



# Rethinking Resilience

# The Future is:

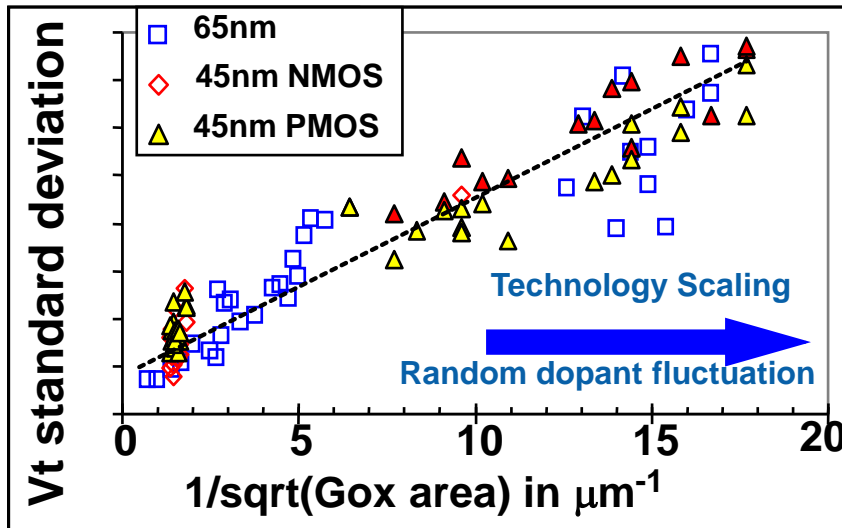
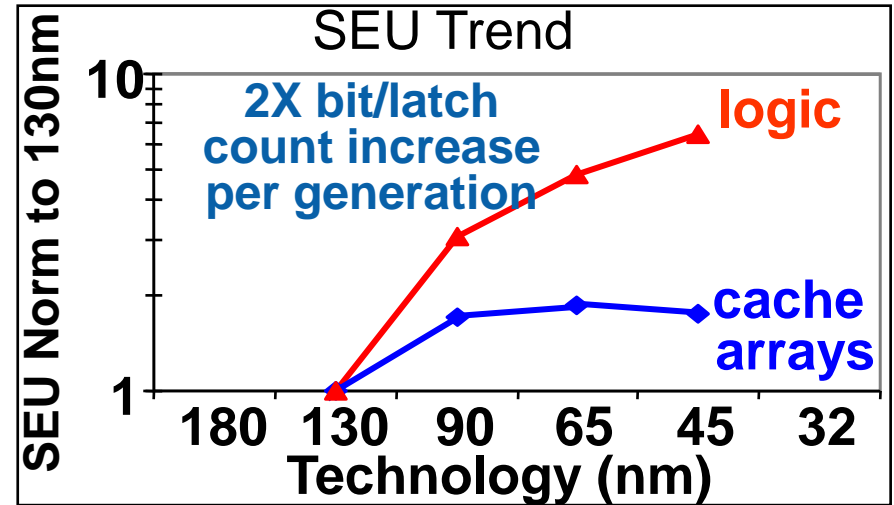
- High variation
- Significant fault rates
- Continued feature size scaling
- Slower Vdd scaling
- Slower clock rate scaling

How do we use all the transistors we can build?

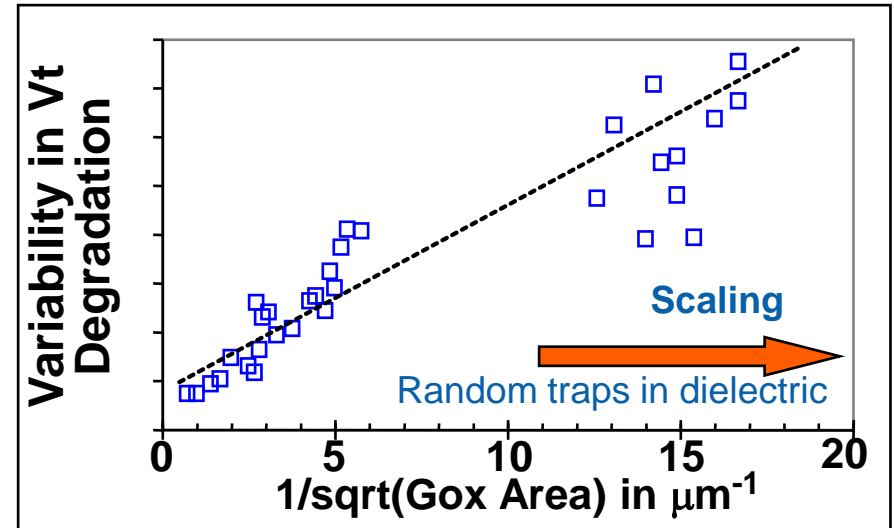
**Resilience!**

# Preaching to the Choir

- Fabrication variation, aging, soft error rate *per chip* all getting worse with scaling



**T=0 Variation Increasing**

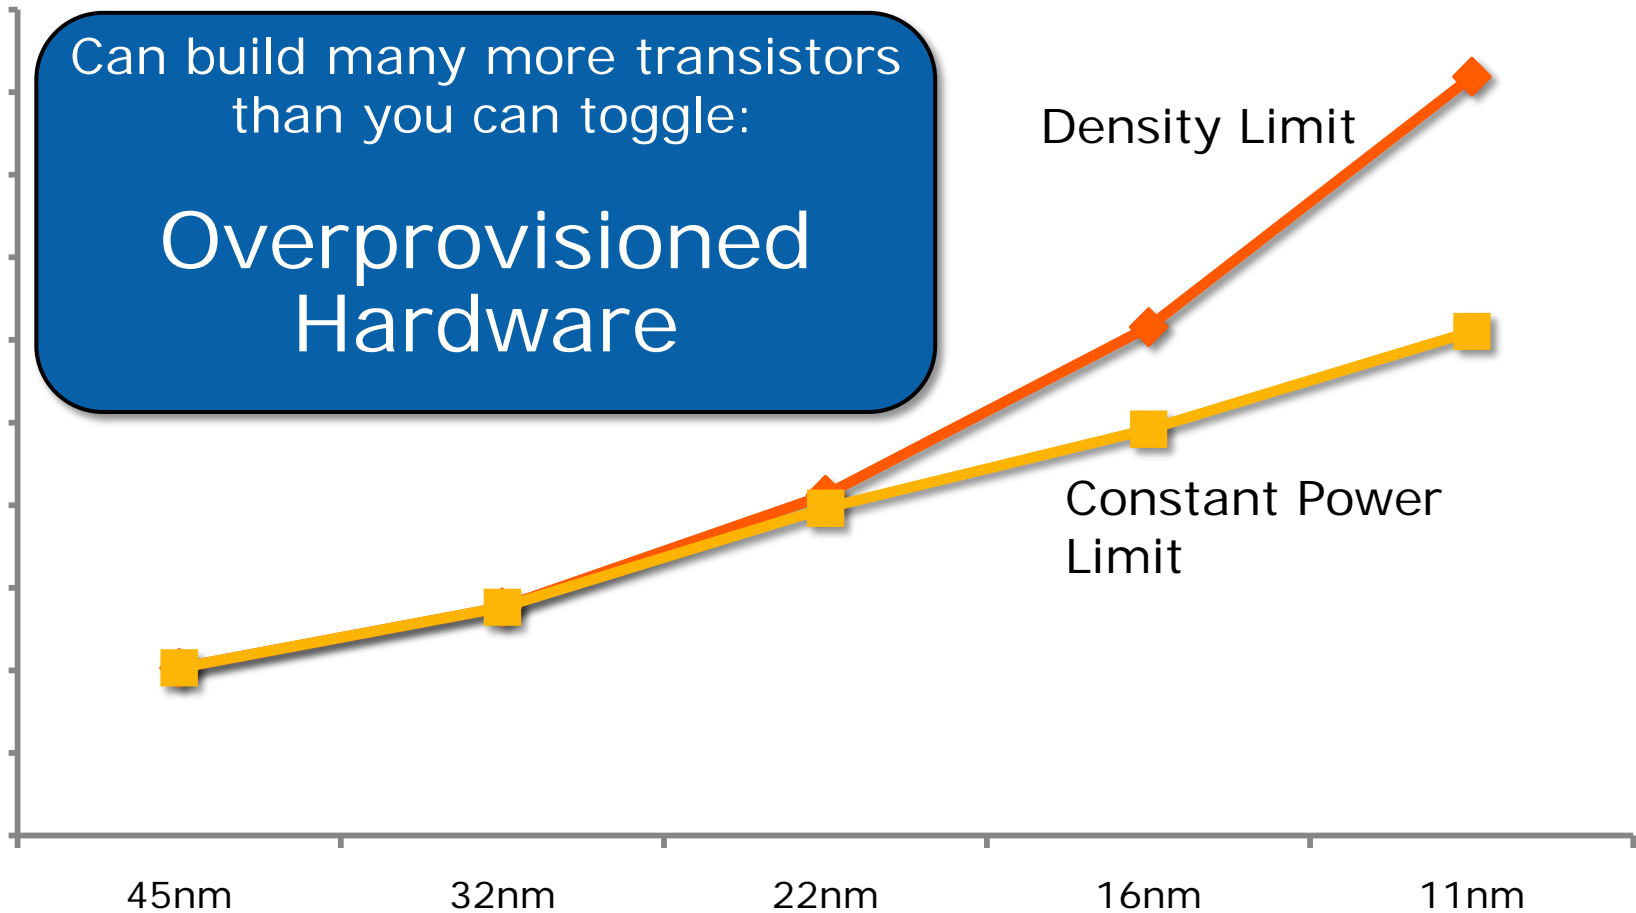


**Aging Causes More Variation (T=EOL)**

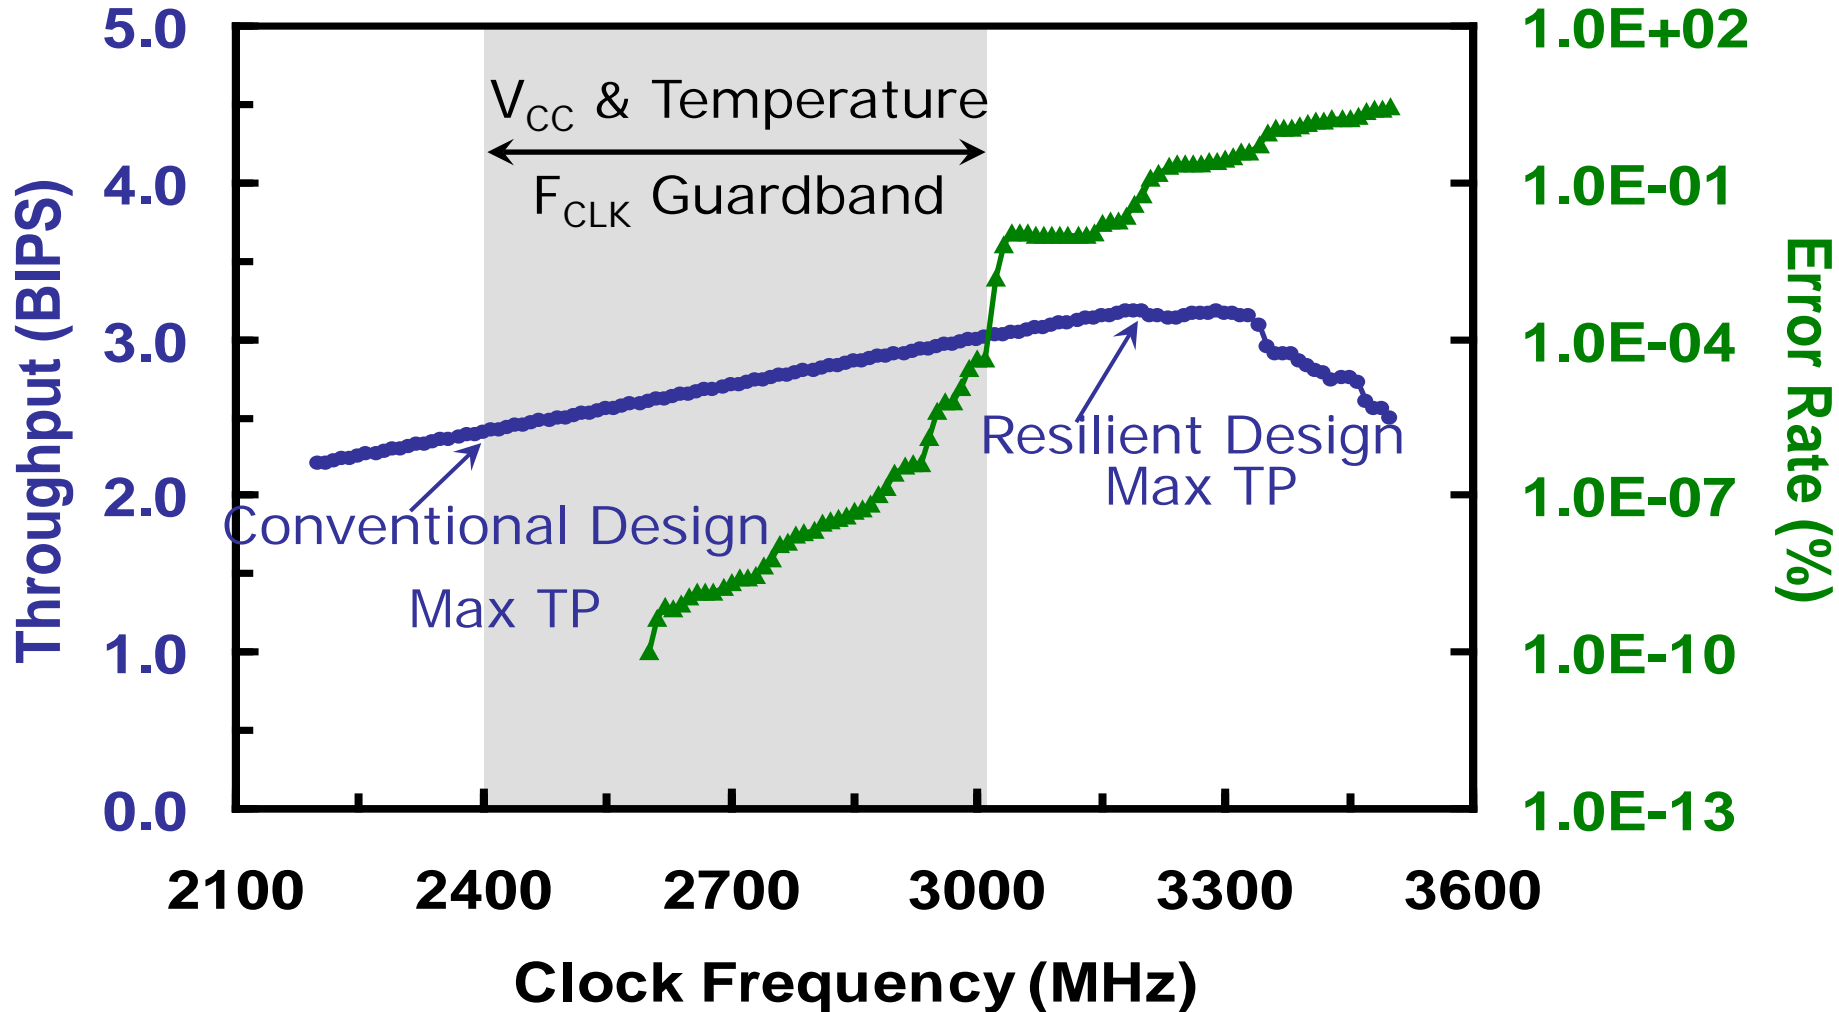
# Technology Outlook

High Volume Manufacturing	2008	2010	2012	2014	2016	2018	2020	2022
Technology Node (nm)	45	32	22	16	11	8	6	4
Integration Capacity (BT)	8	16	32	64	128	256	512	1024
Delay Scaling	>0.7			~1?				
Energy Scaling	~0.5			>0.5				
Transistors	Planar			3G, FinFET				
Variability	High			Extreme				
ILD	~3			towards 2				
RC Delay	1	1	1	1	1	1	1	1
Metal Layers	8-9	0.5 to 1 Layer per generation						

# How Many Transistors Can I Build?

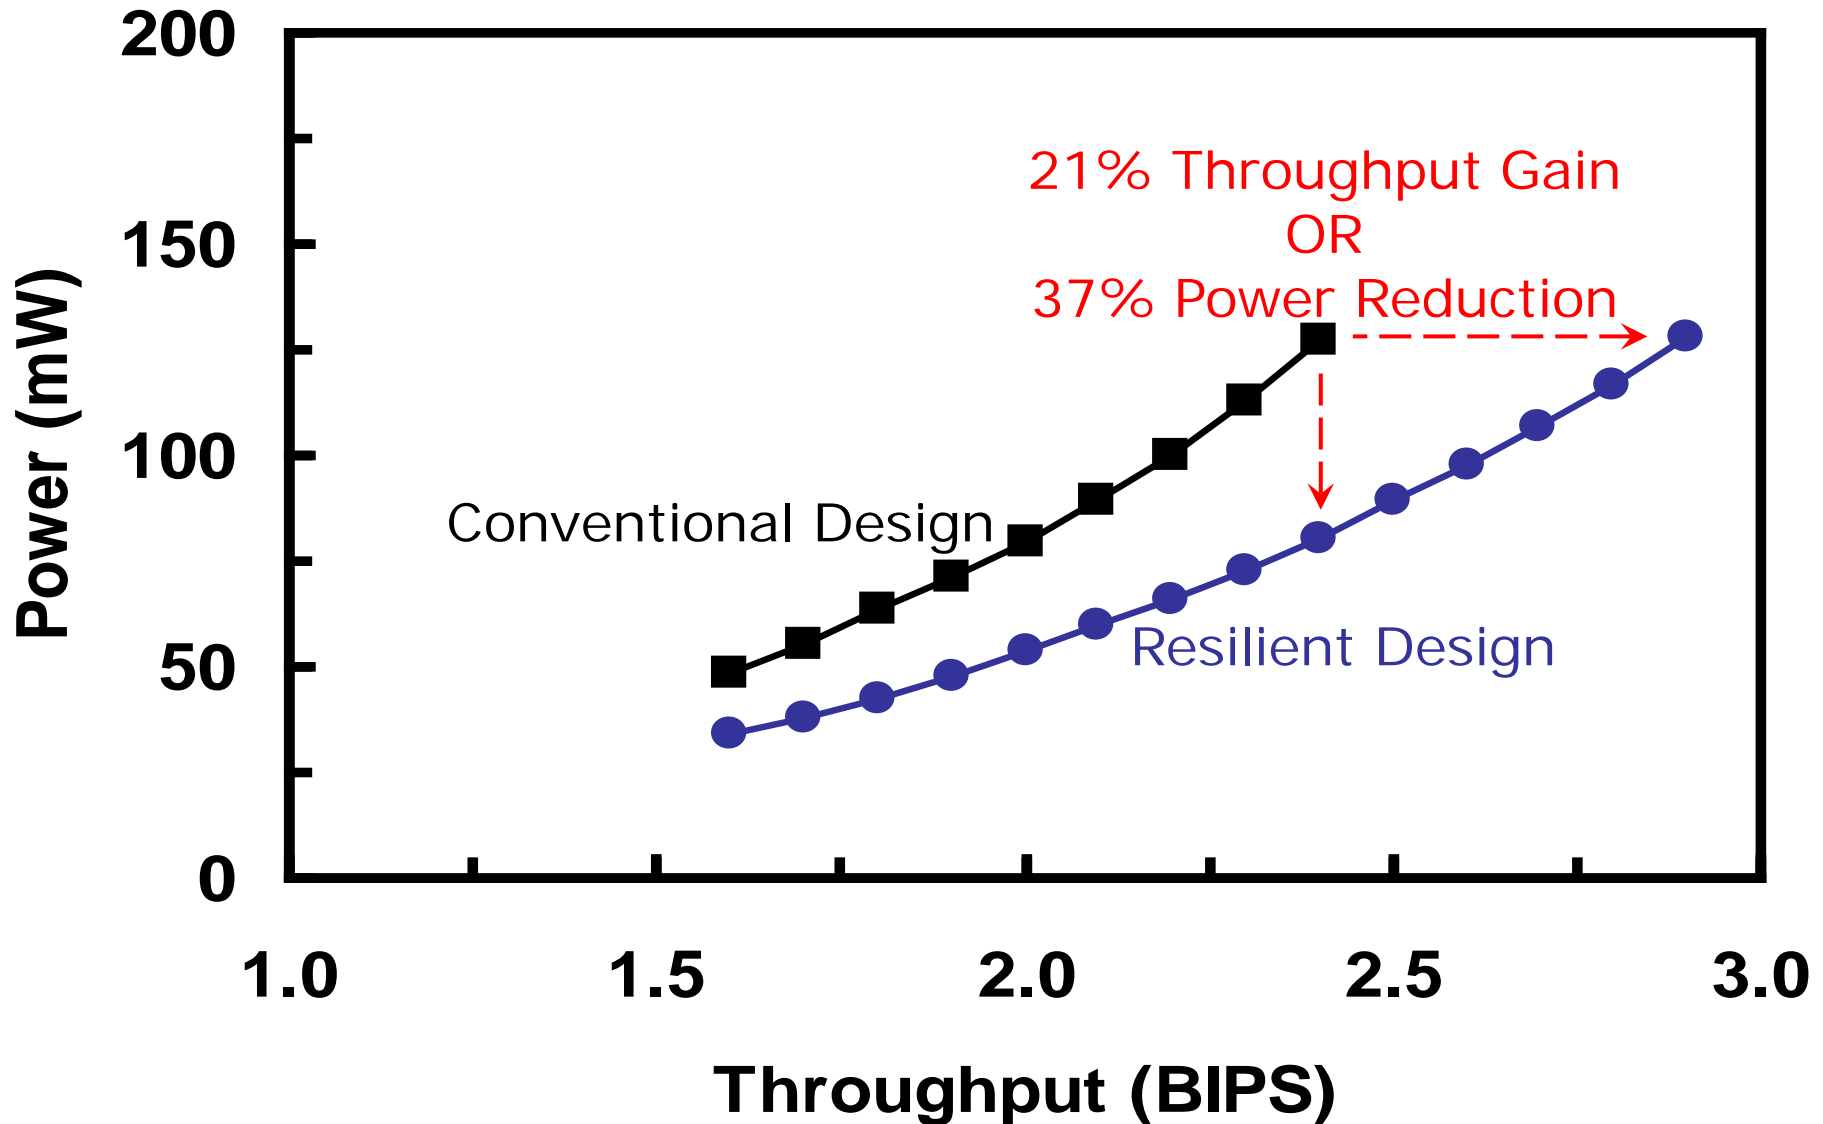


# Operate at Average-Case via Resilience (not Worst-Case)



Source: K. Bowman, et al.,  
*ISSCC, 2008.*

# Power vs. Throughput



Source: K. Bowman, et al.,  
*JSSC*, 2009.



# Implications for Resilient Systems

- Excess hardware resources
  - Can sacrifice some unit performance to reduce complexity
  - Can afford to disable some working HW to reduce overhead
- Power is everything
  - Analysis needs to consider entire cost of any HW for resilience
    - Checkpoint/restore a great example of a place where it's easy to count the wrong costs
- Range of system sizes increasing
  - Want resilience mechanisms that scale from cell phone to supercomputer



# Strawman Resilient System

- Adaptive software diagnostics determine system state and aging
- Small set of hardware mechanisms to detect errors
- Reconfigure at coarse grain (ex: core)
  - Finer-grained reconfig. only where it's very cheap (e.g., cache line disable)
- Divide error analysis/recovery/reconfiguration between HW and SW
  - Tune system for error rate that maximizes power efficiency
  - Build cheapest set of HW mechanisms that allow that error rate
  - Push everything else into SW