

Exascale Computing Roadblocks: Scale, Energy, Resilience

Marc Snir

Director, MCS Division, Argonne National Laboratory

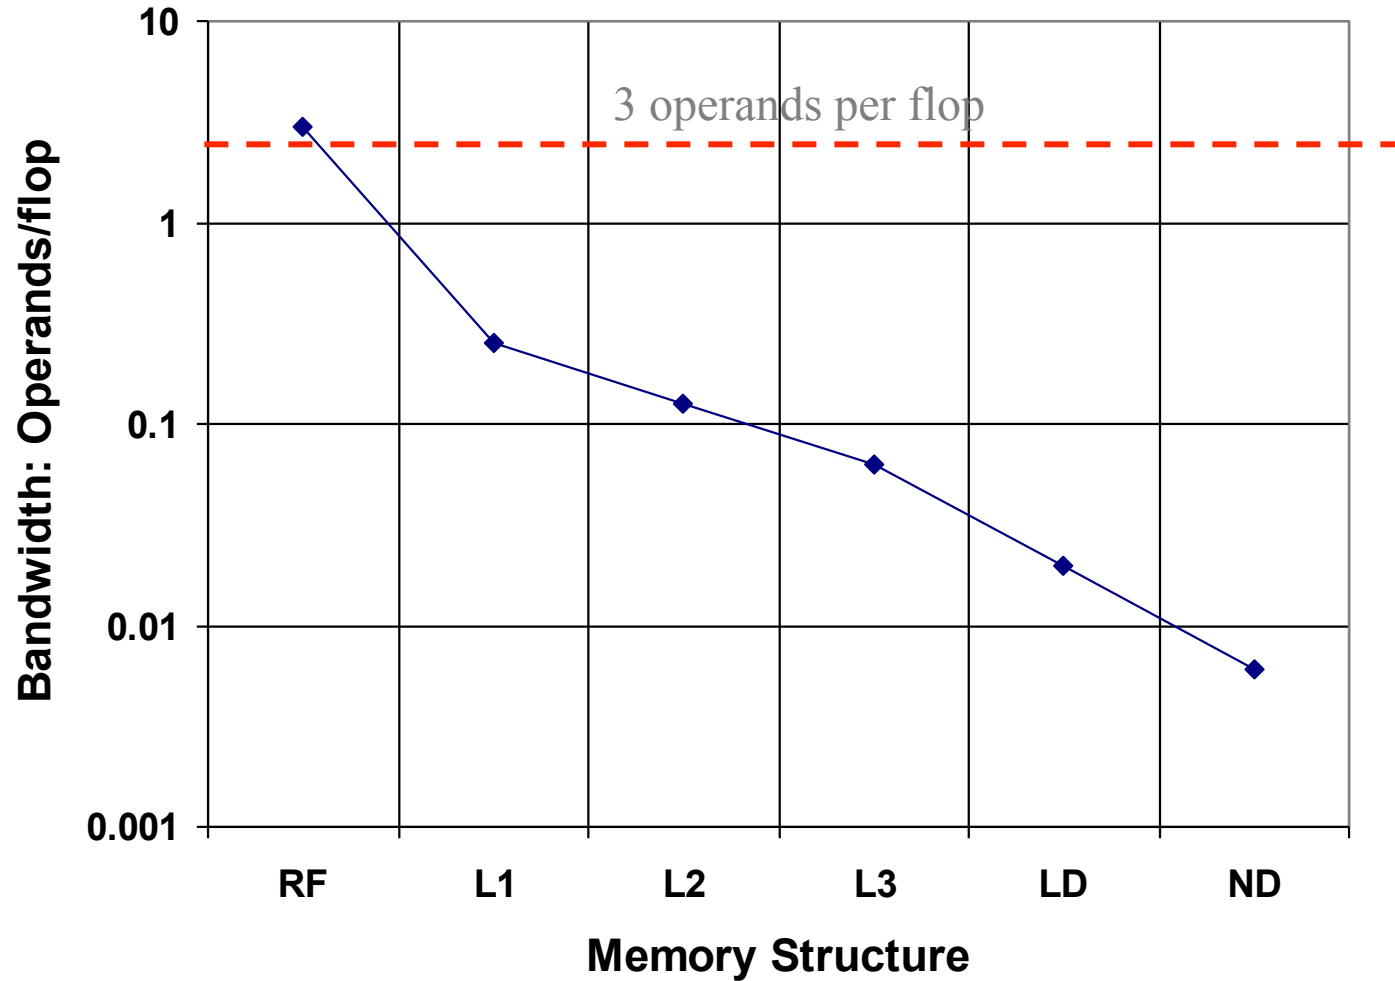
Professor, Dept. of Computer Science, UIUC

Where technology goes

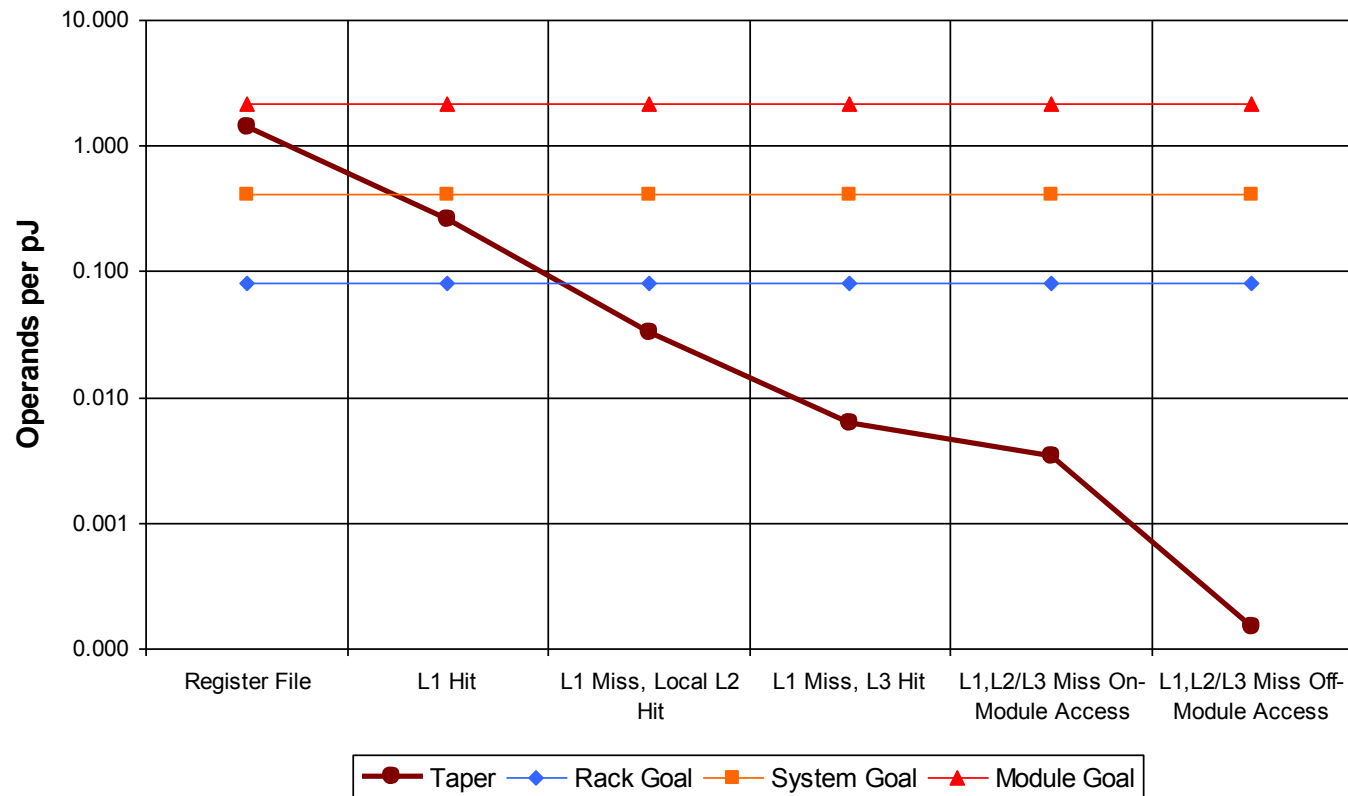
- Clock speed is not increasing
 - Performance increase = increase in parallelism
- Energy consumption is major performance bottleneck
 - Due to cost of communication
- Faults are becoming more frequent as transistors shrink
 - Latches in CPUs are vulnerable
 - 2 bit upsets more likely
 - Silicon ages much faster



Bandwidth Taper [2015] (courtesy Kogge)



Energy Taper Results [2015] (courtesy Kogge)



- < 10 pJ per flop core consumption (10 MW per exaflop)
- > 475 pJ per flop total energy consumption (500 MW per exaflop for Linpack)
- *It's all about communication!*

Evolution of fault rate



Exascale (in ~2020)

- Need to manage 100M – 1B concurrent threads
- Need to reduce communication
- Need to build reliable systems from unreliable components



Communication Theory

- On-chip communication model (back to VLSI complexity?)
 - Off-chip (optics) is easier – cost is constant per hop
- Communication complexity results – especially for numerical computations
 - Relation between physics of simulated system and communication
- Is there a trade-off between compute time and energy consumption?
- Can one save communication by recomputing?



Resilience Theory

- Tradeoff between frequency of failures of components and total computation time and/or energy consumption
 - Assuming fail-safe components
 - Assuming undetected soft errors (bit flips)
- Model 1: nodes are failing with certain probability. How fast can one compute?
- Model 2: data bits are flipped (in CPU) with certain probability; storage is stable. How fast can one compute?
 - Assume “other bits” are well protected
- Back to von Neumann?



Thank You