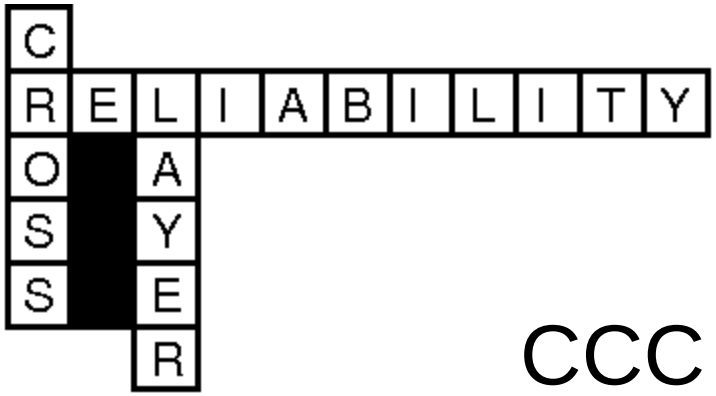

CCC Visioning Study:

System-Level Cross-Layer Cooperation to Achieve Predictable Systems from Unpredictable Components

www.relxlayer.org

Executive Summary

- Continued scaling → unpredictable components
- Traditional solutions are too expensive
  - **Key problem:** spend energy, but designs energy limited
- Cannot solve efficiently at any single level
- Opportunities are cross layer, exploiting information
  - Memory systems offer inspiration
  - Logic: bigger challenge, but big payoff
- **Our goal:** community consensus on
  - Potential vision
  - What can be done
  - Research vector: make-a difference research questions and solution capabilities that help realize the vision
  - Articulate for funders, congress, lay public

Where are we?

- Past:
  - March meeting: get started, turn up ideas and issues
  - Constituency groups: identifying challenge
  - Focus groups on Roadmapping and Metrics
  - July meeting: challenges, metrics
  - Draft outline of vision story
- Now:
  - pull together remaining challenges, roadmap, and formulation
  - Refine vision outline
- Next:
  - Digest/distill recommendations → report
  - Communicate and educate funders
  - Continue to invigorate and organize community

What's a Vision Story

- What are we doing?
- Why now?
- How done today?
- Trends?
- What can we offer? Accomplish with investment?
- What's new?
- Why?

- Science needed? (priorities)
- Mission Impact?
- Metrics and Goals?
  - Measure and manage programs?
- Examples and Scenarios
- Organize research?
- Slogan(s)/simple story?

Roughly outline for rest of introduction

What trying to do?

- Allow continued scaling benefits
  - Reduce energy/operation
  - Reduce $$/gate
  - Increase ops/time with limited power-density budget
  - Avoid stalling this engine of economic growth
- While maintaining or improving safety
- Navigate inflection points in
  - Energy and reliability

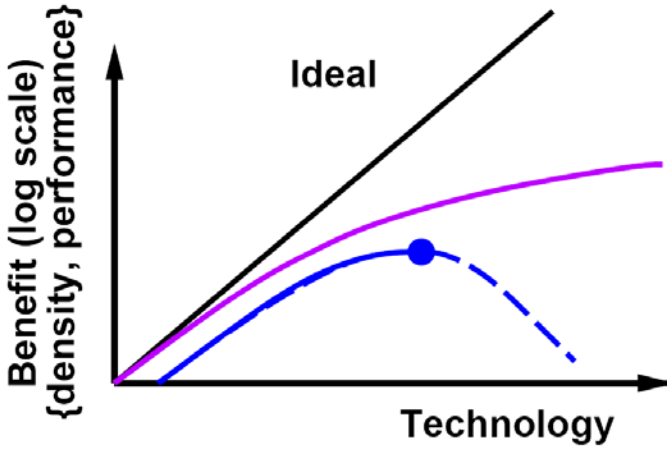

Why Now?

Increasing Error Rates

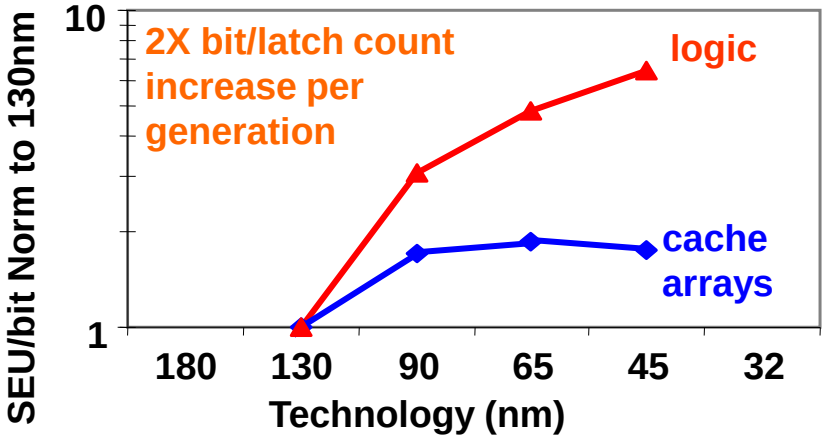

Power Limits Integration

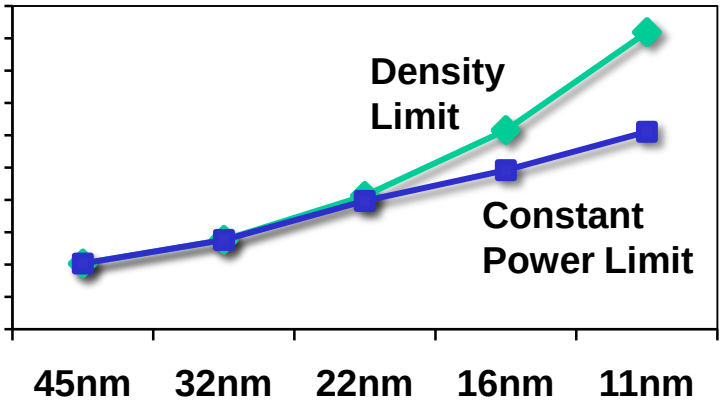

Variation ↑, Predictability ↓

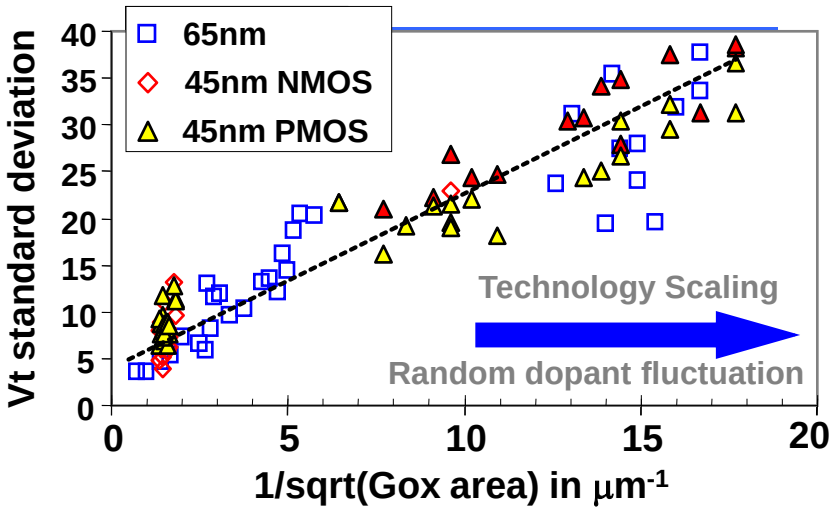

Guard Bands Hurt Power, Perf.

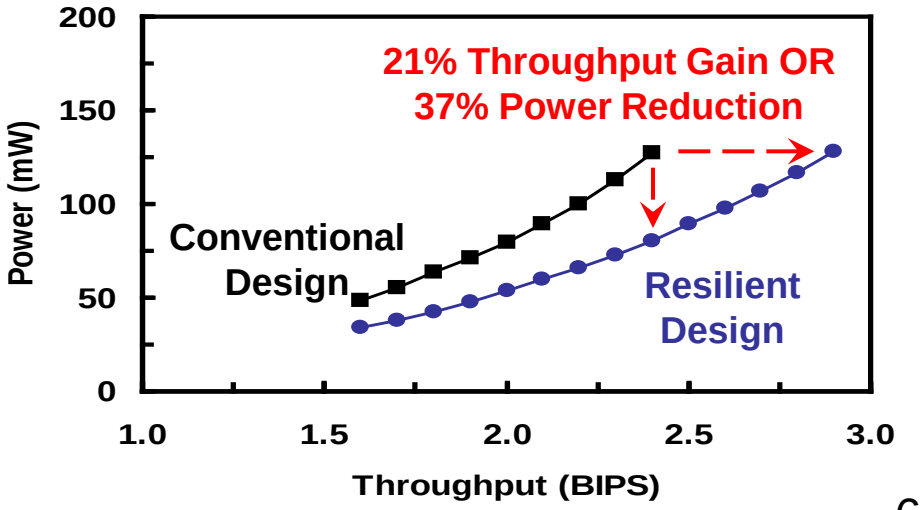

Why now?:
End of Energy Scaling?

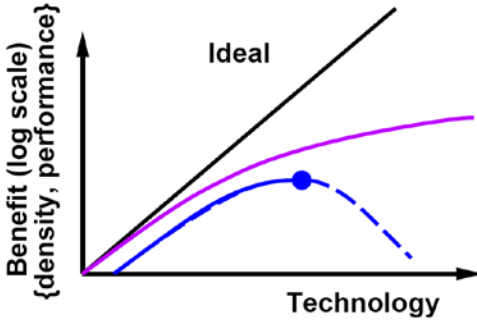

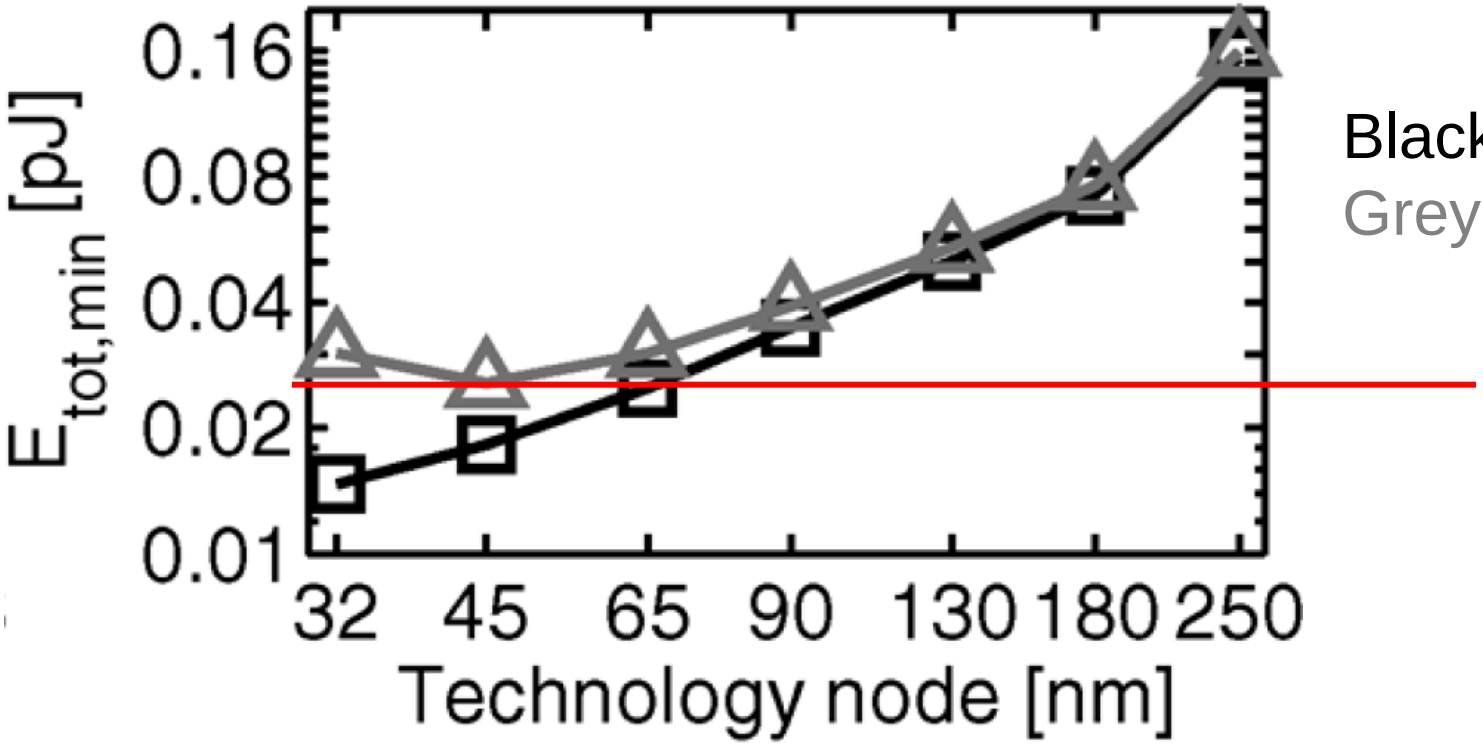

Black nominal

Grey with variation

[Bol *et al*., IEEE TR VLSI Sys 17(10):1508—1519]

How is it done today?

- Demand reliable, consistent **device** operation
- Margin for worst-case device effect
  - Of billions, over multi-year lifetime
- Discard components when devices fail
- System-level redundancy
- Niches where above “not good enough” are small but important (avionics, medical)
  - Spend considerable $$, energy for reliability
    - *E.g.* Brute-force replication
    - Many-year performance lag behind commercial systems

Trends?

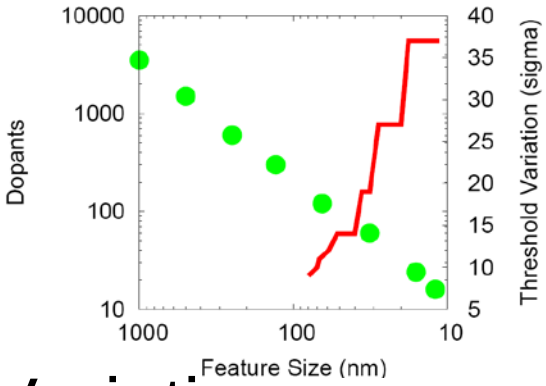

- Power-density limited components
- Fewer dopants, atoms → increasing Variation
  - Must margin over wider range of devices
- Increasing Transistors / chip
  - More things to go wrong, sample extreme devices
- Increasing System Sizes
  - data centers, supercomputers
- Decreasing critical charge
  - Increasing upset rates
- Decreasing opportunities for burnin
- Increasing wear-out effects
- Computations increasingly deployed into critical infrastructure and life-critical roles

More to come from Roadmap Group.

What can we accomplish?

- Build reliable systems from unreliable components
  - Efficiently compensate for unpredictable devices through cooperation at higher levels of system stack
  - Replace energy margins with information margins
- To quantify: how much more efficiently (less energy, less $$) can we make it?

What's New?

Ubiquitously/pervasively exploit:

1. Design prepared for repair
2. Cooperative filtering of errors at multiple levels
3. Cross-layer codesign --- Multi-level tradeoffs
   - Generalization of hardware/software
   - Tools to automate
4. Strategic redundancy
5. Differential reliability
6. Scalable and adaptive solutions

Hints of these abound, but as point solutions
rather than systematic approach.

Why do this?

- Allow scaling to continue without sacrificing safety
  - Continued reduction in energy/op
  - Continued reduction in $$/op
  - Maintain or extend component lifetimes
  - How much further?
- Reliability matters for everything
- Allow construction of larger, dependable systems
- Make infrastructural technology worthy of the trust we place in it → save lives
- Eliminate conservative overdesign
  - Close gaps in processing capabilities in areas of critical need
  - Without sacrificing reliability
  - While containing human burden
- Challenges areas of pain from two focus groups today --- roundup tomorrow

Big Science Questions (Priorities)

1. How do we organize, manage, and analyze layering for cooperative fault mitigation?
2. How do we best accommodate repair?
3. What is the right level of filtering at each level of the hierarchy?
4. Can we establish a useful theory and collection of design patterns for lightweight checking?
5. What would a theory and framework for expressing and reasoning about differential reliability look like?
6. Can a scalable theory and architectures that will allow adaptation to various upset rates and system reliability targets be developed?

Mission Impact

- **Aerospace:** near-commercial performance in space (order of magnitude gain in perf./{size,weight,power}
- **Supercomputers/Datacenters:** continue scale up machine size
  - Bigger science, tame more data
- **Security:** prevent compromise by physics
- **Cyberphysical:** manage safety in life and economically critical roles
- Others?

Metrics, Goals, Measure and Manage Progress

- Energy/Op at noise rate + perf. target
  - Noise rate: defects, variation wear, transients
- Post-fab adaptability to range of noise rates
- Timeliness and quality of adaptation
- Refined by metrics focus group

Slogans, Soundbytes, Promises

- How communicate (market) this succinctly to stakeholders?
  - Funders, students, congress, public
  - In a suitably positive way
- What does success (continued scaling) buy them?
  - Sexy things consumers want
  - Capabilities mission agencies want

How use our time this meeting?

Today

- Constituency Group Challenges
  - Life Critical
  - Infrastructure
- Roadmap update
- Research Organization
- Breakouts
  - Constituency, slogans/promises
- Get organized on Final Report

Tomorrow

- Common challenge roundup
- Working time
- Strategy discussion