Future electronic system design

larssv@chalmers.se

Modern?

- Challenges:
  - Complexity
  - Performance
  - Development time
  - Power dissipation
  - Manufacturing cost
  - …

Future!

- Challenges:
  - Complexity
  - Performance
  - Development time
  - Power dissipation
  - Manufacturing cost
  - …

…only worse…

Challenge #1: Complexity

- Consumers expect more every day

Trend: ever larger projects

- Larger project teams
  - Broader sets of specialized design skills
    - How avoid Tower-of-Babel problem?
- More CAD support
  - Higher levels of abstraction
  - New tasks

Challenge #2: Performance

- Consumers will require higher bandwidths, more pixels, additional frames per second, …
- Industrial processes, truck engines, etc need tighter control
  - Economy
  - Environment

Trend: parallelism

- Microprocessors go multicore / manycore
  - Software must realize performance potential!
  - How provide predictable performance?
- Graphics processors already massively parallel
- Large-scale parallelism needs connectivity
  - Bandwidth requirements increase, at all levels

Challenge #3: Development time

- Market pressure will not abate… but
  - Product complexity increases
    - Larger projects, project teams
- More coordination

Trend: “platform-based design”

180920

- Recall: technology platform = physical tech + IP + tools
  - Extend technology platform with higher-level IP (processors, etc) to create “design platform”

Large effort

    - Choose subsets, combinations for each product

Smaller effort

- “Design platform” includes software components
  - Enables product differentiation
  - Complicates HW/SW co-verification
- Example: same SoC in iPhone, iPad

Challenge #4: Dissipation

- Higher performance, higher integration
  - Difficult to deliver high-quality supply power
  - Difficult to remove heat
- Causes equipment to be expensive, bulky, unreliable, environmentally unfriendly

Trend: specialization

- Uniform multicore full-chip processor cannot operate at full speed due to power constraints (“dark silicon”)
  - How productively use the available silicon area?
- Approach: asymmetric multicore processors w/ same ISA
  - Use the fast one or the frugal one

ARM big.LITTLE

- Next step: replace some cores with special-purpose “accelerators” (not the same ISA)
  - Even higher application performance while still inside power limits

181018

Challenge #5: Manufacturing cost

- Historically, improved through increasing levels of integration
  - Fewer parts to stock, assemble
  - Less weight to cart around
  - Better dependability at PCB level (fewer solder points that may fail)
- … but NRE costs force long manufacturing runs
  - … and spares must be manufactured for each part

Trend: Reconfigurability

- Allows hardware-design re-use with high integration
- First FPGAs simple, uniform
  - Increasing specialization
    - Adders, multipliers, memories, processors …
- Next: “platform ASICs” (cf. “design platform”)
  - Domain-specific multicore processors
  - Specialized peripherals, compute engines
  - Reconfigurable interconnects, networks-on-chip
  - Allows software-defined behavior

Cf. SoCs, 181018

The driving trend

Dr. Moore, again.

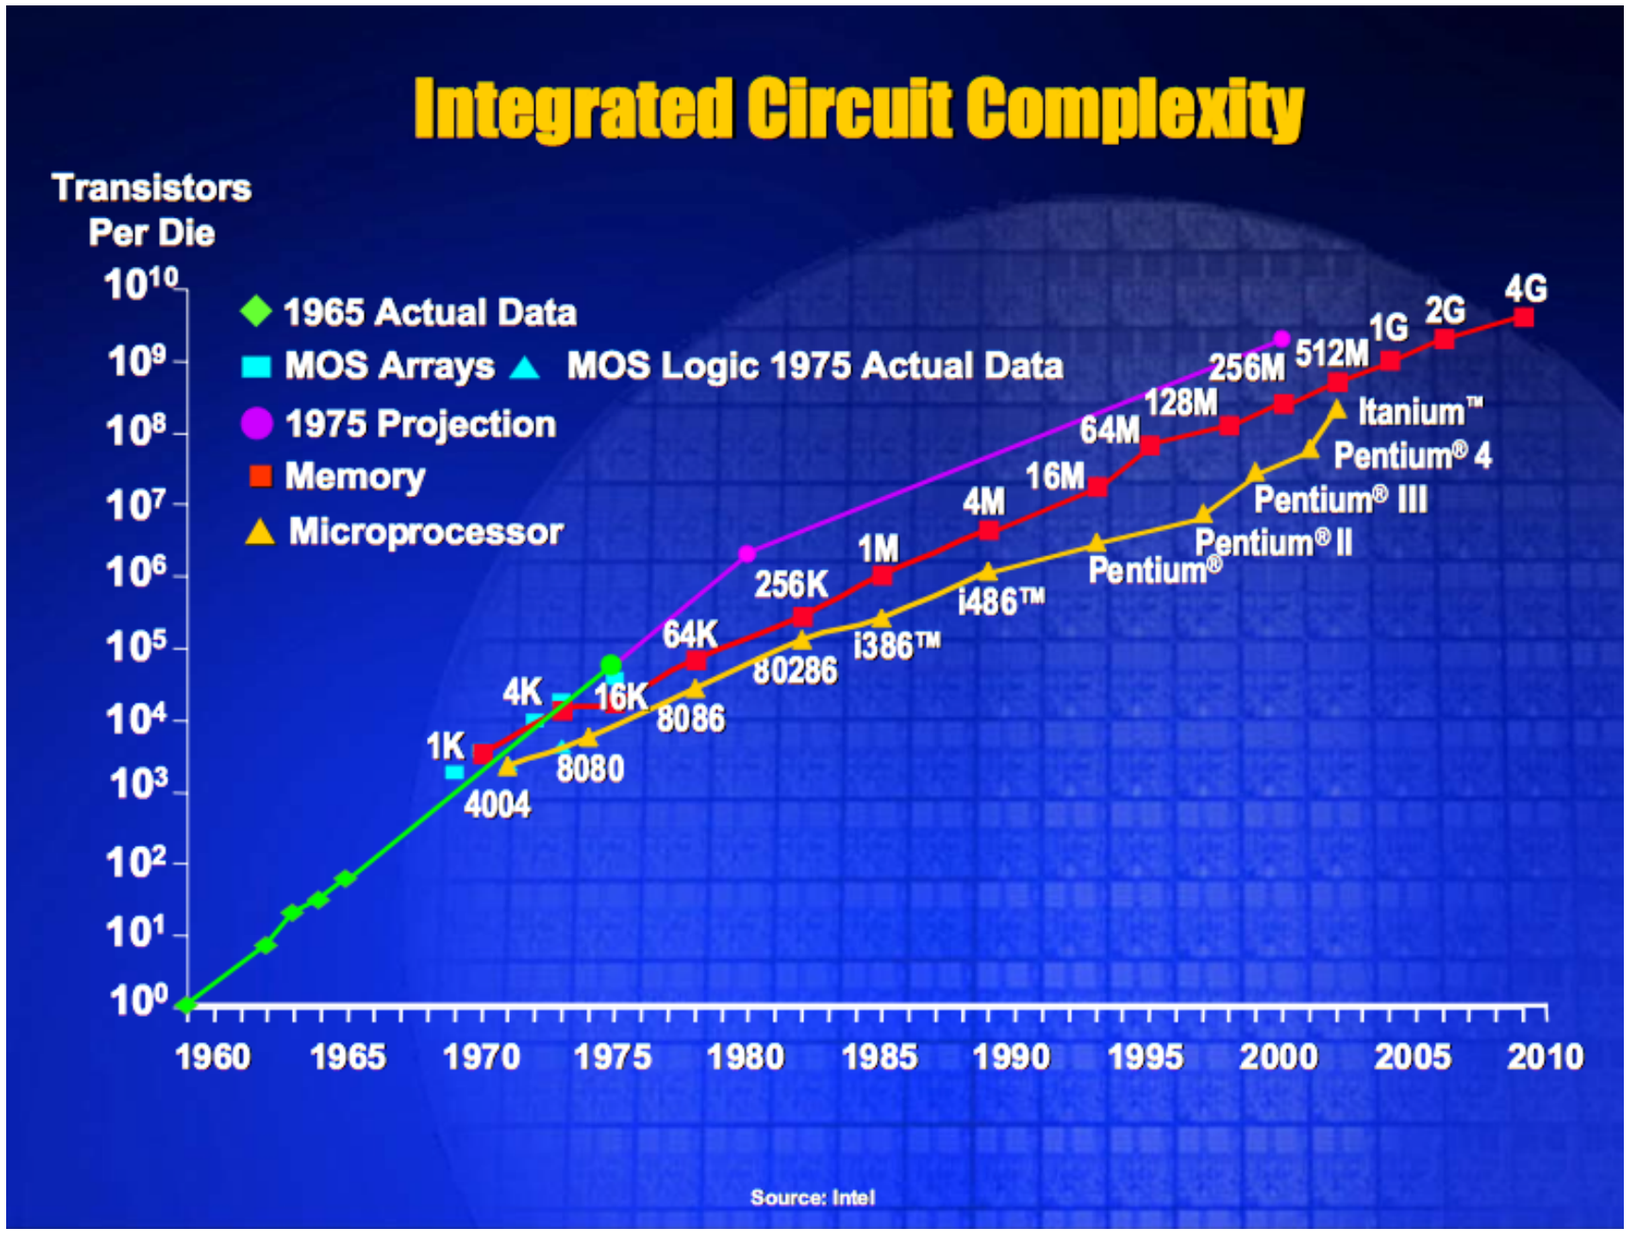

NO EXPONENTIAL IS FOREVER . . .

Future of Moore?

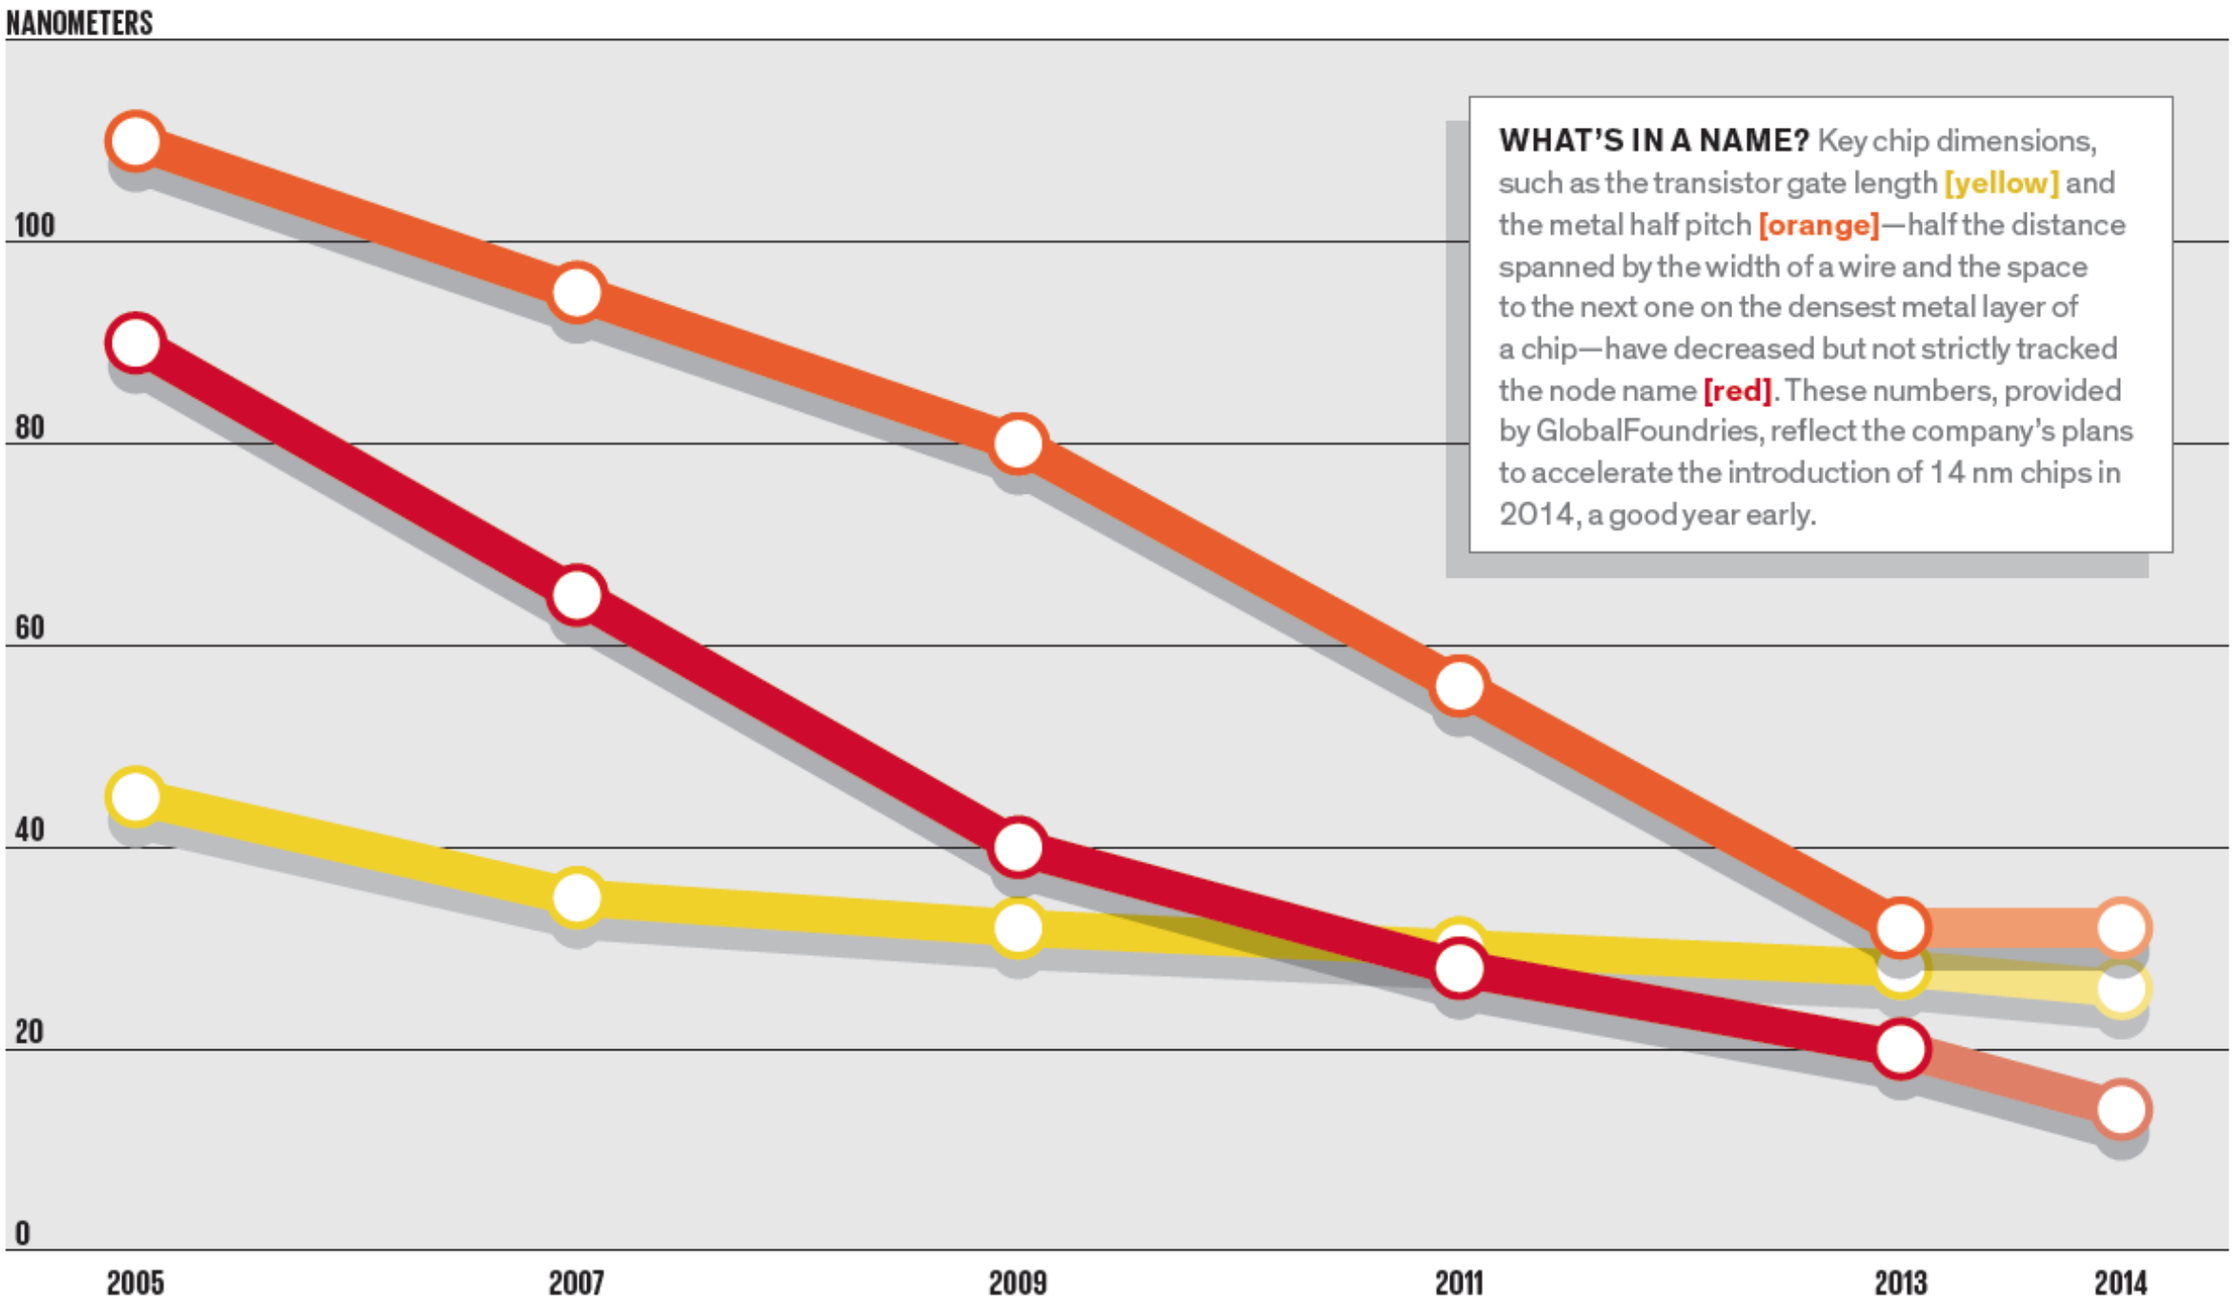

[Courtland: "The end of the shrink", IEEE Spectrum, Nov 2013]

International Technology Roadmap for Semiconductors

- Detailed elucidation of Moore
- Regular releases
  - Odd years
- Updates
  - Even years

ITRS example: lithography

	First Year of IC Production	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026
	DRAM ½ pitch (nm) (contacted)	36	32	28	25	23	20.0	17.9	15.9	14.2	12.6	11.3	10.0	8.9	8.0	7.1	6.3
	MPU/ASIC Metal 1 1/2 pitch (nm)	38	32	27	24	21	18.9	16.9	15.0	13.4	11.9	10.6	9.5	8.4	7.5	7.5	7.5
45	193nm Imm																
32	193 nm DP																
22	EUV																
	193nm MP	Narrow Options															
	ML2 (MPU)																
16	EUV																
	193nm MP																
	ML2						Narrow Options										
	DSA + Litho																
	Imprint																
11	EUV higher NA / EUV + DP																
	ML2																
	DSA + Litho									Narrow Options							
	EUV (new wavelength)																
	Imprint																
	Innovation																

MPU / DRAM time line

- Start with many alternatives
- Invest in a few, drop the rest
- Select one and then go for production in another 2? years

[ITRS 2012 update]

Variability + Yield

What is the problem?

- How predict performance of manufactured device?
- Will it work at all?

- Complex processing – difficult to model
- Intrinsic statistic variations

Yield

- Important issue in all manufacturing
- Ratio of successful items to total items manufactured
  - The rest is the yield loss
    - High yield is what makes your profit!

- Advanced tech is always immature
  - Expect yield to improve with time as manufacturing technology matures

Yield loss types

- Defect (“process”) yield loss
  - Item is unusable / unsellable
  - No resale value!
- Performance (“parametric”) yield loss
  - Item performance does not reach spec
  - Some resale value?
  - Example: microprocessor/FPGA speed grades

Example yield loss mechanisms

- Metal line width and spacing variations
  - Defects: opens or shorts
  - Parameters: R/C variations
- Device gate uniformity
  - Parameters, defects: W, L, V_T

Design For Manufacturing

- Need to understand variability
  - Causes
  - Effects
- Then, possible to <u>design</u> so as to <u>manage</u> variability
  - Minimize defect yield loss
  - Control (and accept some) parametric yield loss
- DFM contributes to design rule-set explosion
  - 100–200 in 130nm, ~1000 in 65nm, …
  - … and the 22nm rule book is > 800 pages …

Statistical exercise!

Variability categories

1. Static variability
   - Essentially manufacturing spread
2. Dynamic variability
   - Influence by environment, context

- Not perfect separation…

1. Static variability

- Theoretical categories:
  - Systematic variations (predictable, but not necessarily predicted)
    - Truly random fluctuations
  - In practice, unpredicted variations resemble random fluctuations
  - Continuing model refinements (improved predictions) reduce “random” spread

V_T variability (from device simulations)

- 25 nm bulk MOSFET
- Averages
  - ± 2 std devs
    - 95% coverage
- Expect some very leaky devices!

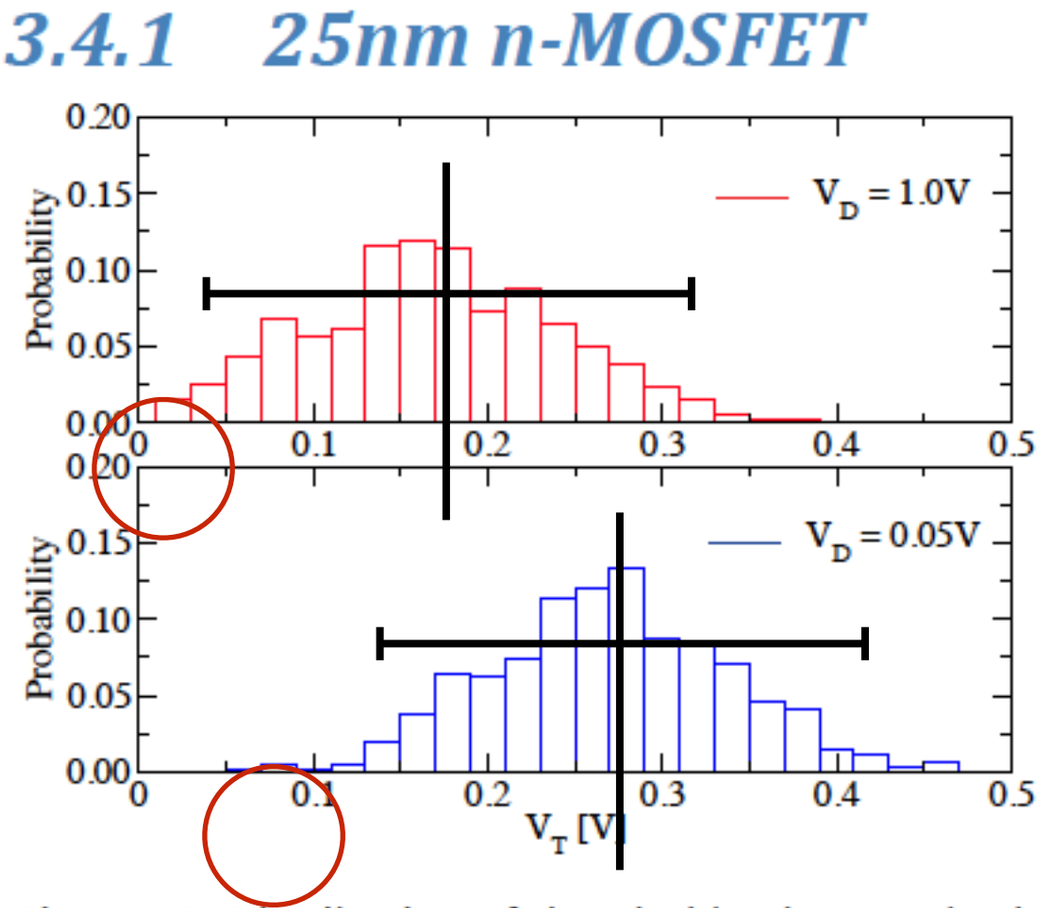

Figure 42: Distribution of threshold voltage at both low drain and high drain biases in the 25nm n-MOSFET.

2. Dynamic variability

- Supply voltage variations
  - Switching logic, clock gating, V_{dd} gating
- Temperature variations
  - Hot spots increase device leakage
- Aging
  - V_T drift, metal electromigration, …

V_{dd} design margins

- Supply voltage must support performance requirements, always
- V_{dd} includes PVT margins (process, voltage, temperature)
- Larger (%) margins needed in more advanced processes
- Performance / power cost

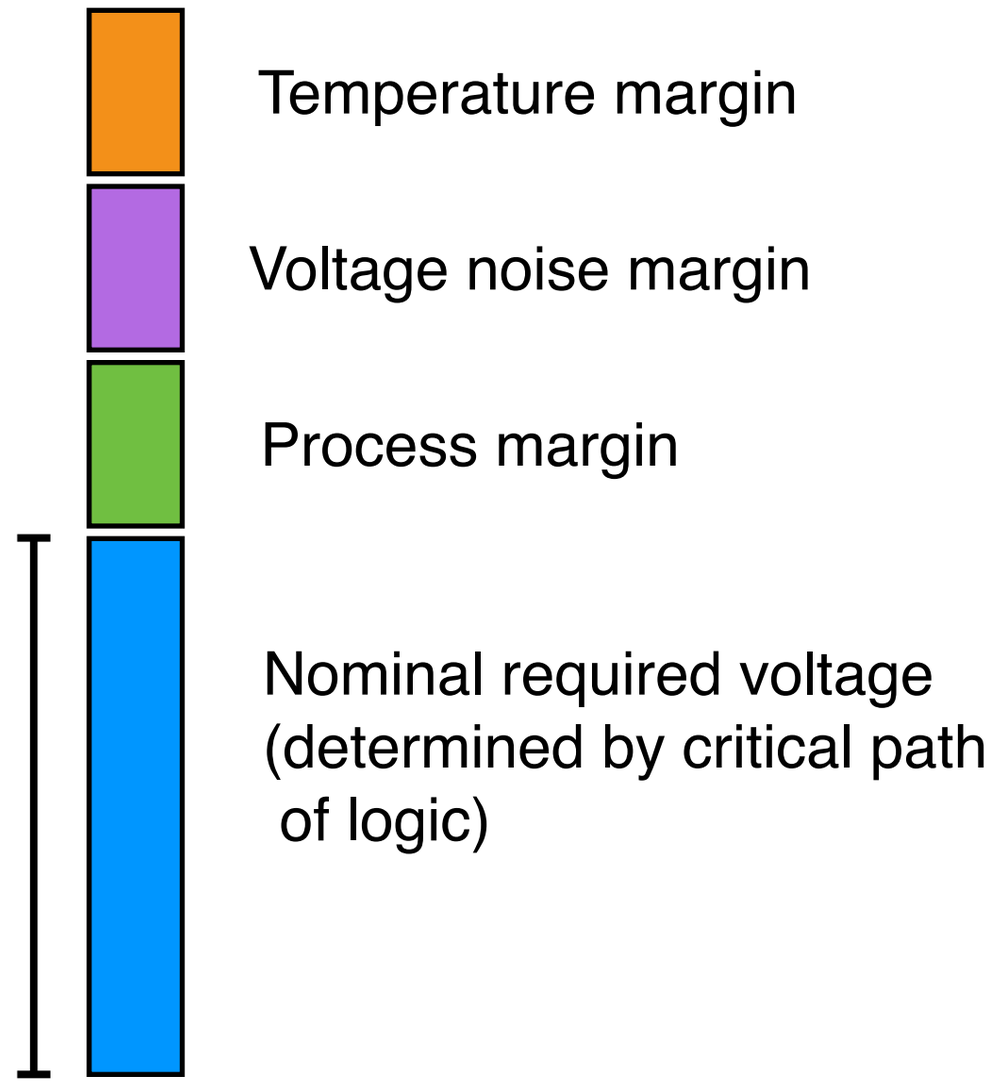

Present yield loss remedies

- Defect yield loss:
  - Redundancy
  - Error-Control Codes (ECC)
  - Built-In Self Test (BIST)
- Performance yield loss:
  - Adaptivity (e.g. detect setup violations, increase voltage, try again)
- Verifiability issues!

Future yield loss remedies

- More regularity (layout, logic, etc)
  - Proven structures, predictable performance
    - Area overhead
- More redundancy
  - Turn off inactive circuits (leakage)
- Reconfiguration, SW/HW adaptability

Cost of regularity

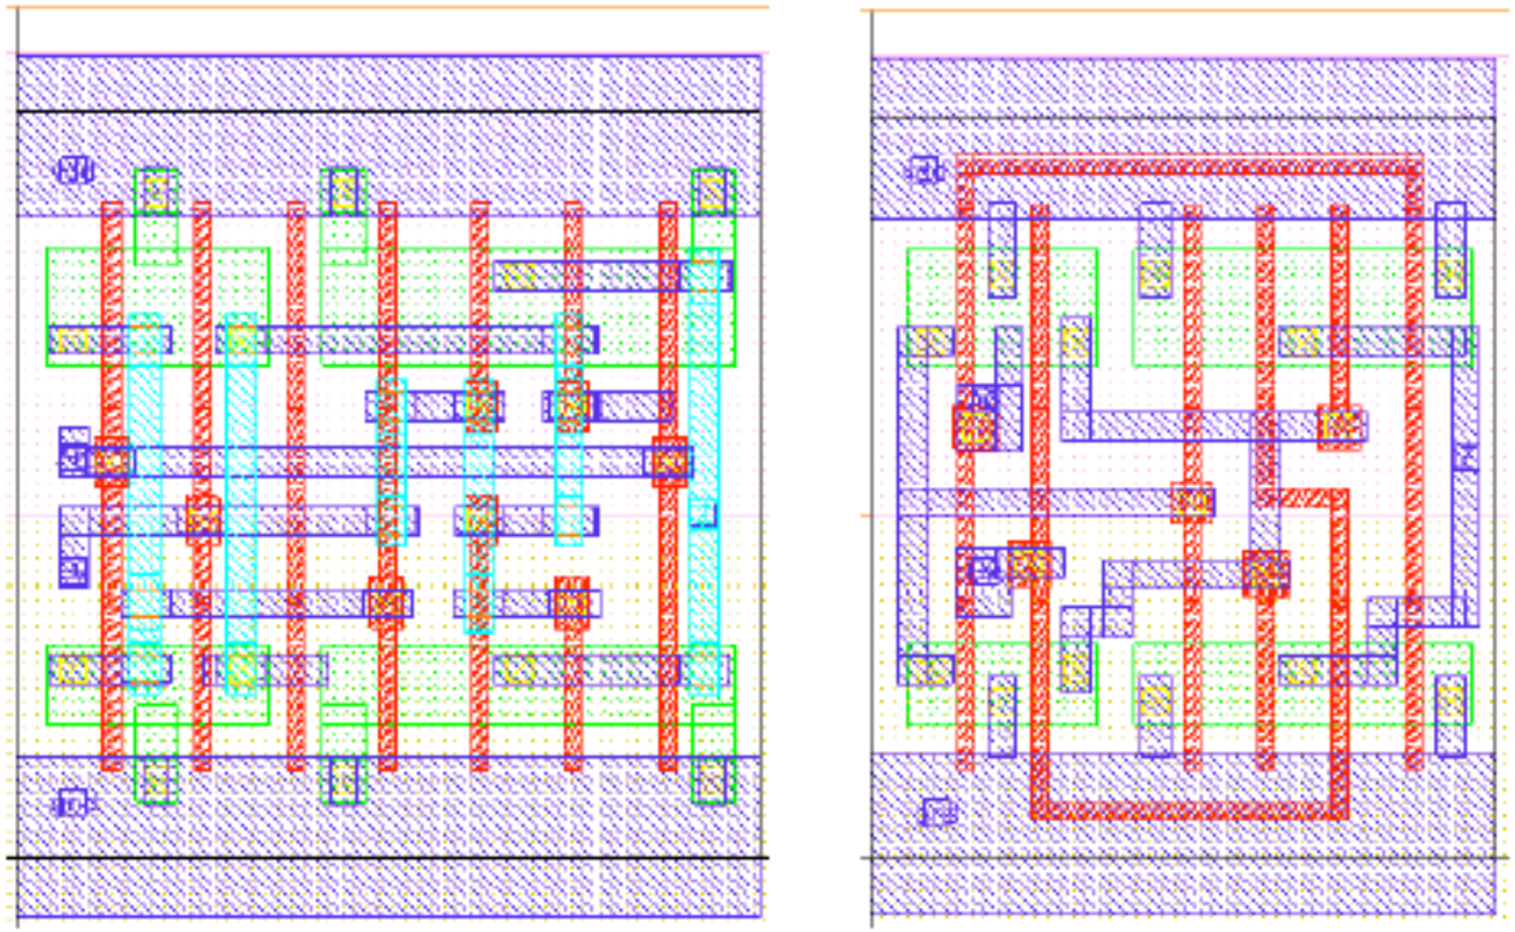

(a) Ultra-regular (b) Semi-regular

- Highly regular cells are larger
- Also cost in routing, placement…

[K Subramaniyan, ASQED 2012 *(best paper award)*]

Statistical design methods

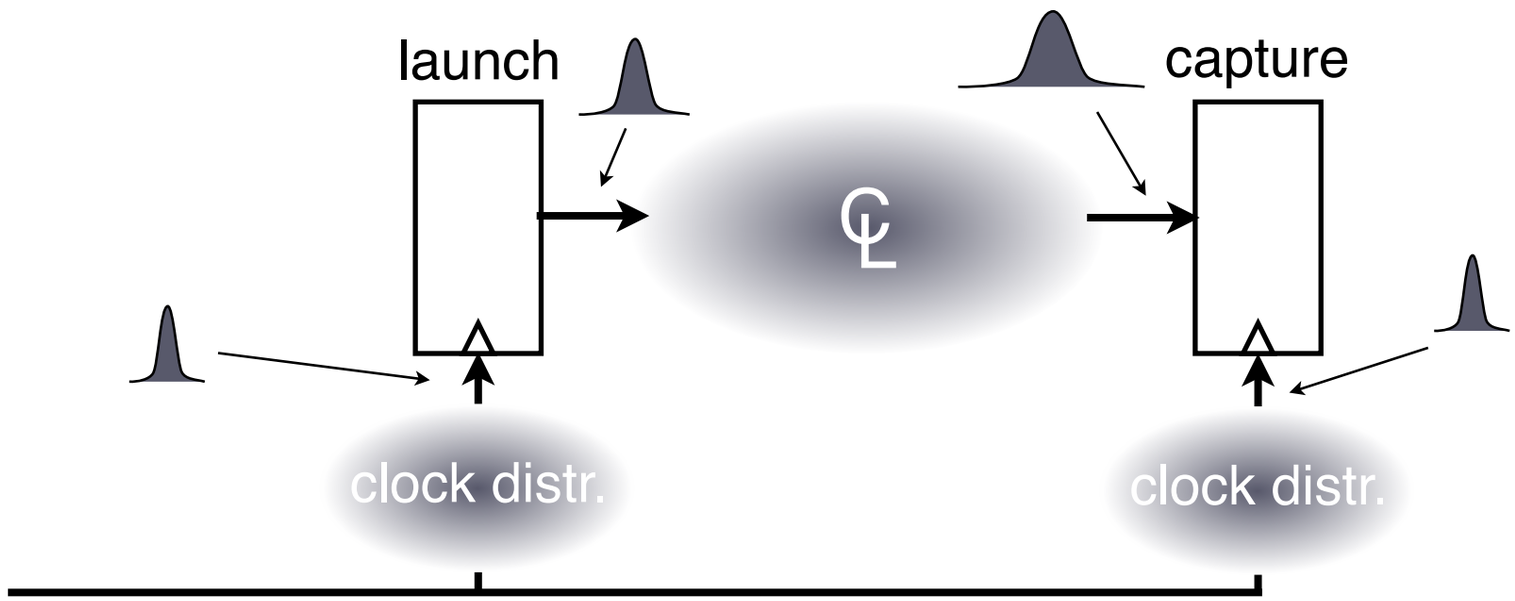

- Example: static timing analysis (STA)
  - Max and min delay through CL must conform to setup, hold reqs at capture FF
  - Min and max clock delays constrain problem
- Statistical STA: With normal distributions, not go/no-go!
- Expect similar reasoning at higher abstraction levels

Example: Redundancy in multicore processors

- Assume at least *M* cores needed for performance

- Design with *N* > *M* cores to have some spares

- Cost grows with *N*. How choose redundancy, that is, *N* – *M* ?

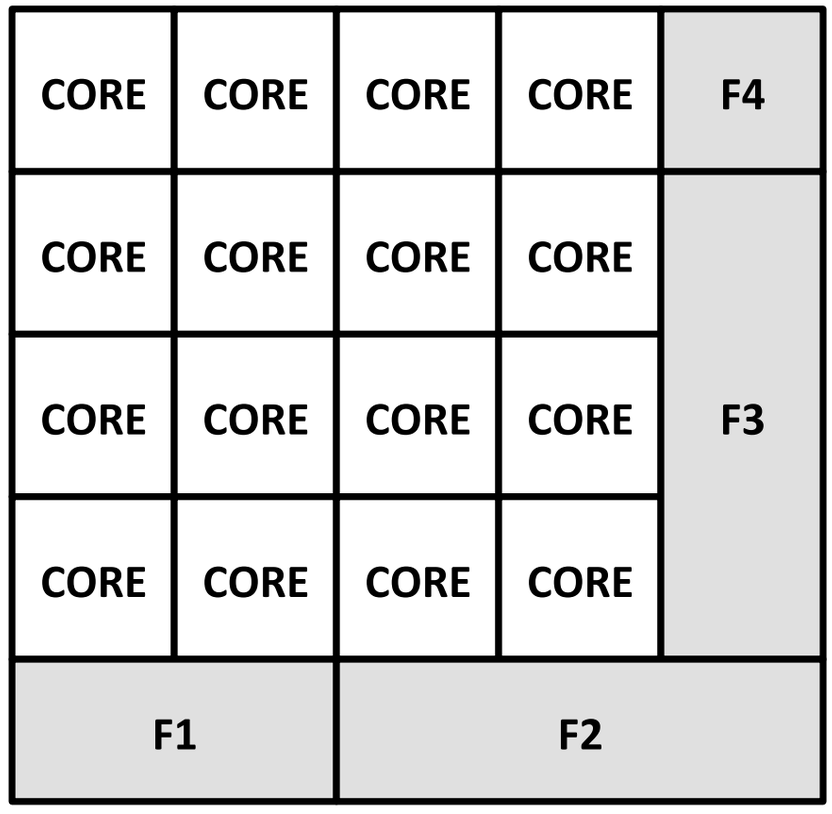

Figure 1: Chip with 16 cores ($N = 16$).

[Mikael Andersson: Integrated Circuit Yield Enhancement – Redundant Multi Core DSP Cluster. MSc thesis, 2010]

Chips per wafer

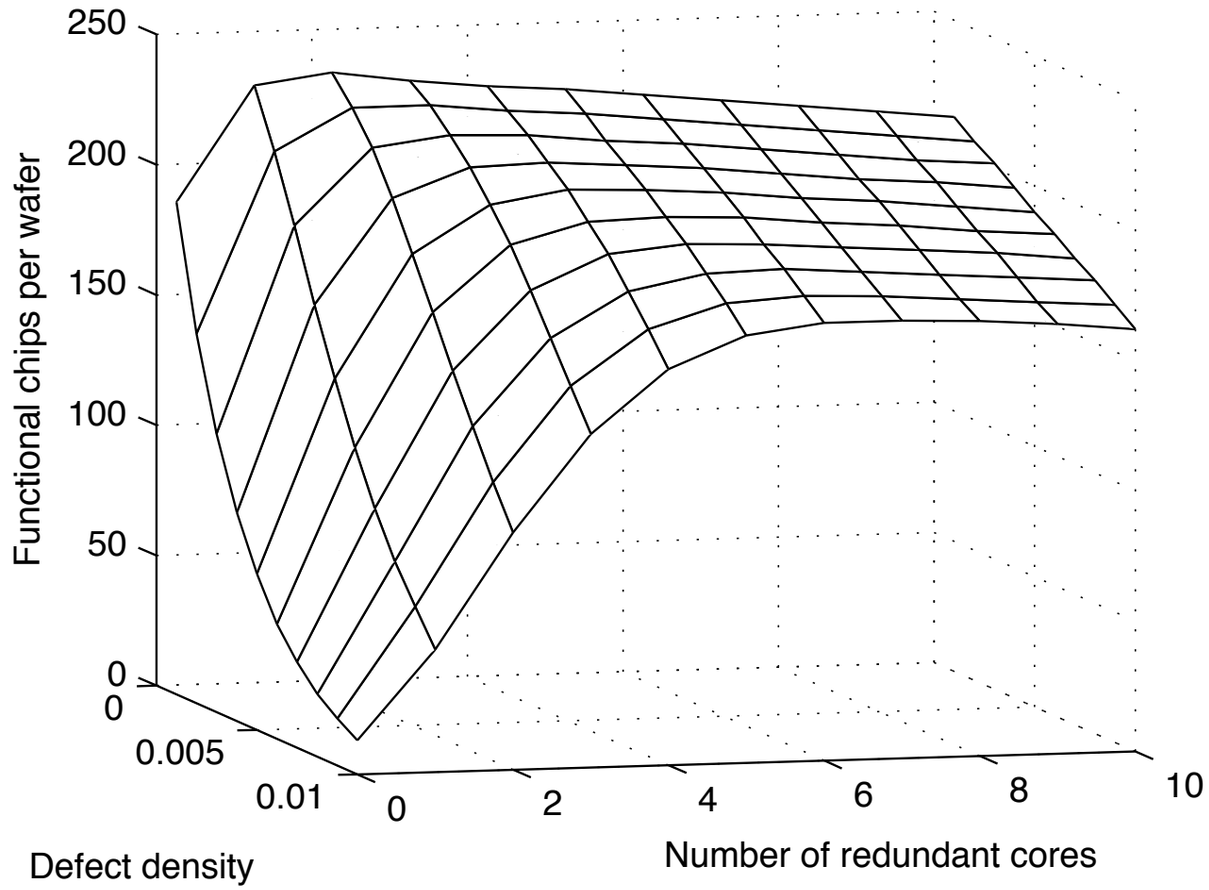

Figure 2: Functioning chips per wafer for different defect densities and numbers of redundant cores.

- "Best" value of *N* – *M* depends on defect densities
- If in doubt, use more redundant cores…

Yield over time

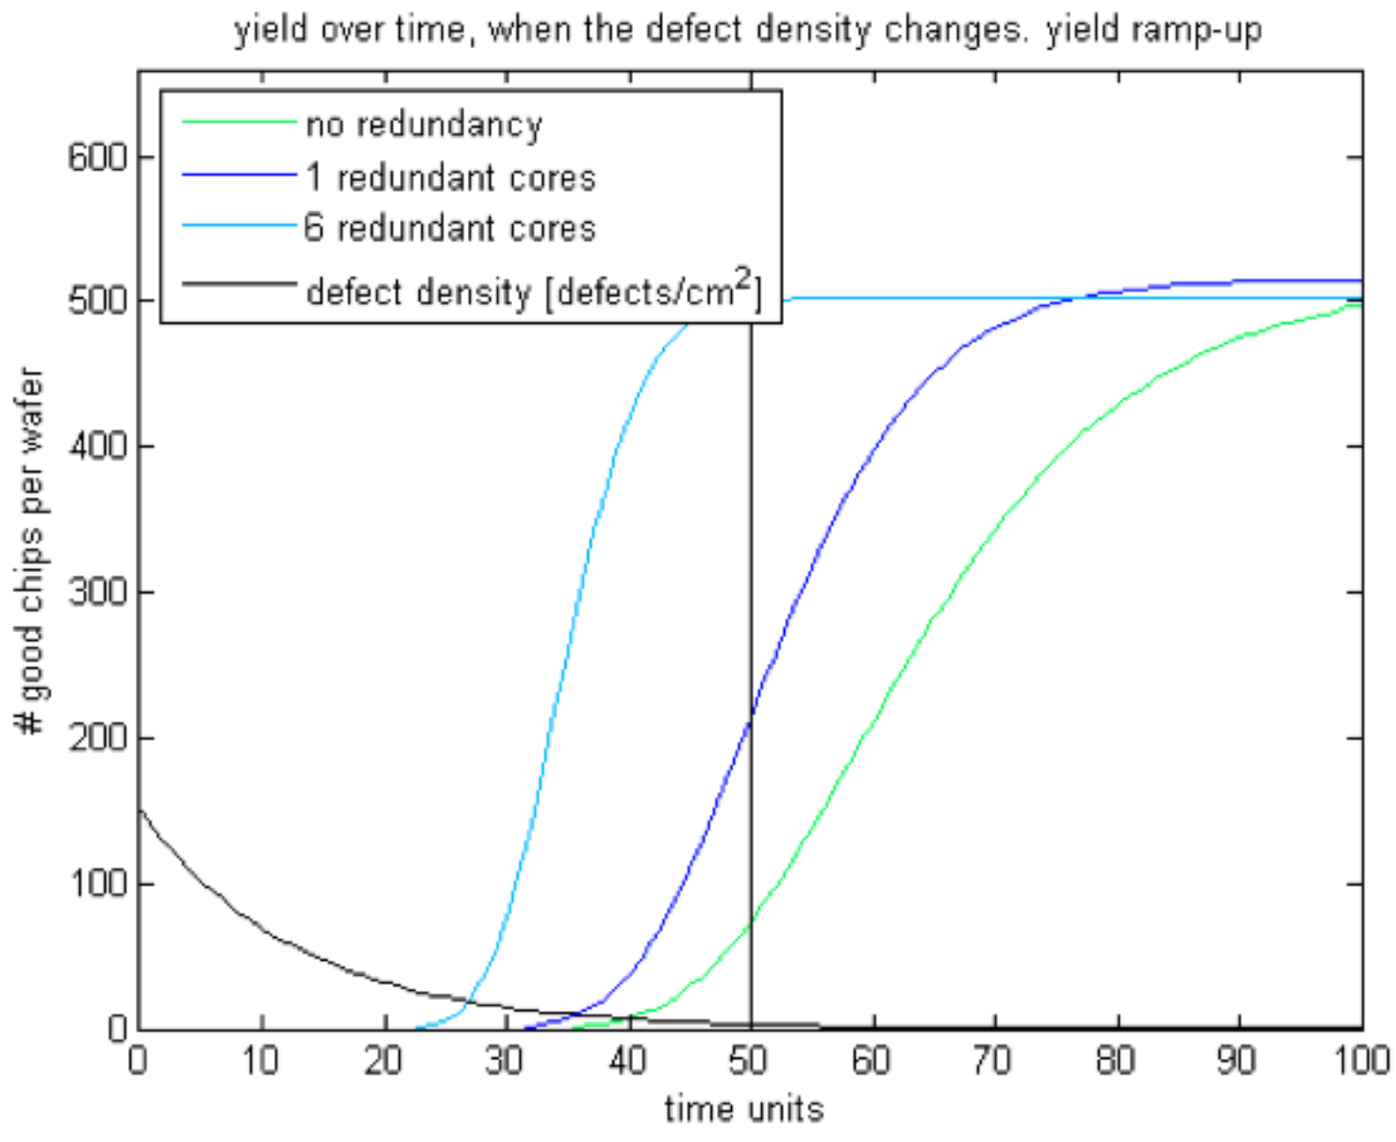

- New fab: many defects -> low yield
- Higher redundancy allows earlier production

The End is Nigh! (?)

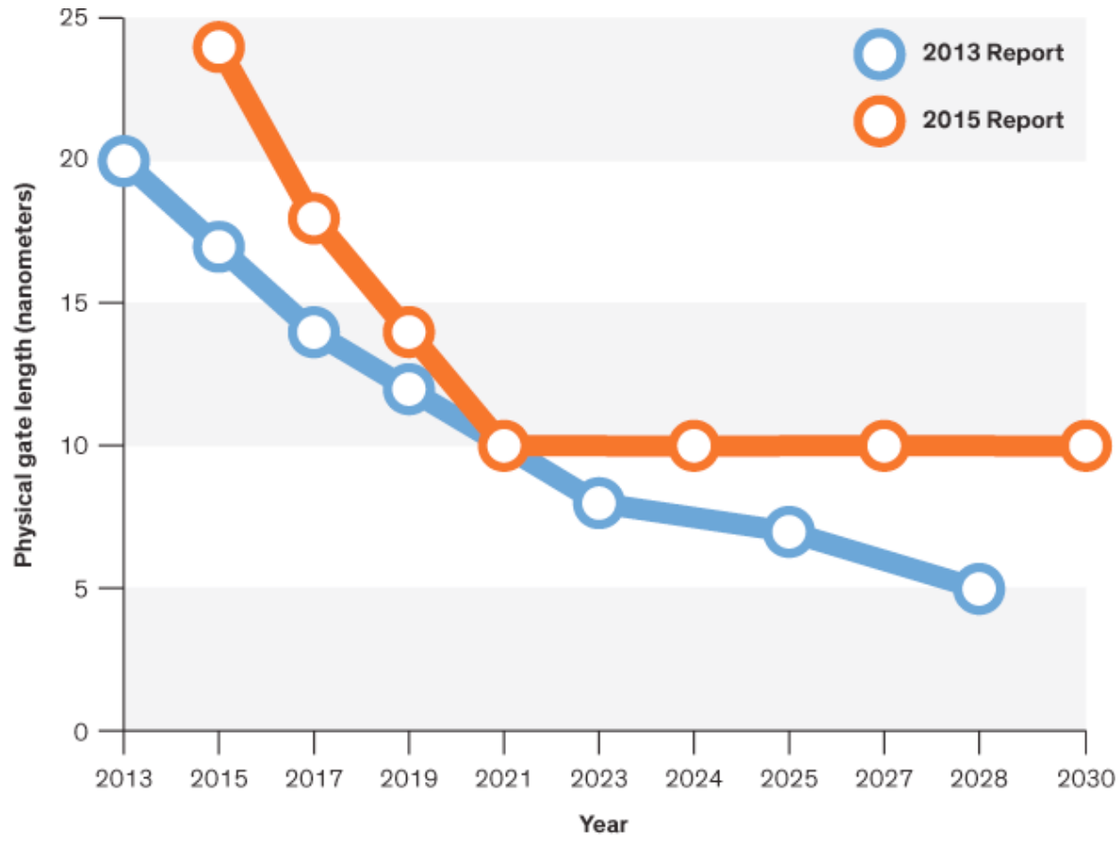

- ITRS 2015 predicts scaling will bottom out within 5 years (“the end of the Moore as we know it” 😨)
- ITRS 2016 report will be the last one
  - “ITRS 2.0”
- Future density increases must go vertical
  - How remove heat from inside of 3D chip stack?

[Courtland 2016 (uploaded)]

Summary

- Challenges remain (and get worse)
  - Other problems come to the fore
- Variability important for near- to mid-term future
  - Statistical methods will emerge also at higher levels
- No more Moore requires 3D construction
- Exciting times ahead!