THE UNIVERSITY OF TEXAS AT AUSTIN

Toward Exascale Resilience

Part 7:

What about clouds?

Mattan Erez

The University of Texas at Austin

July 2015

Very large systems

- Very distributed computing
- Not typically single cohesive problems

Main cloud concerns

- Availability
- Management
- Sharing

Main cloud applications

- Storage
- Search
- Localized analytics

- Many independent requests
- Throughput matters, but **latency caps**
- Limited communication
  - Very relaxed consistency

Main exascale applications

- Cohesive computations
- Full-scale analytics
- Checkpointing

- Few (one) large application
- Throughput paramount

- Heavy communication
  - BW and latency important

Resilience implications

- Interrupts much less critical in the cloud
- Better hardware not good enough for customer storage anyway
- Most computation is approximate already

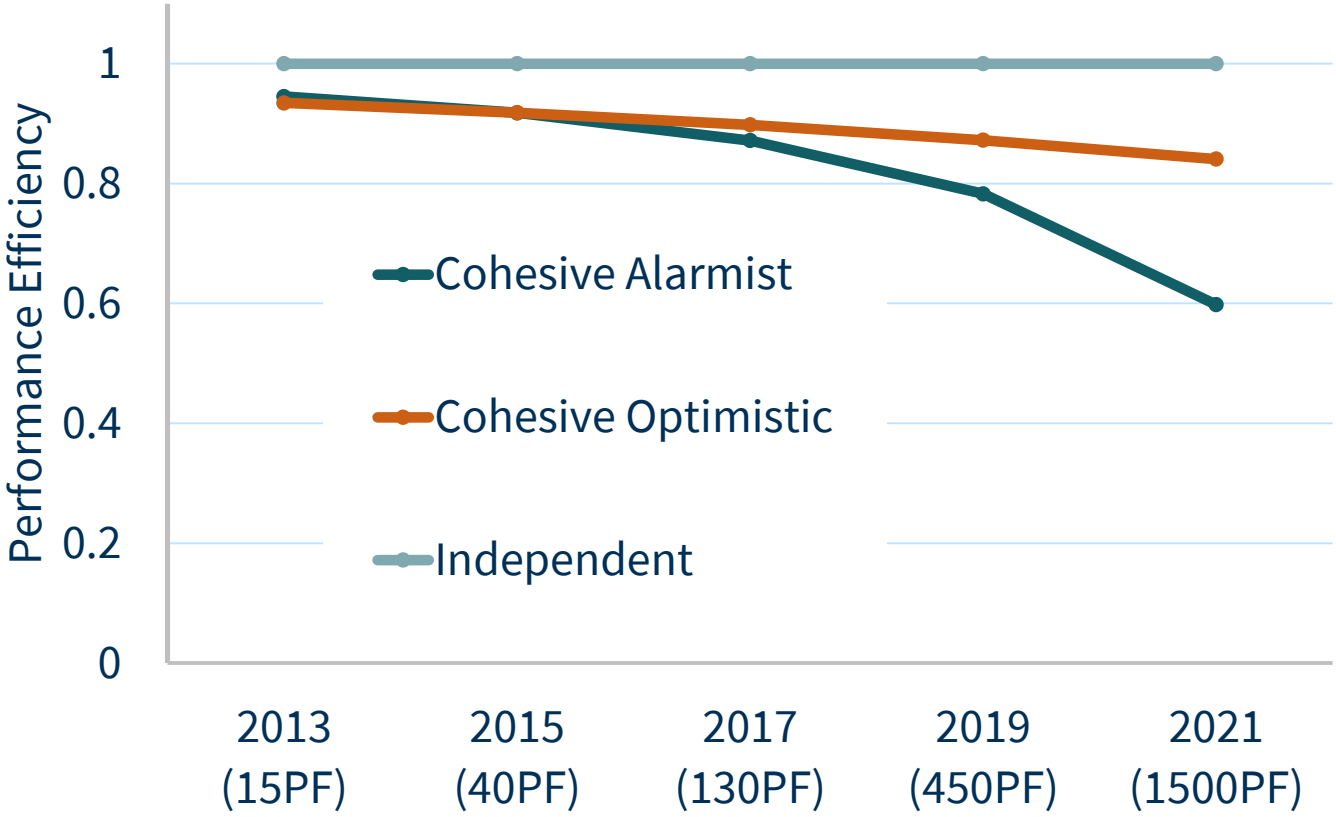

Today, similarities abound

- Xeons, mostly
- ECC memory
- Dense packaging

But very important differences

- Virtualization vs. bare metal
  - Lots of performance lost
  - Aavailability and management gains
  - Not really different hardware, but shows focus not perf.
- Ethernet vs. specialized interconnects
  - Differences in workloads
- Limited throughput-computing in the cloud
  - Only in special-purpose components (groups of racks)
    - Because latency actually important
  - Exascale will need everywhere

Possible cloud/exascale divergence

- That will not be good
- Likely scenario: exascale will have to make use of cloud-oriented processors

- Many differences are really software stack

What can we learn from the cloud?

- Make jobs look more independent
- Manage jobs and sharing
- Be ready to approximate

Making jobs look independent

- Switch algorithm
  - E.g., Monte Carlo
- Switch programming model
  - Tasks based data flow
  - Transactional approaches
- Utilize hierarchy
  - Containment domains

Manage jobs and sharing

- Preemptive migration
  - Failure forecasting
- Incremental and non-blocking (overlapped) CP
- Alternatives to global coordinated recovery

A word on system noise

Prepare to approximate

20MW / 1 exa-FLOP/s

Energy ≤

20pJ/op

50 GFLOPs/W sustained

Best supercomputer today: **~300pJ/op**

I/O

Arithmetic

Comm.

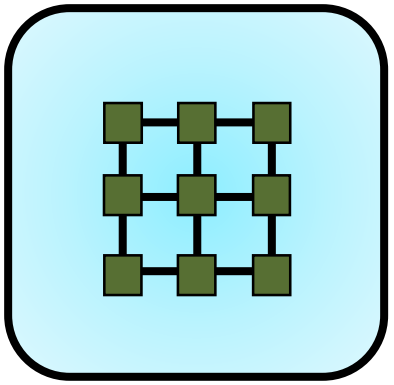

Actual processing

Control

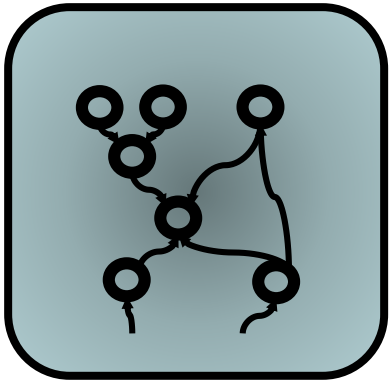

Memory

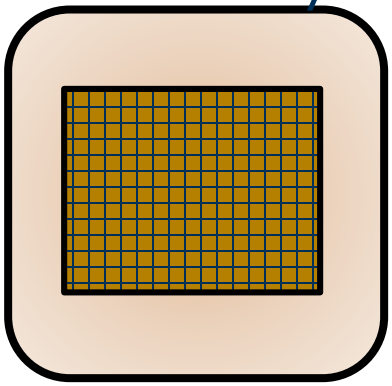

Idle/margin

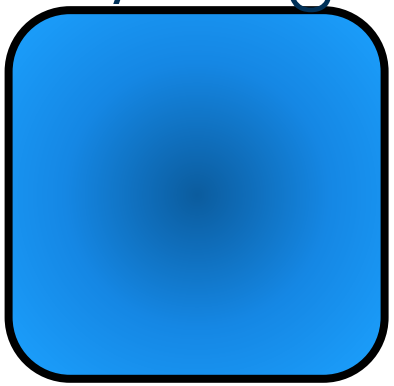

How much of each component?

Arithmetic

64-bit floating-point operation

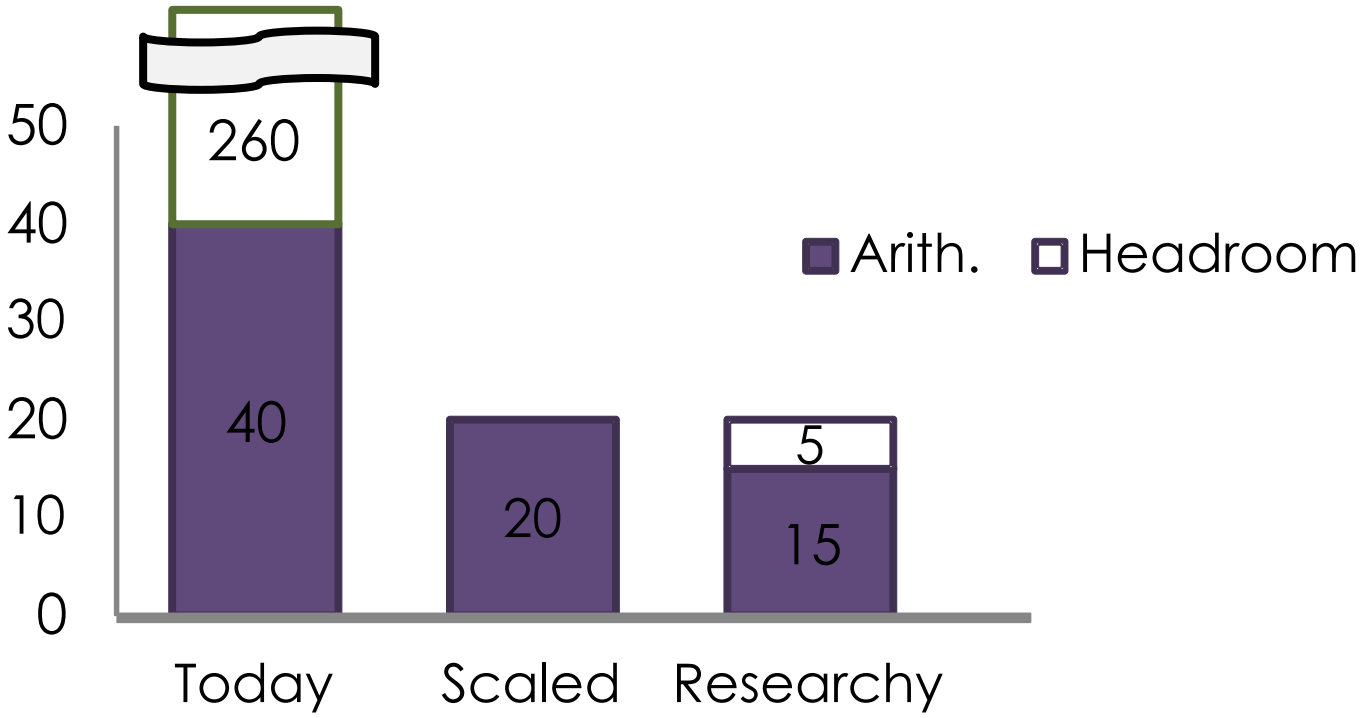

Rough estimated numbers

Enough headroom?

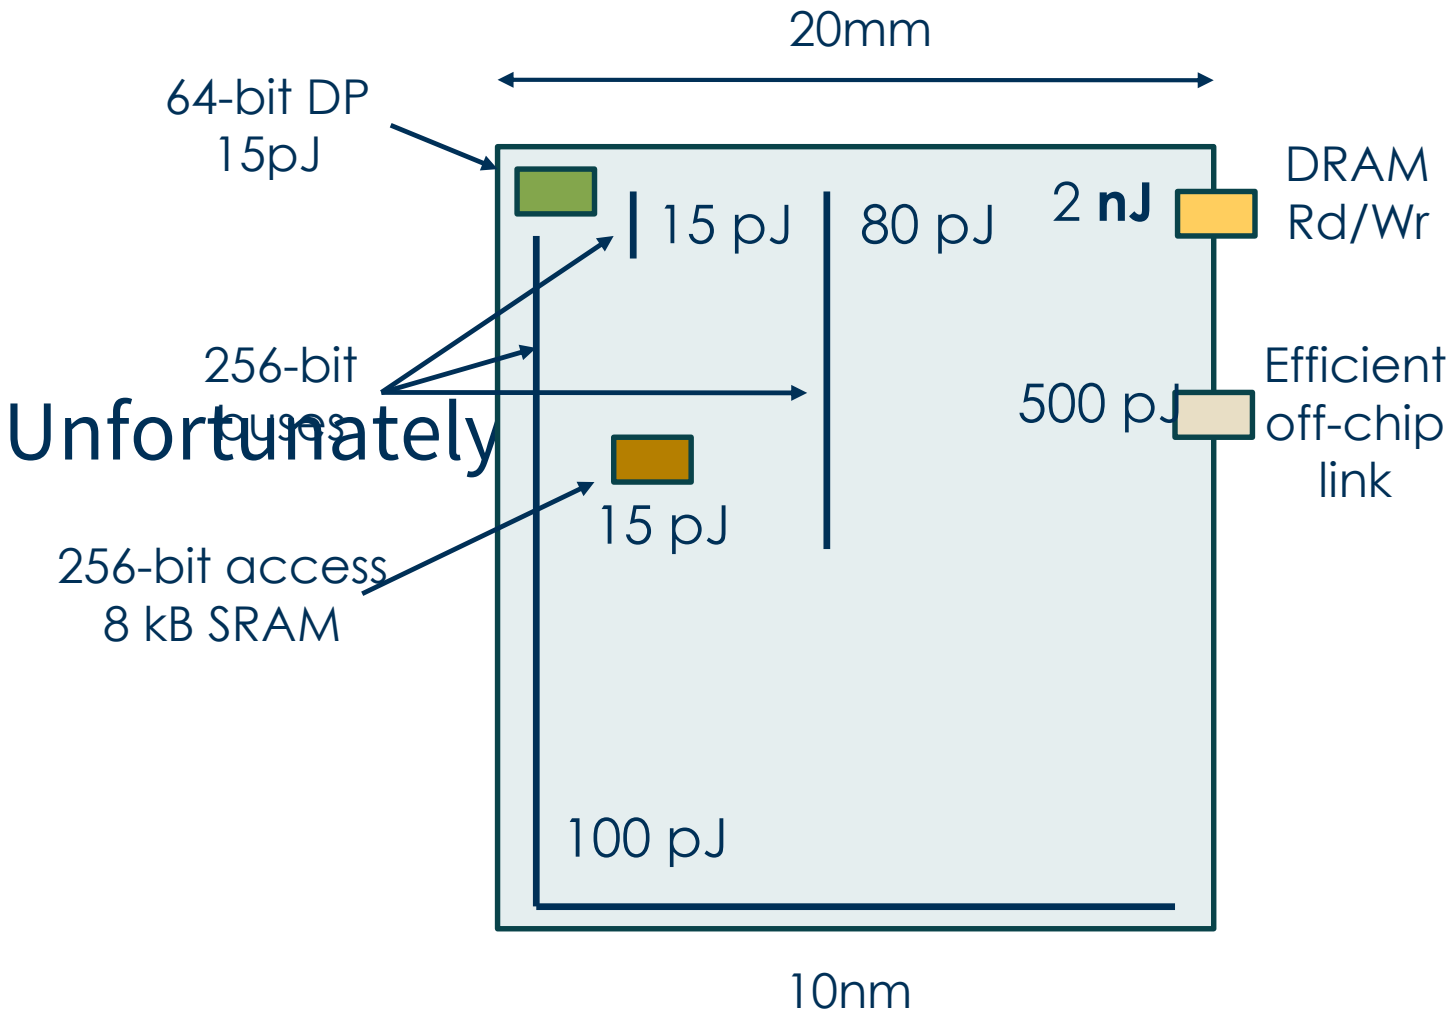

Unfortunately

Need **more headroom**

– Minimize waste

Do we care about single-unit performance?

Must all results be equally precise?

Must all results be correct?

Lunacy?

Relaxed reliability and precision

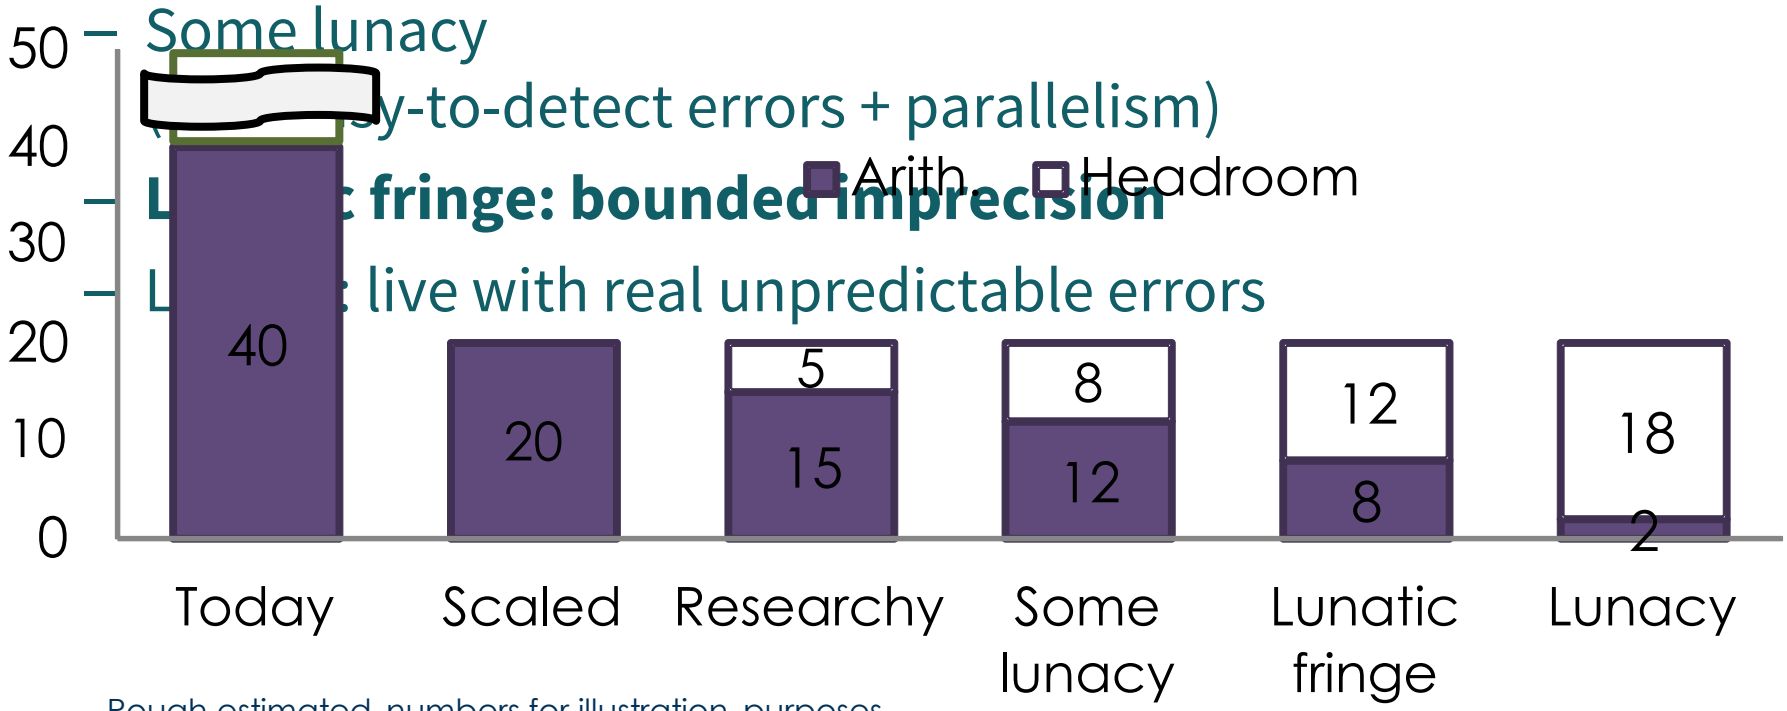

Rough estimated numbers for illustration purposes

Bounded App[illegible]n

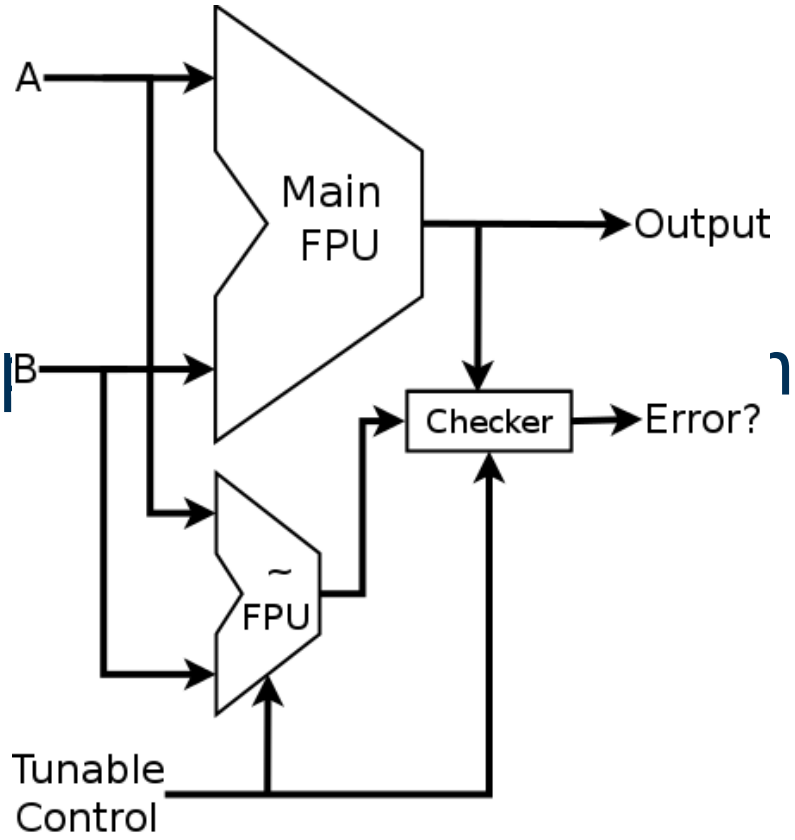

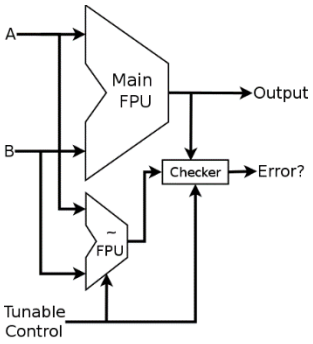

Bounded Approximate Duplication