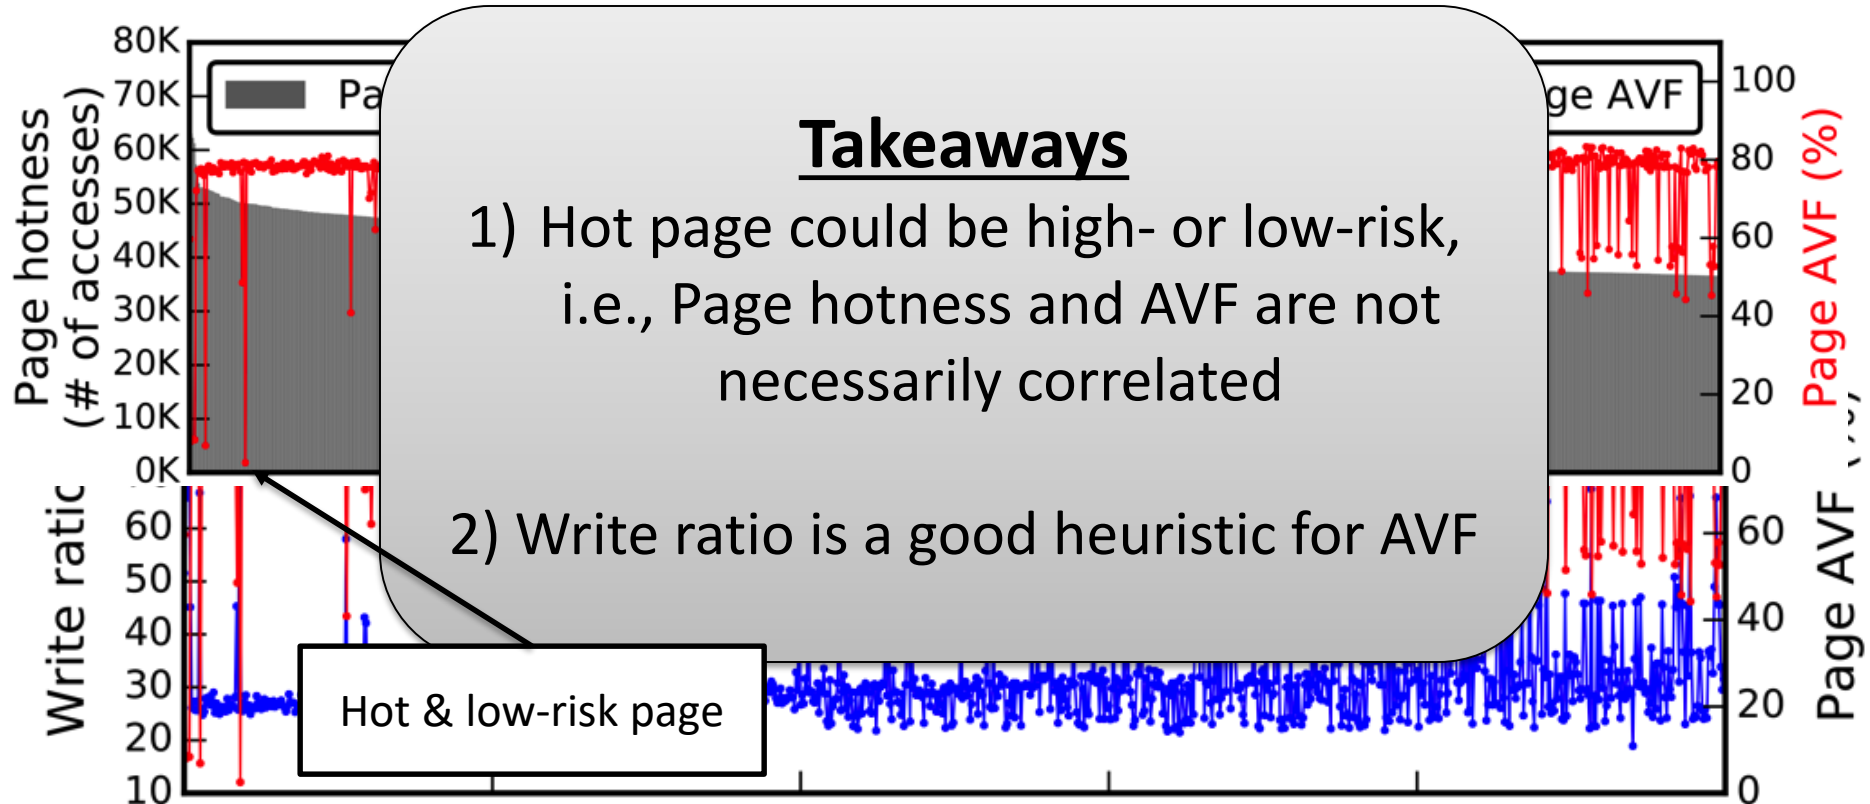
The goal is to find hot & low-risk pages for HBM

Outline

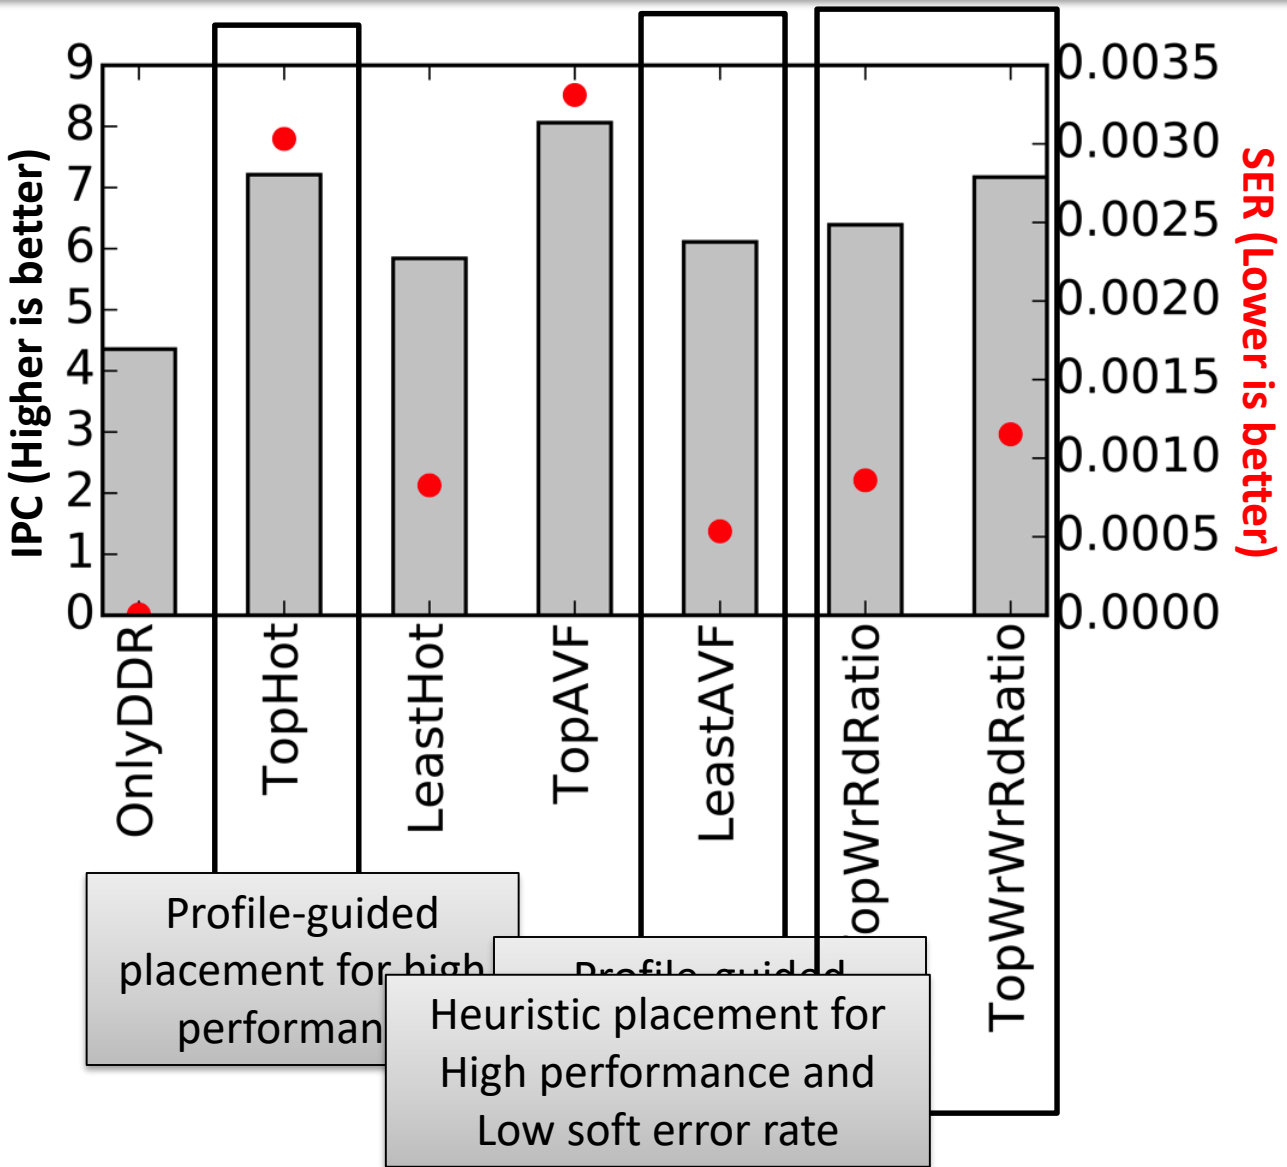
- Motivation
- Background
- Estimating Data Vulnerability using AVF
- Data AVF vs. Hotness
- Evaluation Methodology
- Results
- Summary

Data (Memory Page) AVF vs. Hotness

Is hotness correlated with risk (AVF)?



Profile-guided Data Placement (One Workload)



Takeaways

- 1) There is a nice trade-off curve in between IPC and SER
- 1) We have operating points with low SER and high IPC

Outline

- Motivation
- Background
- Estimating Data Vulnerability using AVF
- Data AVF vs. Hotness
- Evaluation Methodology
- Results
- Summary

Evaluation Methodology

DRAM Failure Data and Simulation Tools

- Jaguar Cluster [1] with 2.69M DRAM devices
- FaultSim [2] for memory failures and different ECCs
- Ramulator [3] for performance simulations

Evaluation and Results

- On homogeneous and mixed 16-core multi-programmed workloads created using SPEC2006 benchmarks
- We show IPC and SER for different placements averaged for homogenous, mixed, and all workloads

[1] A Study of DRAM Failures in the Field, Sridharan et al. SC 2012

[2] FaultSim: <https://github.com/Prashant-GTech/FaultSim-A-Memory-Reliability-Simulator>, Nair et al. TACO 2016

[3] Ramulator: <https://github.com/CMU-SAFARI/ramulator>, Kim et al. IEEE CAL 2015

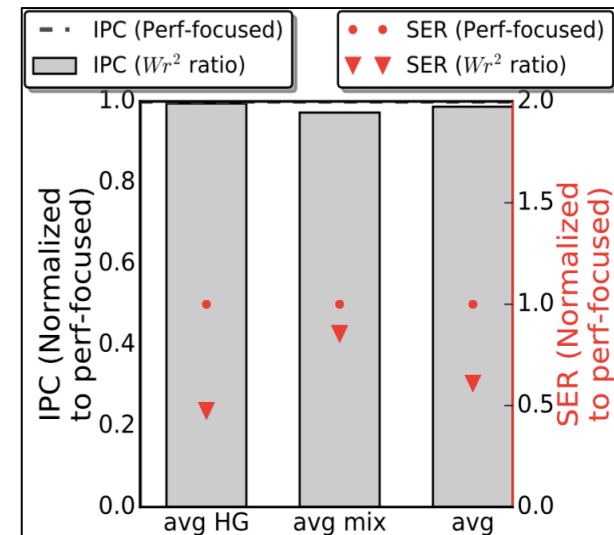
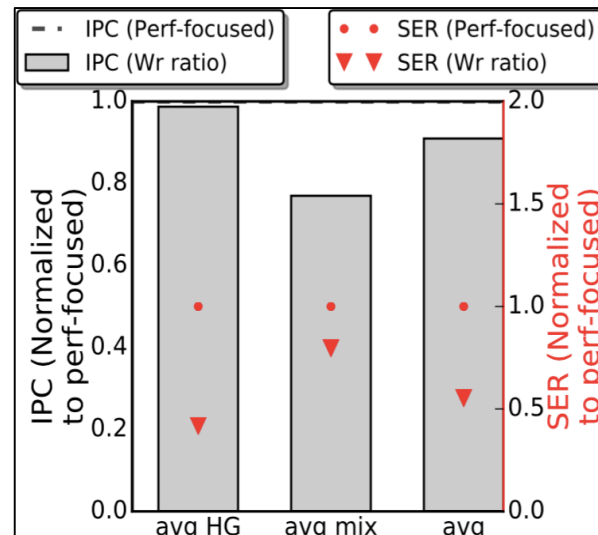

Outline

- Motivation
- Background
- Estimating Data Vulnerability using AVF
- Data AVF vs. Hotness
- Evaluation Methodology
- **Results**
 - Profile-guided Data Placement
 - Dynamic Data Placement
 - Program Annotations
- Summary

Profile-guided Data Placement

Goal

Reduce **SER** to as low as possible
Keeping **IPC** as close as possible to performance-focused IPC



AVF-focused

SER reduction: 5x

IPC loss: 17%

Top Wr/Rd Heuristic

SER reduction: 1.8x

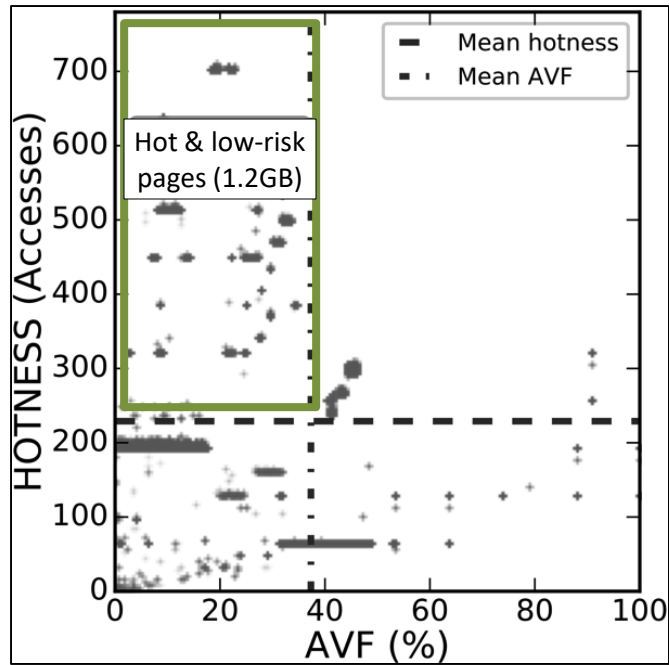
IPC loss 8.1%

Top Wr²/Rd Heuristic


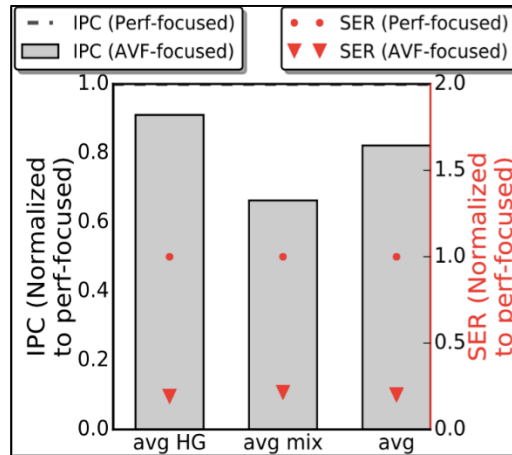
SER reduction: 1.6x

IPC loss: 1%

Homogenous vs. Mix Workloads (AVF-focused)




Homogenous workload
Small hotness and AVF spans
Page hotness 0 to 700

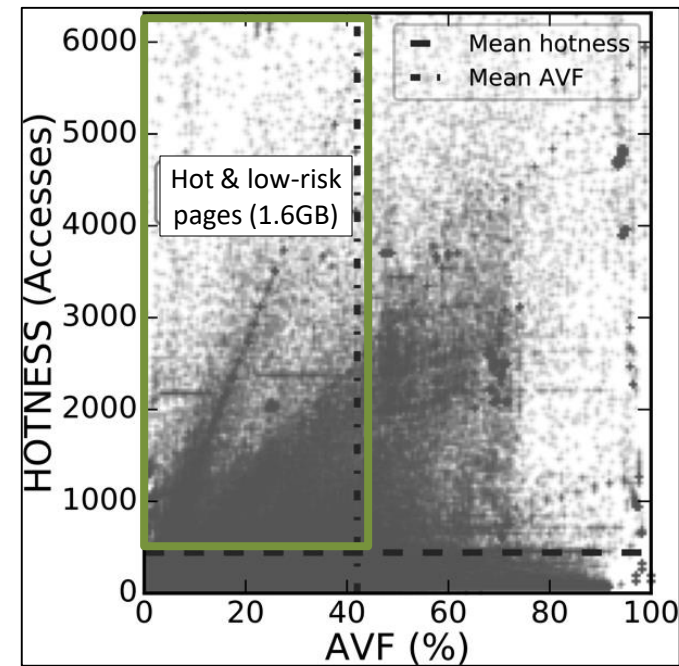



Mix workload
Large hotness and AVF spans
Page hotness 0 to 6000

Homogenous vs. Mix Workloads (Wr^2/Rd Heuristic)



Homogenous workload
Small hotness and AVF spans
 Page hotness 0 to 700

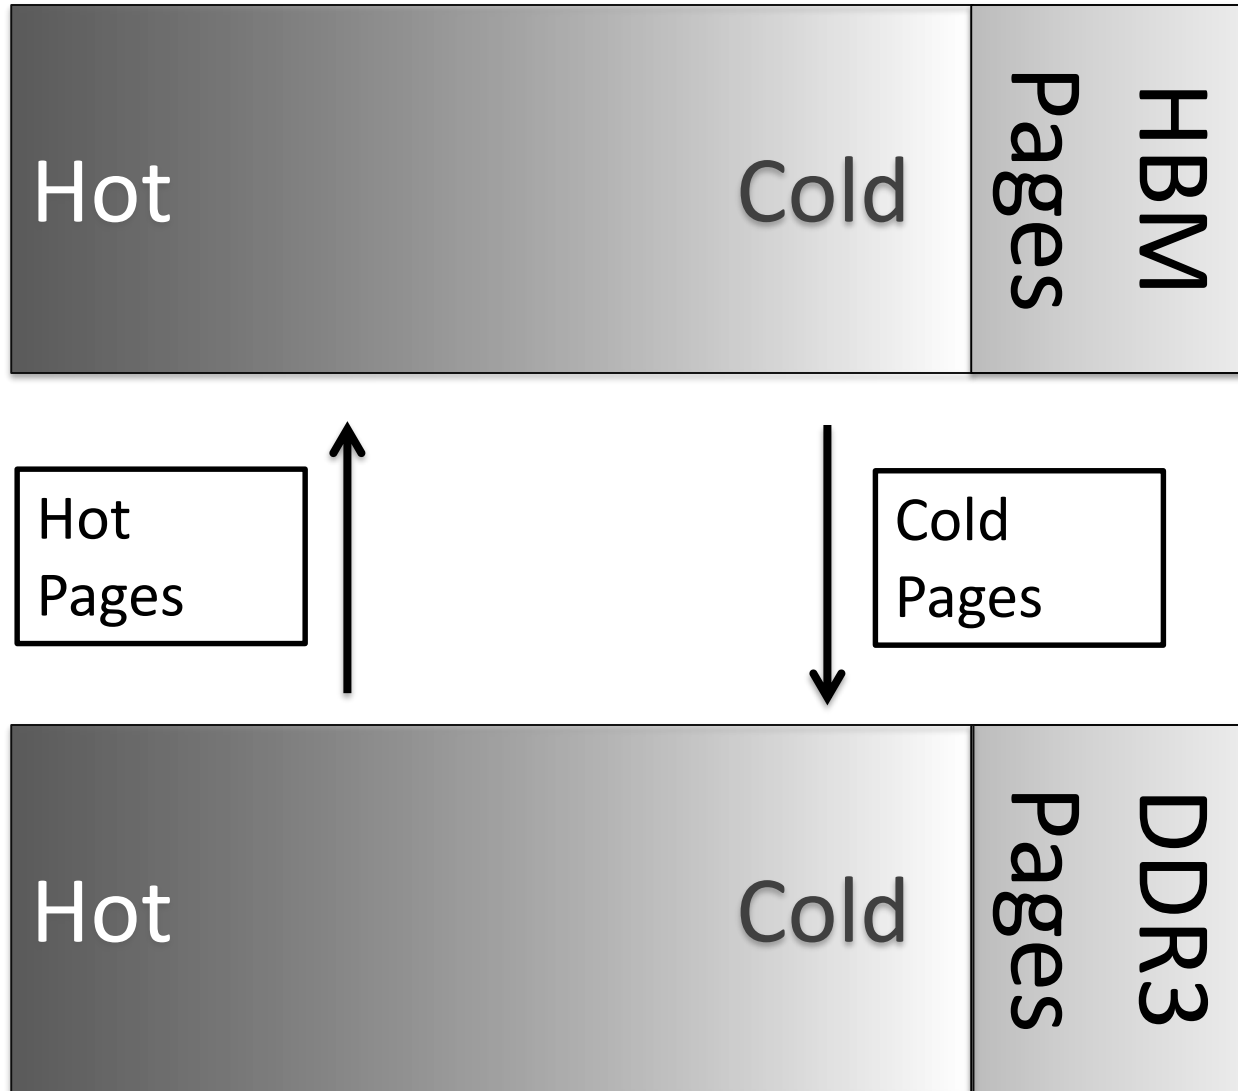


Mix workload
Large hotness and AVF spans
 Page hotness 0 to 6000


Outline

- Motivation
- Background
- Estimating Data Vulnerability using AVF
- Data AVF vs. Hotness
- Results
 - Profile-guided Data Placement
 - **Dynamic Data Placement**
 - Program Annotations
- Summary

Performance-focused Dynamic Migration



Vulnerability-aware Dynamic Migrations




Hot &
low-risk
pages


An upward-pointing arrow is positioned to the right of the text box.

Cold &
high-risk
pages

A downward-pointing arrow is positioned to the left of the text box.



Reliability-aware Dynamic Migrations [1]




FullCounters [1]

SER reduction: 1.8x

IPC loss 6%

Perf-focused migrations using FullCounters [2]

- FullCounters for hotness
- One counter per page
- Base case for migrations
- 4.25 MB of counters



CrossCounters [1]

SER reduction: 1.5x

IPC loss: 4.9%

Reliability-aware migrations using FullCounters [1]

- FullCounters for hotness and vulnerability heuristics
- Two counters per page
- **8.5 MB of counters**

Reliability-aware migrations using CrossCounters [1]

- MEA counters [3] for hotness + FullCounters for vulnerability heuristics
- **600 KB of counters**

[1] M. Gupta et al. "Reliability-aware Data Placement for Heterogeneous Memory Architectures" in HPCA18


[2] M. Meswani et al. "Heterogeneous Memory Architectures: A HW/SW Approach for Mixing Die-staked and Off-package Memories" in HPCA15

[3] A. Prodromou et al. "MemPod: A Clustered Architecture for Efficient and Scalable Migration in Flat Address Space Multi-level Memories" in HPCA17

Outline

- Motivation
- Background
- Estimating Data Vulnerability using AVF
- Data AVF vs. Hotness
- Results
 - Profile-guided Data Placement
 - Dynamic Data Placement
 - Program Annotations
- Summary

Program Annotations



- Annotating only one program structure pins ~512 MB hot & low-risk data in HBM
- Results in **SER reduction** of 1.3x at IPC loss of 1.1%
- Thus, minimal program annotation results in improved reliability at marginal performance loss

Summary

Heterogeneous memory architecture are becoming popular

Heterogeneity exists not only in performance, but also in reliability

We discussed techniques to balance both performance and reliability

Disclaimer & Attribution

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTIONS

© 2016 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Other names are for informational purposes only and may be trademarks of their respective owners.

Backup

Thanks



MESL, UCSD



Architecture Lab, UCSD




AMD Research

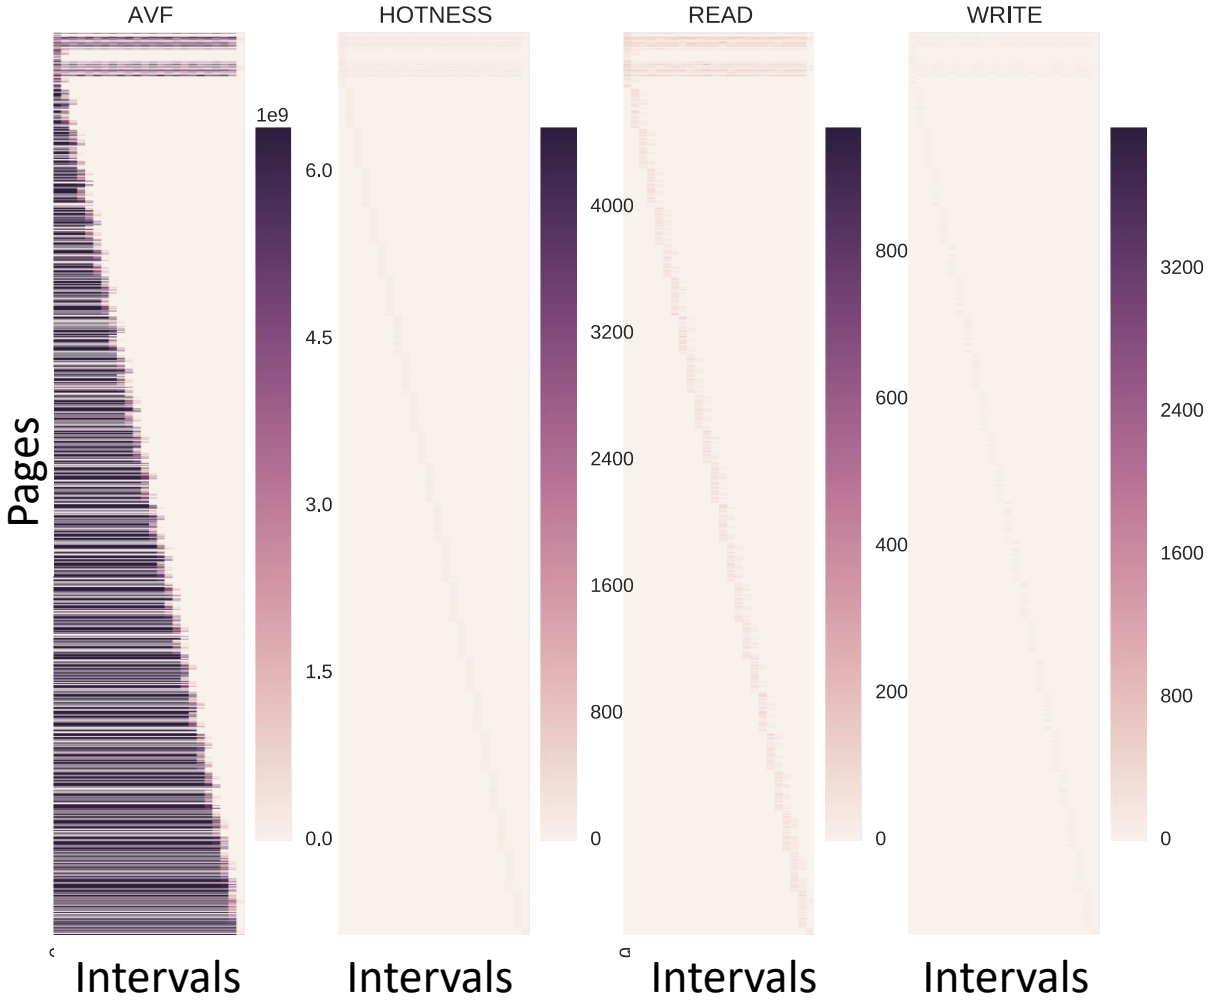


PL, UCSD

Astar Heatmap



Closer Look (Dynamic Migrations)



FaultSim

- Fault Simulation can be using
 - analytical model
 - Interval-based simulations
 - In interval-based simulations, introduces fault in the memory
Components based on FIT rates, apply ECC and report error rates
 - Event-based simulations
 - Failure per device happens rarely. Thus, instead of asking random number generate if there's a fault in this interval. One can ask the random number generate what's the timing difference between the next interval.


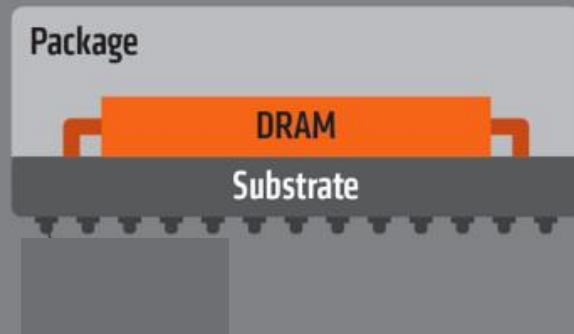
A Study of DRAM Failures in the Field SC 2012

Table I. DRAM Failures per Billion Device Hours (FIT)
[Sridharan and Liberty 2012]

DRAM Chip Failure Mode	Fault Rate (FIT)	
	Transient	Permanent
Single bit	14.2	18.6
Single word	1.4	0.3
Single column	1.4	5.6
Single row	0.2	8.2
Single bank	0.8	10
Multibank	0.3	1.4
Multirank	0.9	2.8

- More than 2000 DRAM devices experience a single fault
- Logging using x86 Machine-check registers to log corrected and uncorrected errors
- 250K errors (corrected + uncorrected) per month. 6.6 errors per node per month
- Transient vs. Permanent separation. Using scrubbing interval

High Bandwidth Memory (HBM) [1]



Conventional DDR Memory

32-bit

Up to 1750MHz

Up to 28GB/s per chip

1.5V

Per Package

Bus Width

Clock Speed

Bandwidth

Voltage

High Bandwidth Memory (HBM)

1024-bit

Up to 500MHz

>100GB/s per stack

1.3V