DRAM Errors in the Wild: A Large-Scale Field Study

Schroeder et al.
In Proc. of SIGMETRICS 2009

Presented by:

John Otto

Irene Haque

What is a DRAM error?

- A bit is read differently from how it was written
- Soft errors
  - Transient
  - Caused by cosmic rays, alpha particles,
    - leakage, random noise
- Hard Errors
  - Permanent hardware problem

What are the real-world impacts of memory errors in terms of:

- Temperature
- Memory/CPU Utilization
- Age
- Technology (DDR1, DDR2, FBDIMM)
- Manufacturer
- Capacity
- Chip Size

Contributions and Statistics

- Large scale and longitudinal study
- Measures soft ***and hard*** errors
- Under real-world conditions

Conclusion 1:

*The **incidence of memory errors** and the range of error rates across different DIMMS are **much higher** than previously reported.*

- **Overall Rate**: order of magnitude higher than previously reported
- **8%** of DIMMs have 1+ **CE**/year
- **32.2%** of machines have 1+ **CE**/year
- **1.3%**: 1+ **UE**/year

- **Impact**:
  - ECC is essential.
  - Must plan for individual machines to be offline (UE's)

Conclusion 2:

Memory errors are strongly correlated.

- Top 20% of DIMMs: >94% of all errors
- Past CE's increase future probability of CE

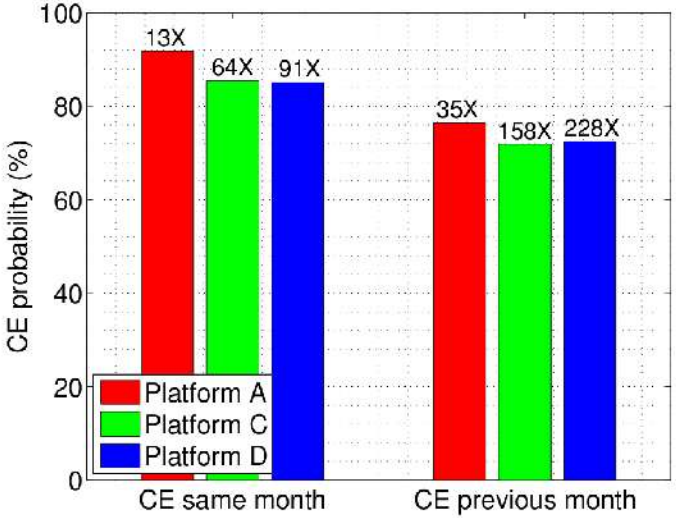

- CE predicts future UE

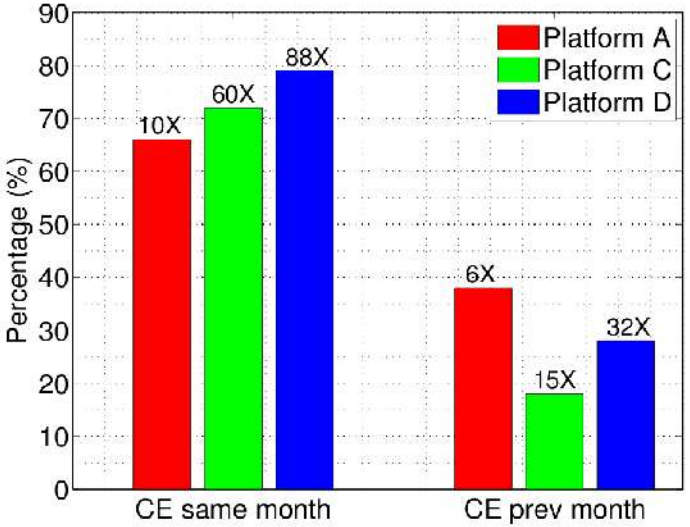

Conclusion 3:

The incidence of correctable errors increases with age.

- # of CE's increases every month

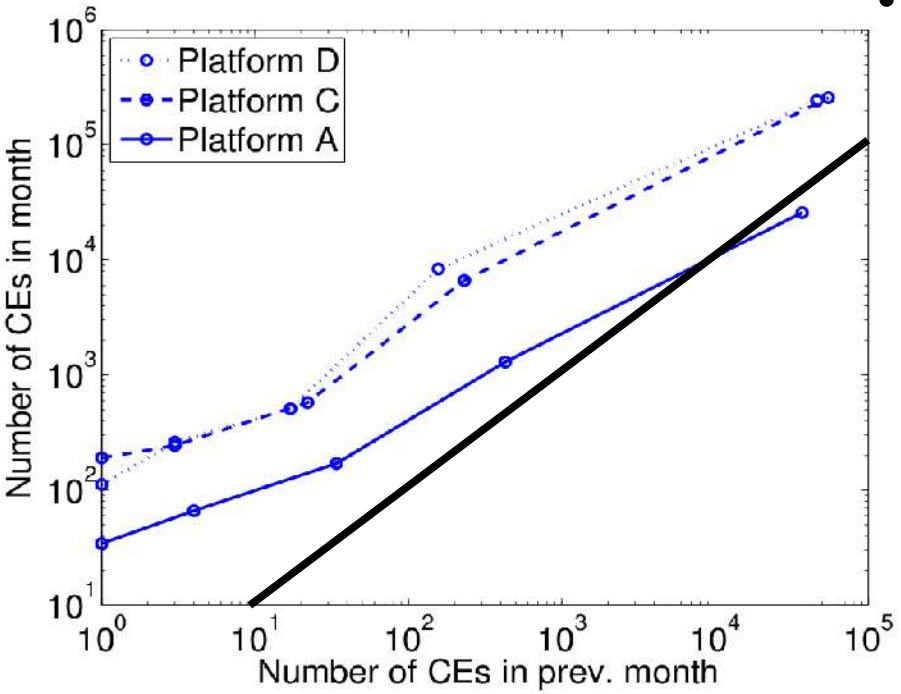

- Error rates increase after 10-18 months
- Low early failure rates indicate burn-in is effective

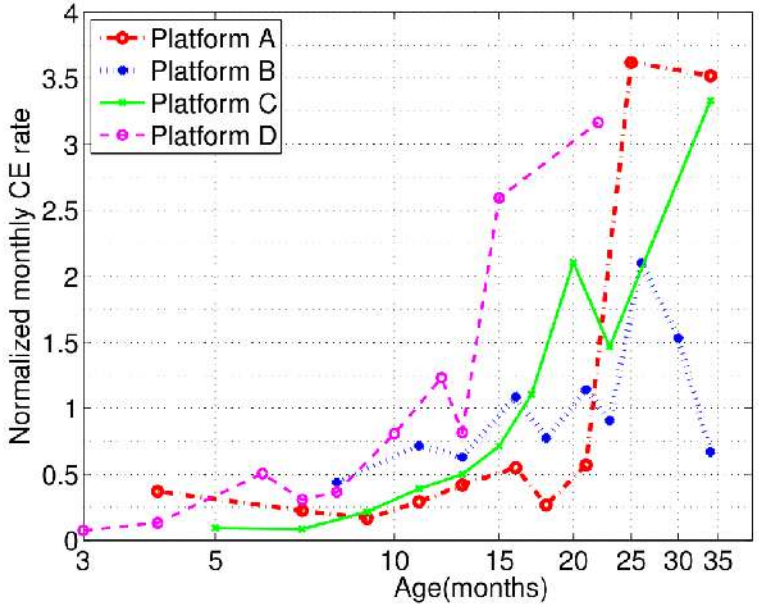

Conclusion 4:

There is no evidence that new generation DIMMs have worse error behavior.

- Newer technologies and higher capacity DIMMS
  - ***Lower*** CE rates compared to older DIMMs.
    - **Older** platforms: **10-15%** of DIMMS have CE's
    - **Newer** platforms: **3-5%** of DIMMS have CE's

Conclusion 5:

Within a certain range, temperature has a surprisingly low effect on memory errors.

- Aggregate:
  hotter = more errors

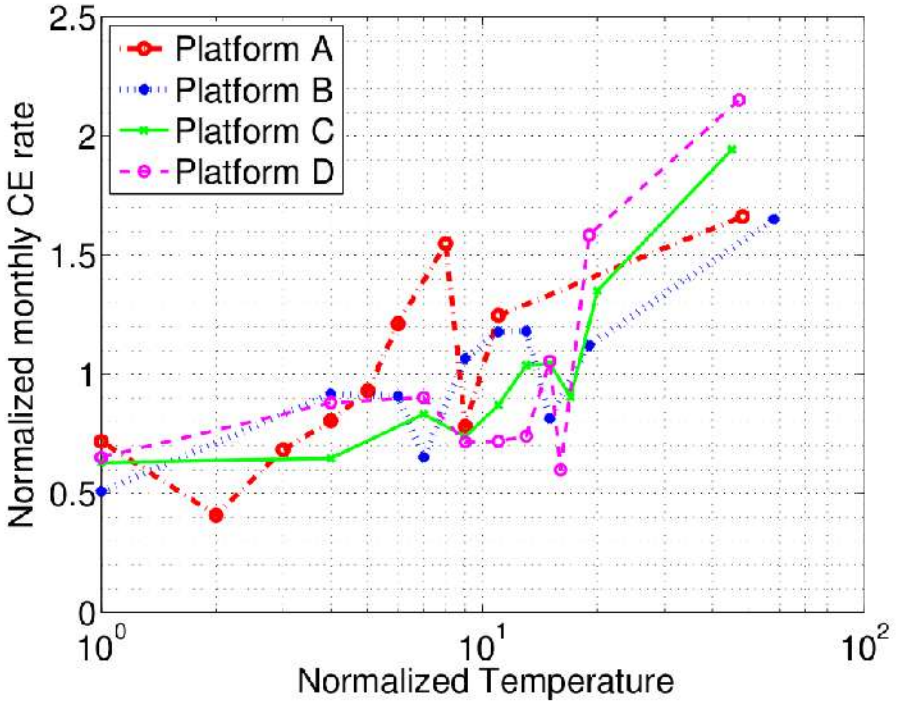

- Accounting for CPU/memory utilization: temperature is not the largest factor

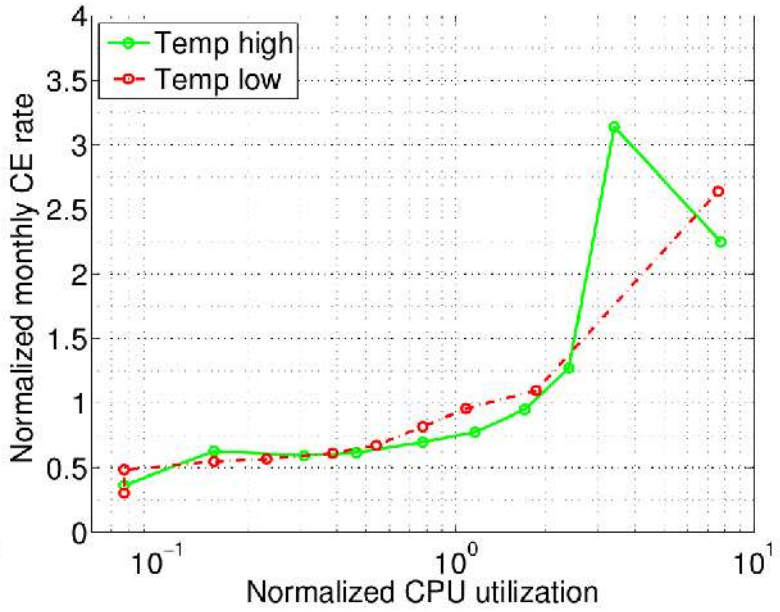

Conclusion 6:

Error rates are strongly correlated with utilization.

- Increased CPU utilization results in more errors

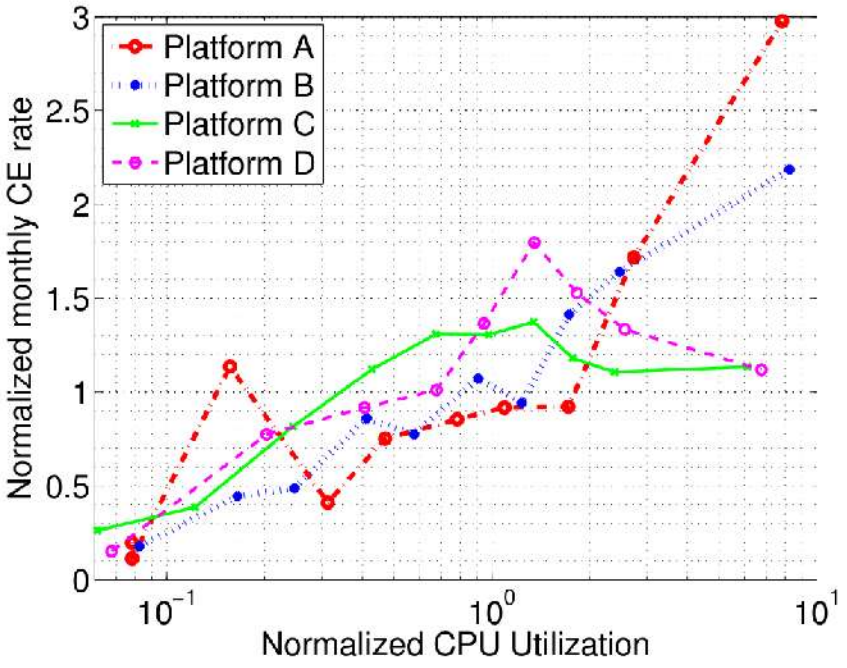

- Across all temperatures, higher CPU utilization gives significantly more CE's

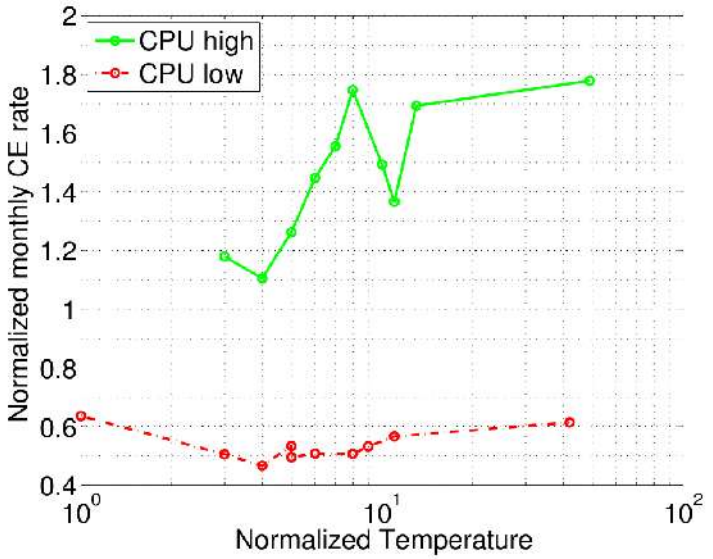

Conclusion 7:

Error rates are unlikely to be dominated by soft errors.

- Soft errors occur randomly
  - External circumstances
- Expect constant error rate regardless of memory utilization

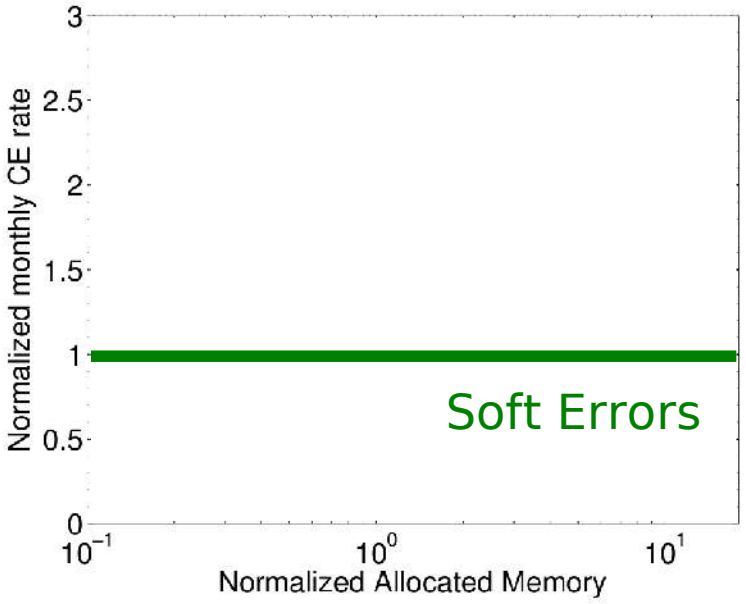

Conclusion 7:

Error rates are unlikely to be dominated by soft errors.

- Hard errors are from flaws in the DIMM
- Increased memory utilization => flaws more frequently revealed

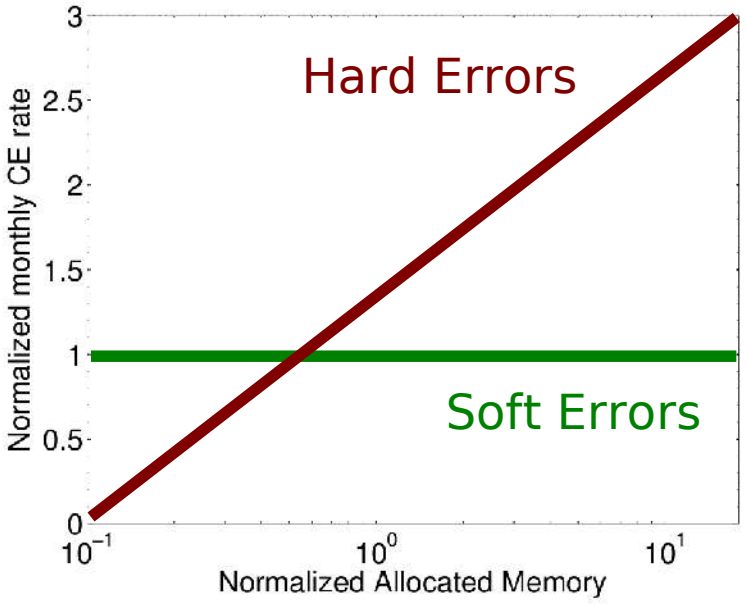

Conclusion 7:

Error rates are unlikely to be dominated by soft errors.

- Memory utilization correlates with CE rate
- Hard errors are dominant

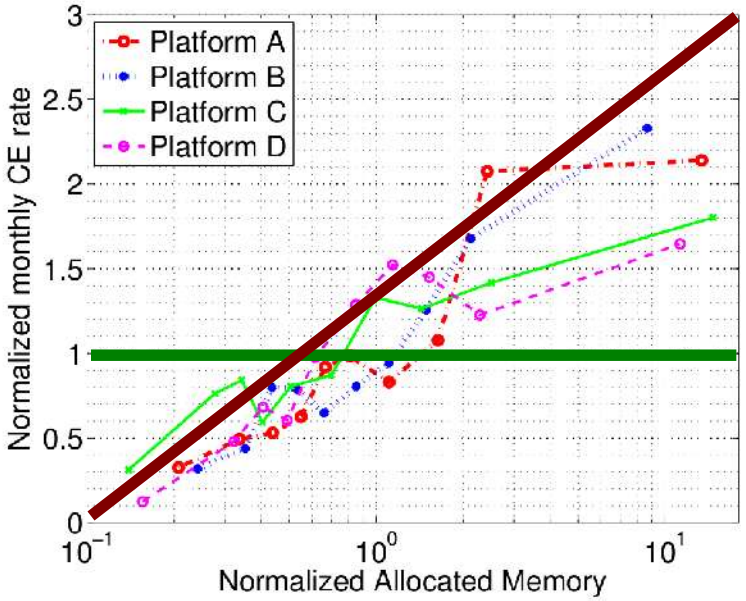

Benefits of the work

- Error rates are not worsening with the new generation DIMMs
  - Good news for chip manufacturers!
- But *real-world error rates are higher* than previously reported

- **Awareness** of real-world memory error rates

Benefits of the work

- **Awareness** of real-world memory error rates
- Enables better forecasting:
  - Predicting actual operating costs
    - DIMM replacement
    - Redundancy planning for failures
  - Enables better trade-offs between cost and technology
    - Better ECC gives lower aggregate downtime