Modeling the Implications of DRAM Failures and Protection Techniques on Datacenter TCO

Panagiota Nikolaou[1], Yiannakis Sazeides[1], Lorena Ndreu [1], Marios Kleanthous [2]

[1]University of Cyprus, [2]MAP S.Platis

MICRO 48, Waikiki, Hawaii, December 5th 2015

Today's Datacenters

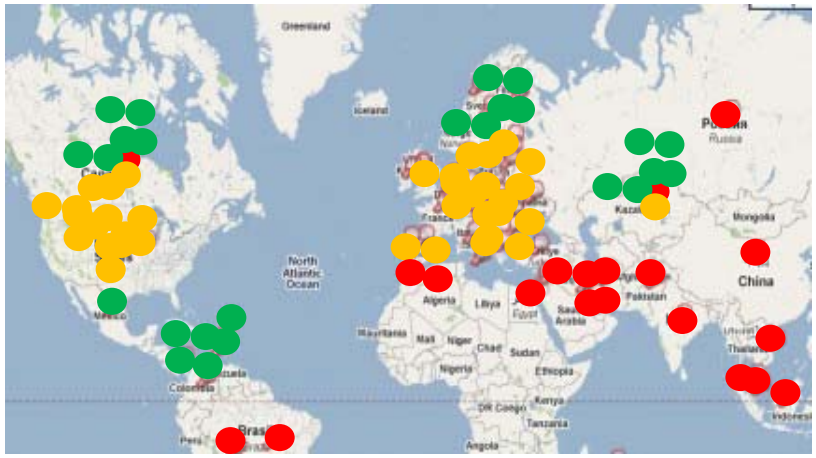

> 510,000 DC in all over the world [Emerson, 2011]

> 285 Million Sqft [Emerson, 2011]

Large scale Datacenters: >10,000 commodity servers

Many Million $ per month

Datacenter Cost

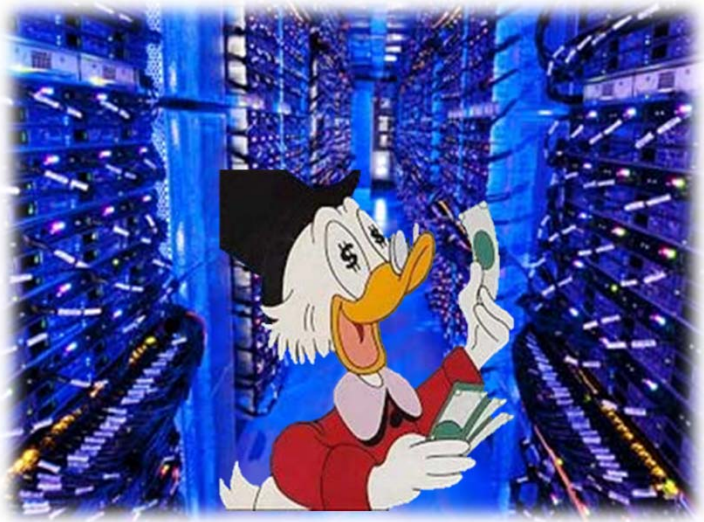

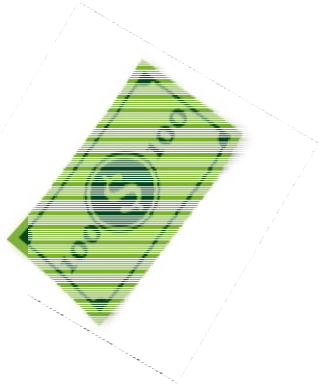

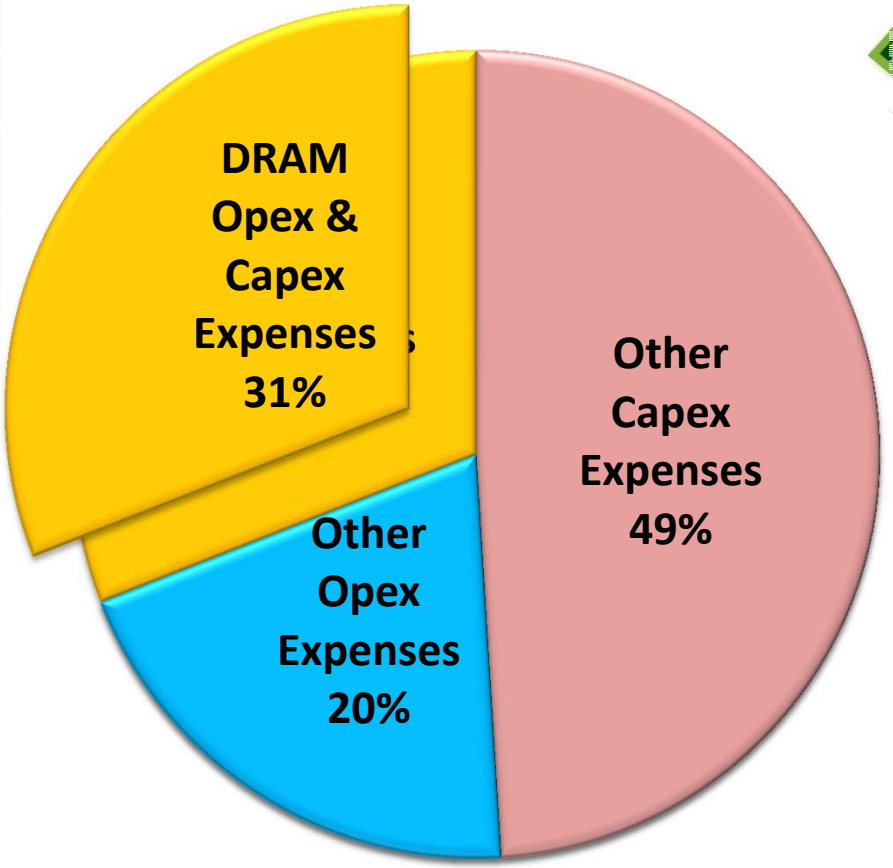

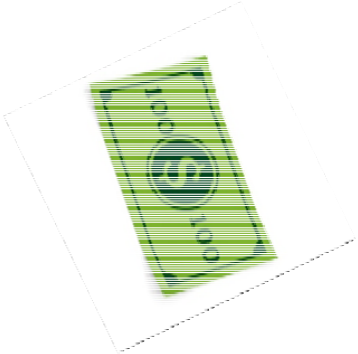

[Analysis using COST-ET tool, D. Hardy 2013]

DRAM Protection Cost

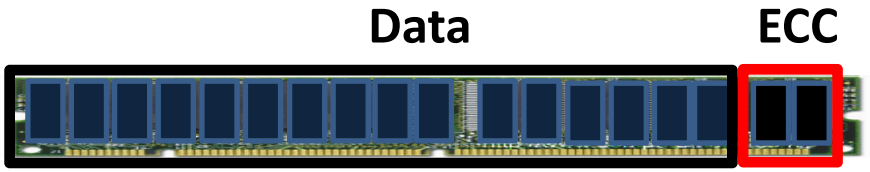

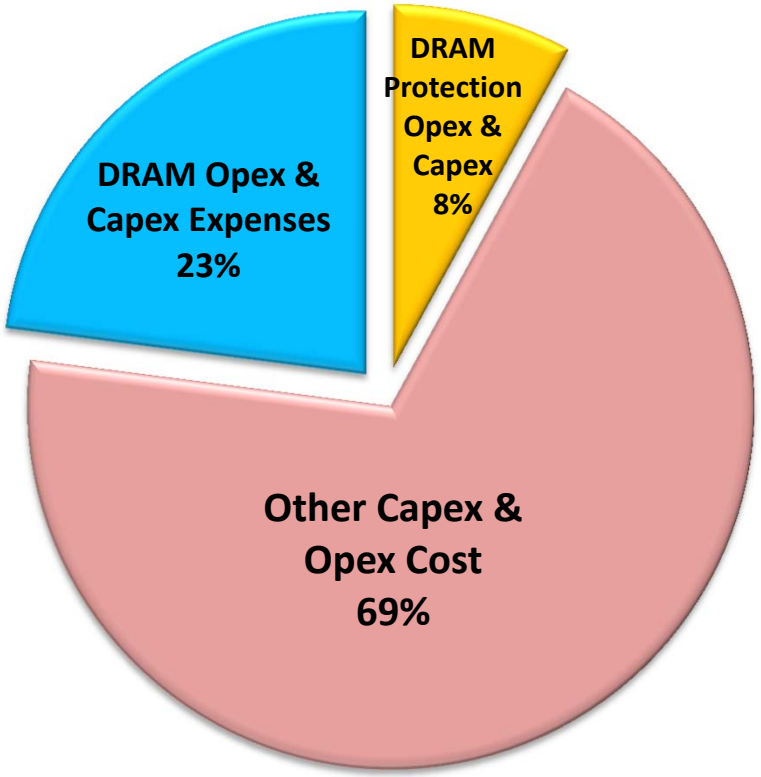

[Analysis using COST-ET tool, D. Hardy 2013]

Do we need DRAM protection?

[Borucki, IRPS 2008], [K. Lim ISCA 2009], [Daniel Bowers, "Server Trends"]

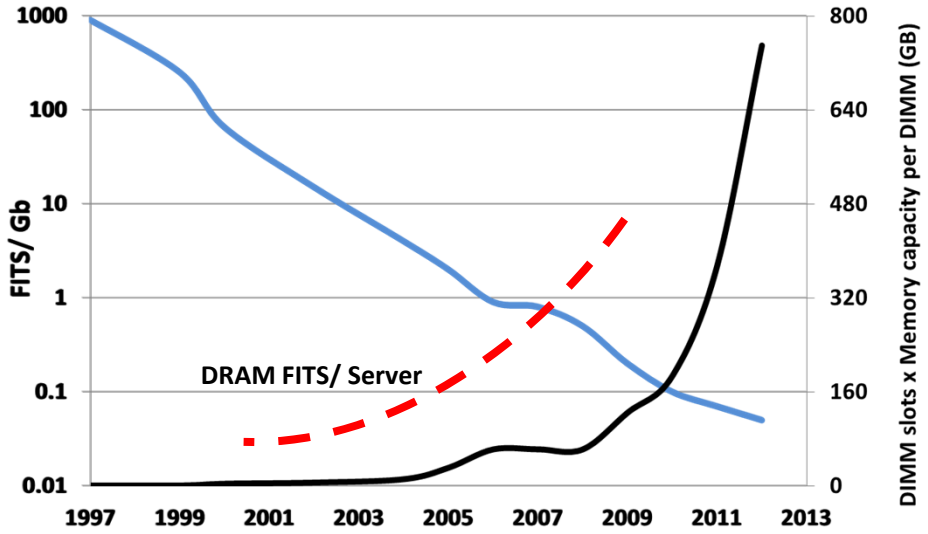

- **Google Failure Study** [Barroso, 2009]
- **DRAM large field studies** [V. Shridharan 2012, 2013]

DRAM protection is essential !!

DRAM protection choices

ChipkillDC

Cost	+
Reliability	+++
Performance	+

ChipkillSC

Cost	++
Reliability	++
Performance	++

SECDED

Cost	+++
Reliability	+
Performance	+++

DRAM protection selection

***Analyzer of Memory Protection and Failures Implications on TCO (AMPRA tool), site: http://www2.cs.ucy.ac.cy/carch/xi/ampra_tco.php**

Our Proposition & our Contribution

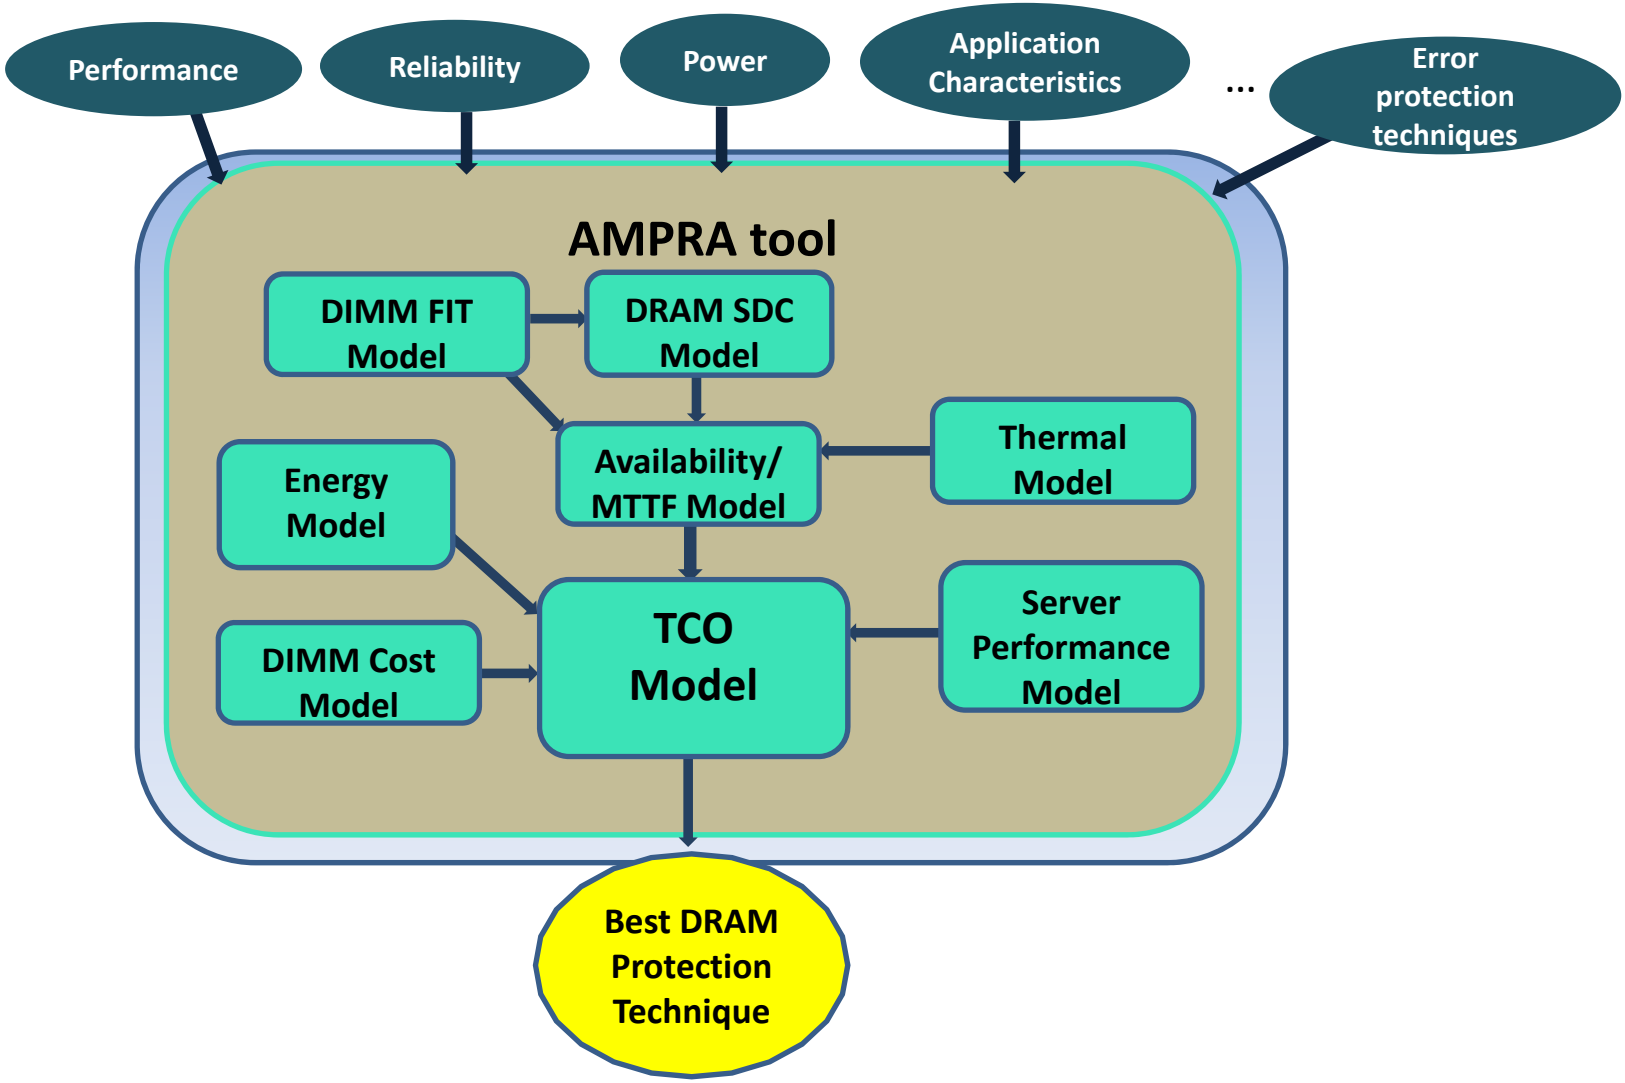

Related work

- **[Y. Luo DSN 2014]** Proposes and analyzes cost of a heterogeneous memory protection scheme

  Differences:
  - Performance, power implications of memory protection techniques
  - Co-located services
  - Datacenter cost

- No other related work considers various parameters


Outline

- Proposed Framework (AMPRA tool)
- Use Case
- Experimental Framework
- Results
- Conclusions

Proposed framework (AMPRA tool)

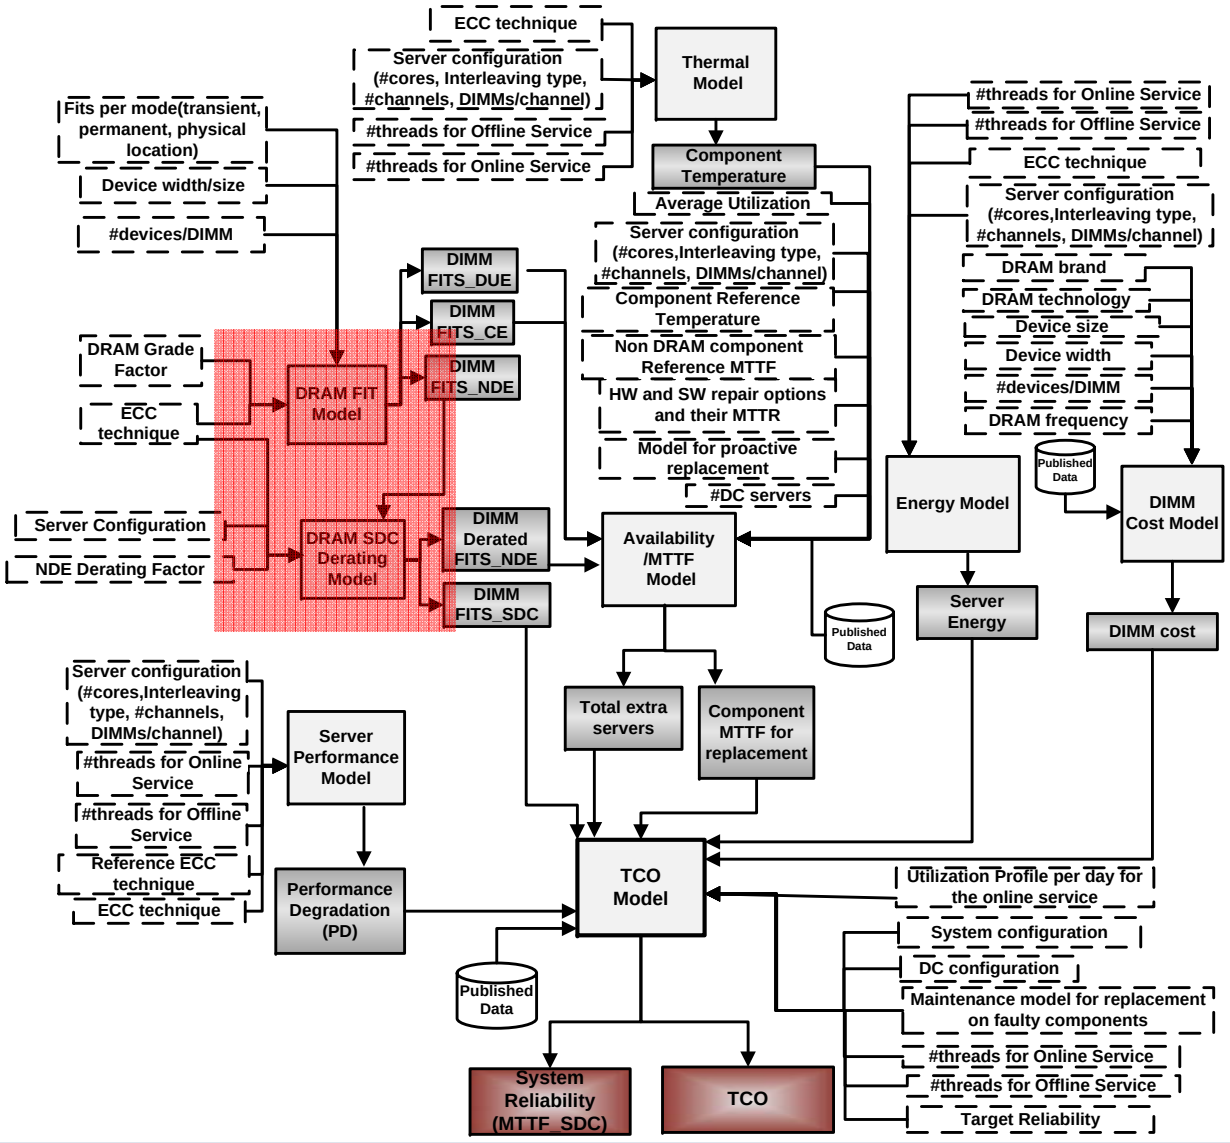


Outline

- Proposed Framework (AMPRA tool)
- Use Case
- Experimental Framework
- Results
- Conclusions

Use Case

Bandwidth vs. Latency vs. Reliability vs. Power

Chipkill with Dual Channel Implementation **(ChipkillDC)**

Chipkill with Single Channel Implementation **(ChipkillSC)**

16 ECC bits for 128 Data bits-144 bit codeword

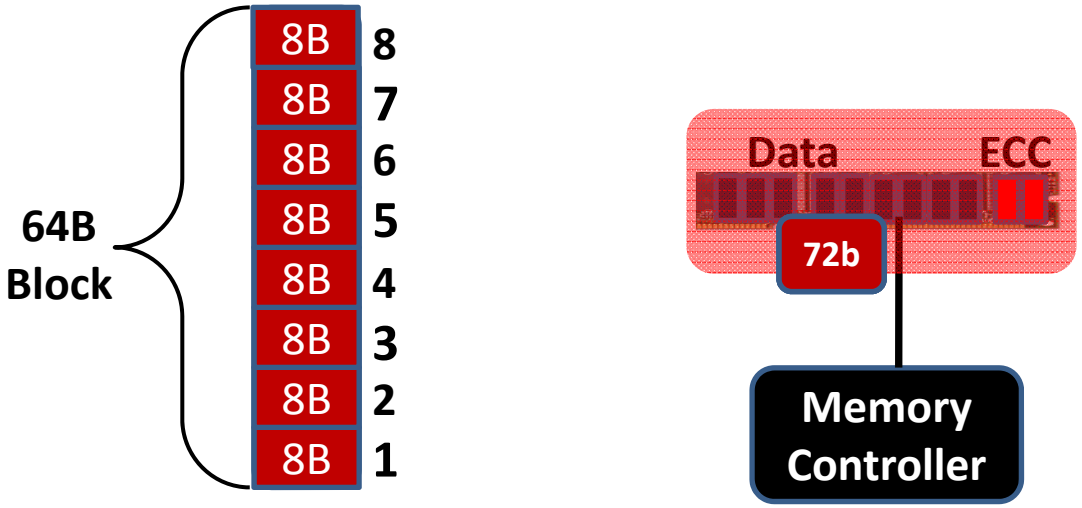

FIT model

- **ChipkillDC**:
  - **Detects** all the errors in 2 devices
  - **Corrects** all the errors in 1 device

- **ChipkillSC**:
  - **Cannot detect** all the errors in 2 devices
  - **Corrects** all the errors in 1 device

ChipkillDC can provide better Reliability than ChipkillSC

Performance and Power model

How it works: (ChipkillDC)

- Read

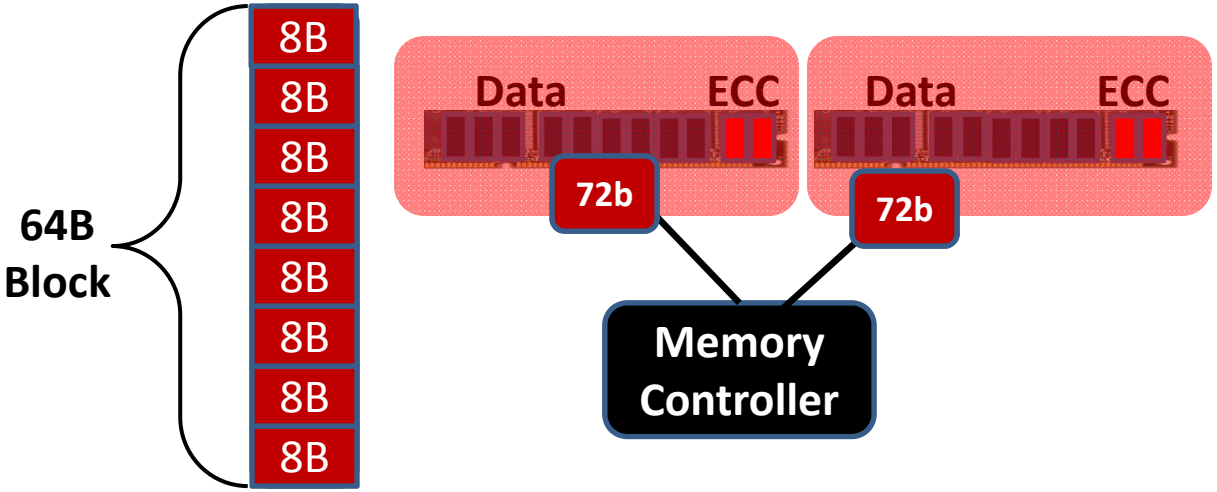

- **Requires accessing two DIMMs**
- **Codeword in a single burst**
- **Latency short**
- **Low Bandwidth**
- **High Power Consumption**

Performance and Power model

How it works: (ChipkillSC)

- Read

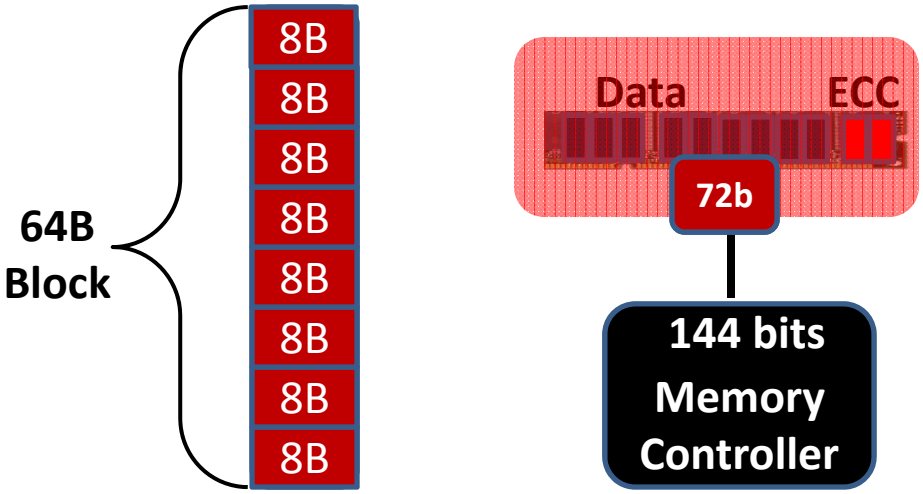

- Requires accessing one DIMM
- Codeword in two bursts
- Latency long ☹
- High Bandwidth ☺
- Less Power Consumption ☺

Design Space

	ChipkillSC	ChipkillDC
Reliability	✓ Cannot detect all the errors in 2 devices Corrects all the errors in 1 device	✓✓ Detect all the errors in 2 devices Corrects all the errors in 1 device
Bandwidth	✓ Access one DIMM	✗ Access two DIMMs
Latency	✗ Codeword in two bursts	✓ Codeword in one burst
Power	✓ Access one DIMM	✗ Access two DIMMs

What happens with the Cost?

- **Application characteristics**
  - **Memory intensive, compute intensive**
- **Co-running applications**

Online and Offline Services

Online Services: High QoS requirements

Offline Services: Do not have QoS constrains

Co-location: Improve server utilization and reduce TCO

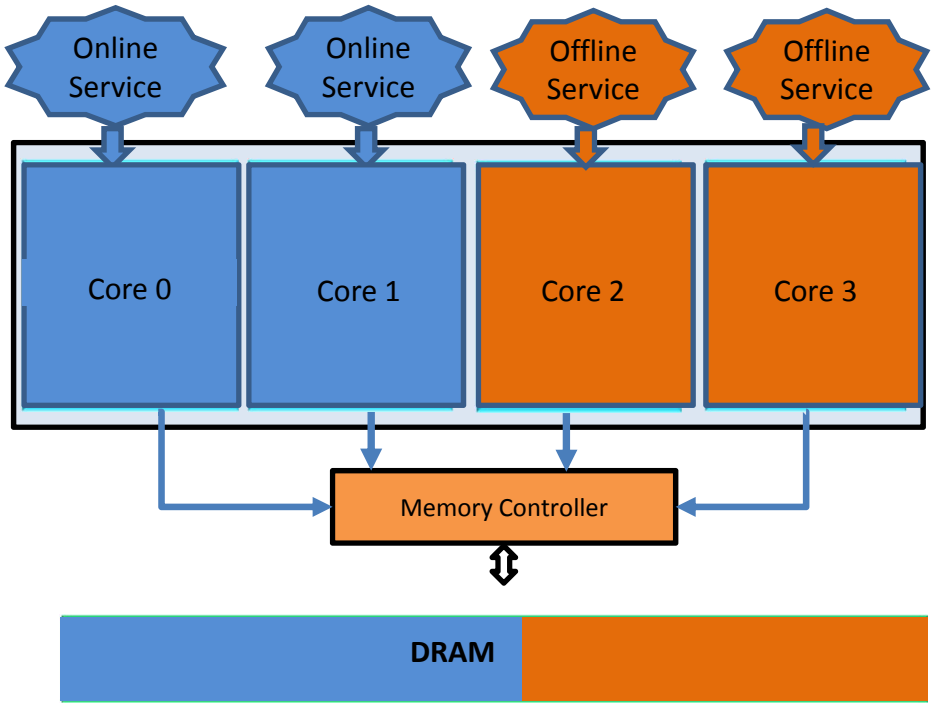


Outline

- Proposed Framework (AMPRA tool)
- Use Case
- Experimental Framework
- Results
- Conclusions

Experimental Framework

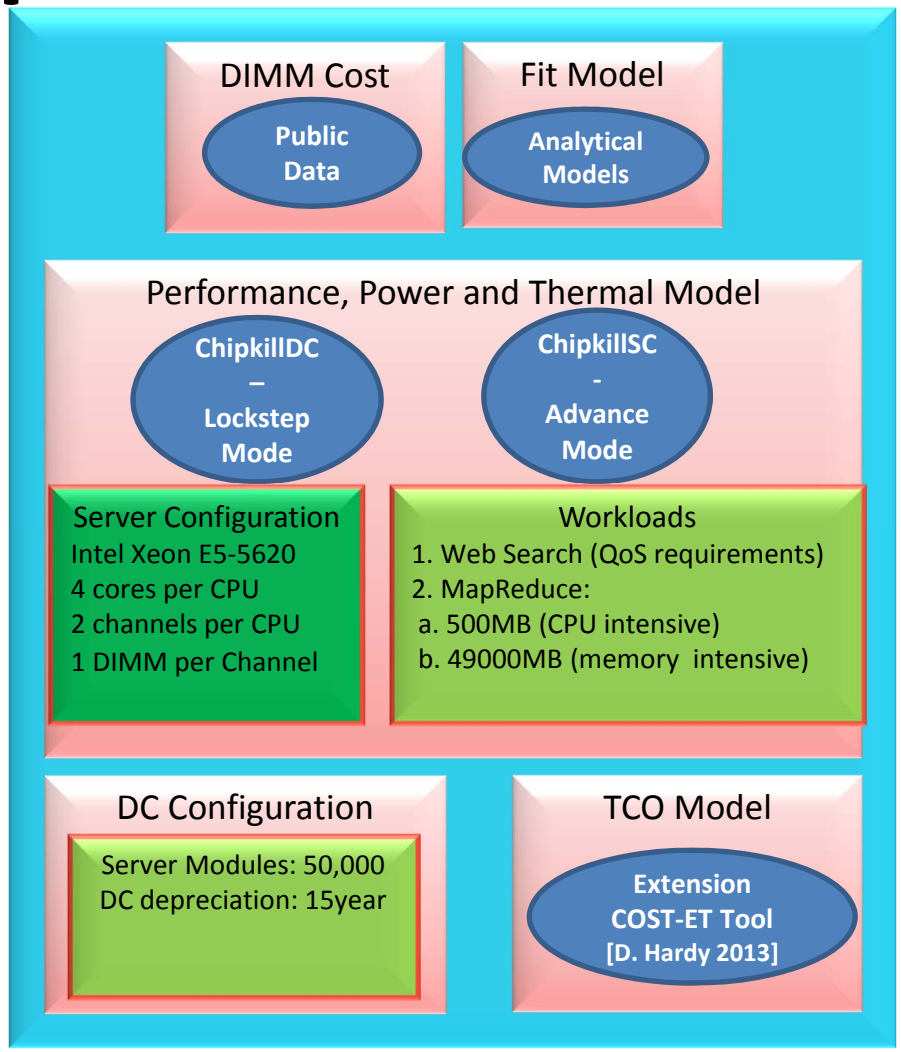

DRAM Protection Implications on Performance

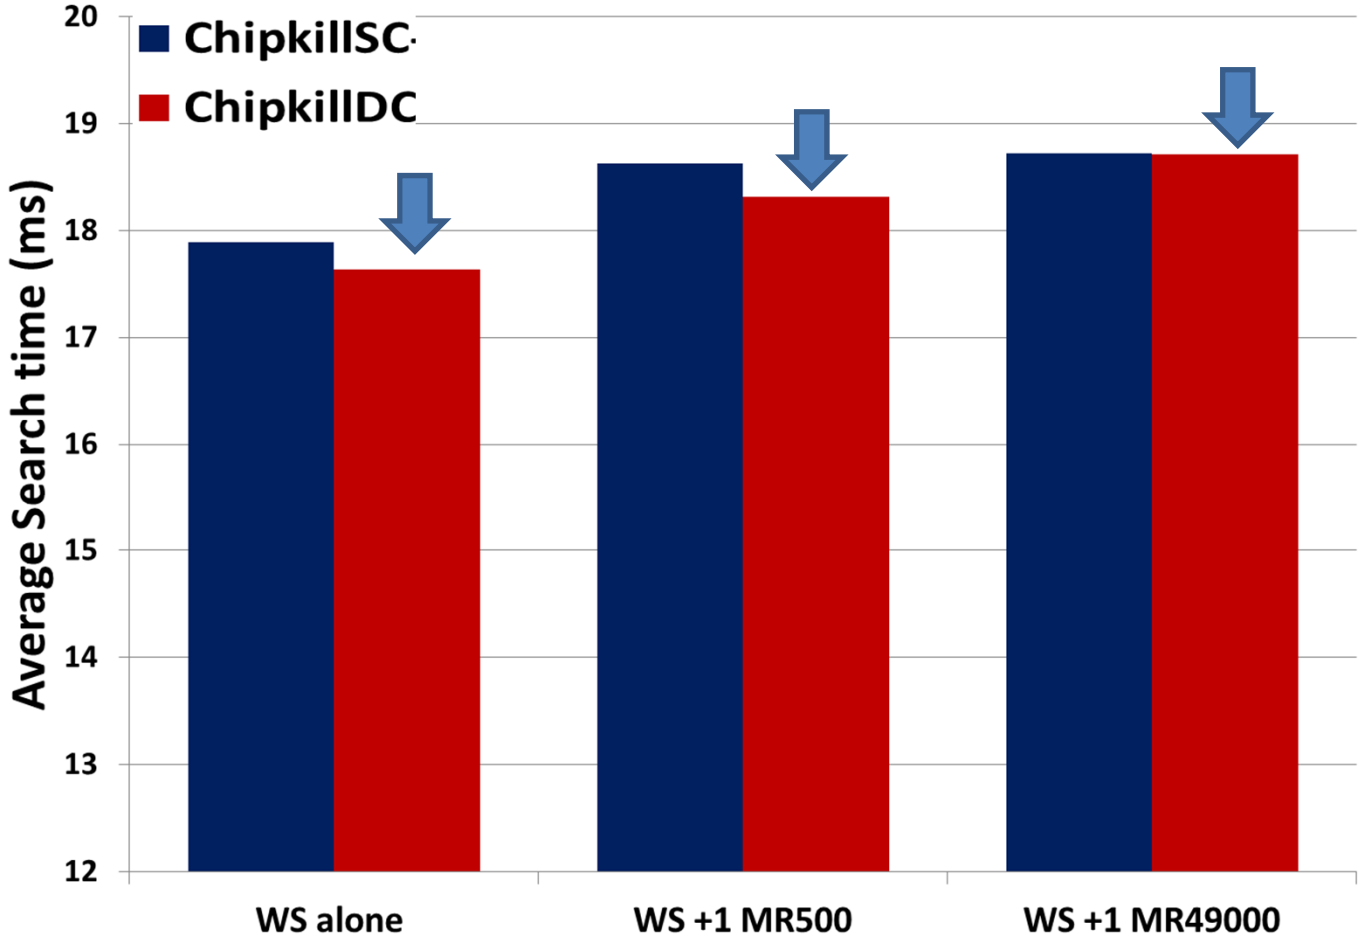

WS: Web Search
MR500: Map Reduce 500MB
MR49000: Map Reduce 49000MB

DRAM Protection Implications on Power

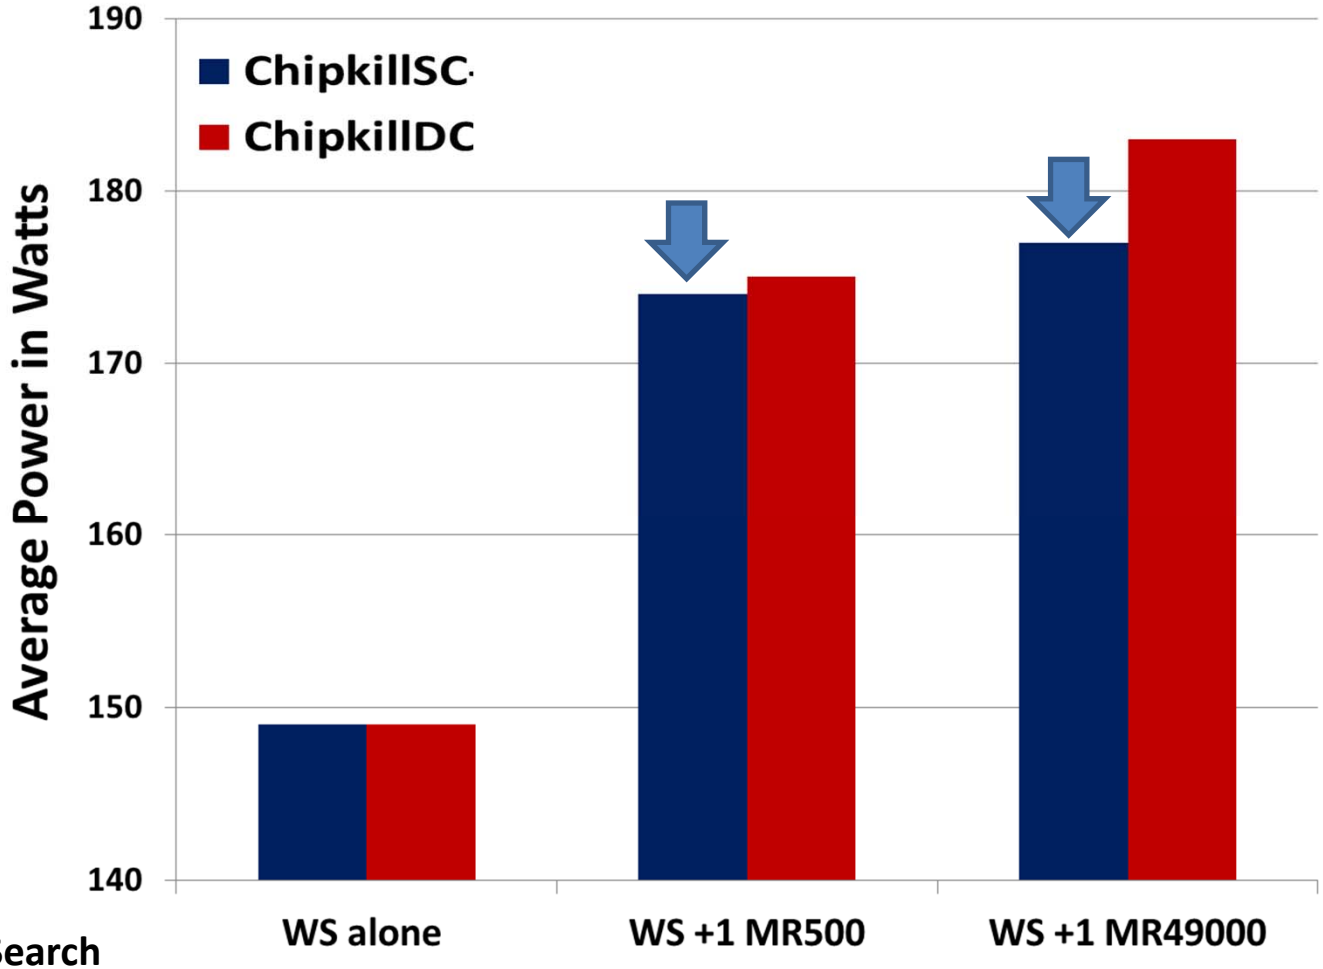

WS: Web Search
MR500: Map Reduce 500MB
MR49000: Map Reduce 49000MB

DRAM Protection Implications on Cost

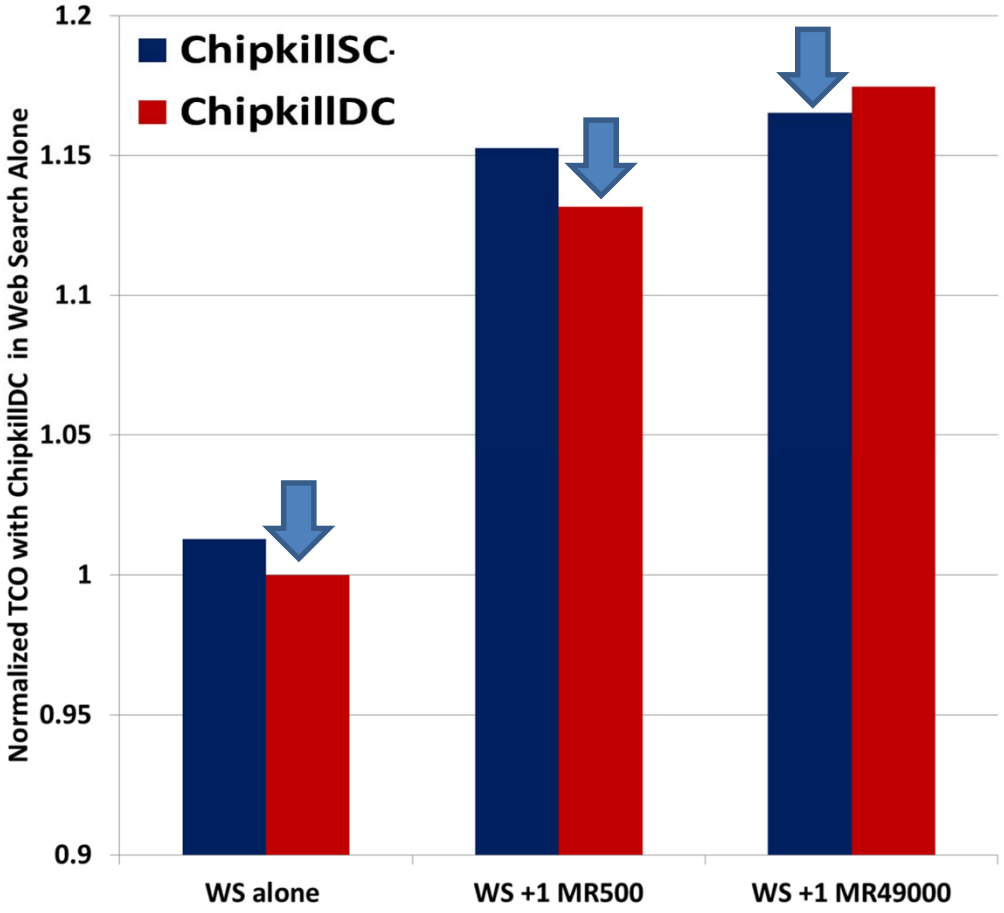

WS: Web Search
MR500: Map Reduce 500MB
MR49000: Map Reduce 49000MB

- **Underlines the importance of understanding the usage and characteristics of all the services to be run in a DC before making memory protection design choices**
- **Highlights the need of proposed framework !!**

Usage

- **Datacenter designers:** Select processor and protection technique

- **Researchers:** Investigate the implications of new ideas related to DRAM failures and DRAM protection techniques

- **Service providers:** Find how to charge for running offline services and to makeup for the increase in TCO due to co-location

More in the paper

- Detailed explanation of each model
- DRAM grades and how affect TCO
- Results for other protection techniques (SECDED)
- Power and performance results for more applications

Conclusions

- DRAM is one of the dominant cost consumers in a DC
- Different protection techniques have different TCO implications
- **Framework** to encapsulates all the parameters and tries to determine the cost-effective protection technique for a DC
- Highlight the need of the framework
  - It is not straightforward to decide which DRAM protection technique is best for a DC setup in the lack of this framework

Future Work

- Evaluate TCO for more online and offline services
- Explore the cost-benefits of new ECC schemes
- Validation of the framework by using detailed logs from a real DC

AMPRA tool download site:
http://www2.cs.ucy.ac.cy/carch/xi/ampra_tco.php