

Moore's Law at 50: Are we planning for retirement?

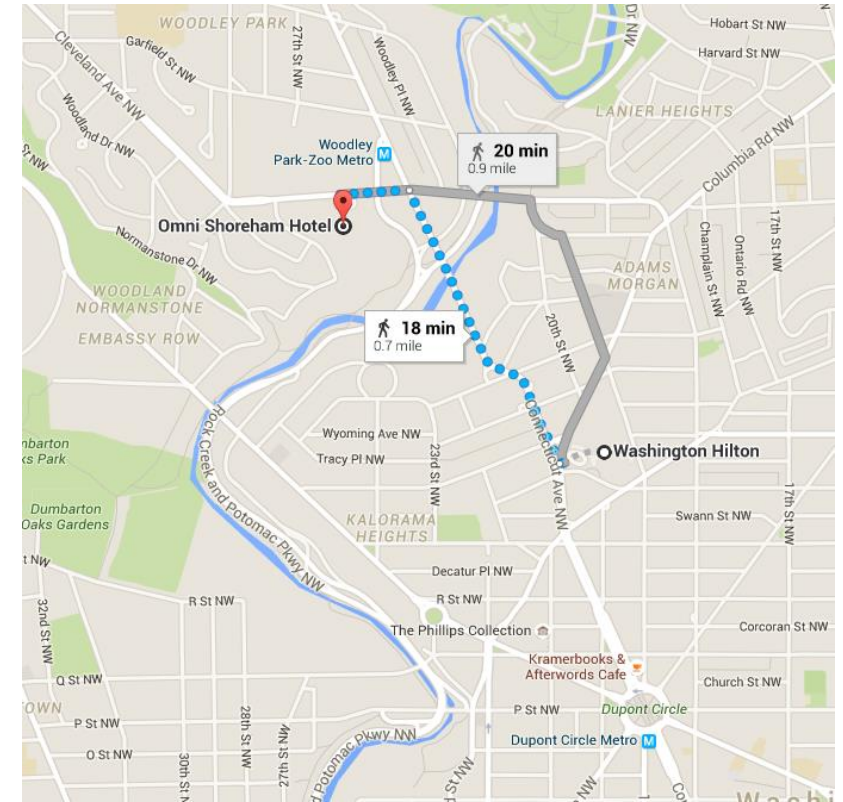
Greg Yeric

ARM

60th Anniversary of IEDM...



1st Annual Technical Meeting on Electron Devices: 1955




...and last in D.C.?

“58 years on” *



Single-word nickel delay line (Science Museum, London)



Punched cards
3.3kB memory
3mS multiply
> 15 kW
~ £ 20,000

1080p video
512 MB memory
1 GHz CPU
0.6W
£ 4

*<http://imgur.com/a/cafly>


2015: Buy a magazine, get a free computer



raspberrypi.org




or, just be a 7th grader
In the UK




1080p video
512 MB memory
1 GHz CPU
0.6W
£ 4

2015: 50 years of an exponential



The Question



Moore's Law: 1965

The experts look ahead

Cramming more components onto integrated circuits

With unit cost falling as the number of components per circuit rises, by 1975 economics may dictate squeezing as many as 65,000 components on a single silicon chip

By Gordon E. Moore

Director, Research and Development Laboratories, Fairchild Semiconductor division of Fairchild Camera and Instrument Corp.

The future of integrated electronics is the future of electronics itself. The advantages of integration will bring about a proliferation of electronics, pushing this science into many new areas.

Integrated circuits will lead to such wonders as home computers—or at least terminals connected to a central computer—automatic controls for automobiles, and personal portable communications equipment. The electronic wrist-watch needs only a display to be feasible today.

But the biggest potential lies in the production of large systems. In telephone communications, integrated circuits in digital filters will separate channels on multiplex equipment. Integrated circuits will also switch telephone circuits and perform data processing.

Computers will be more powerful, and will be organized in completely different ways. For example, memories built of integrated electronics may be distributed throughout the

The author
Dr. Gordon E. Moore is one of the new breed of electronic engineers, schooled in the physical sciences rather than in electronics. He earned a B.S. degree in chemistry from the University of California and a Ph.D. degree in physical chemistry from the California Institute of Technology. He was one of the founders of Fairchild Semiconductor and has been director of the research and development laboratories since 1959.

machine instead of being concentrated in a central unit. In addition, the improved reliability made possible by integrated circuits will allow the construction of larger products. Machines similar to those in existence today will have lower costs and with faster turn-around.

Present and future
By integrated electronics, I mean all the technologies which are referred to as microelectronics as well as any additional ones that result in electronics supplied to the user as irreducible units. Technologies were first investigated in the late 1950s which were to miniaturize electronics equipment into increasingly complex electronic functions in limited minimum weight. Several approaches evolve microassembly techniques for individual component film structures and semiconductor integrated circuits. Each approach evolved rapidly and converges each borrowed techniques from another. Many believe the way of the future to be a combination of these approaches.

The advocates of semiconductor integrated electronics already using the improved characteristics of thin films by applying such films directly to an active or substrate. Those advocating a technology films are developing sophisticated techniques for the present of active semiconductor devices to the past.

Both approaches have worked well and are in equipment today.

Electronics, Volume 38, Number 8, April 19, 1965

The establishment

Integrated electronics is established today. Its techniques are almost mandatory for new military systems, since the reliability, size and weight required by some of them is achievable only with integration. Such programs as Apollo, for manned moon flight, have demonstrated the reliability of integrated electronics by showing that complete circuit functions are as free from failure as the best individual transistors.

Most companies in the commercial computer field have machines in design or in early production employing integrated electronics. These machines cost less and perform better than those which use "conventional" electronics.

Improvements of various sorts, especially the rapidly increasing numbers employing digital techniques, are starting to use integration because it cuts costs of both manufacture and design.

The use of linear integrated circuitry is still restricted primarily to the military. Such integrated functions are expensive and not available in the variety required to satisfy a major fraction of linear electronics. But the first applications are beginning to appear in commercial electronics, particularly in equipment which needs low-frequency amplifiers of small size.

Reliability counts

In almost every case, integrated electronics has demonstrated high reliability. Even at the present level of production—low compared to that of discrete components—it offers reduced systems cost, and in many systems improved performance has been realized.

Integrated electronics will make electronic techniques more generally available throughout all of society, performing many functions that presently are done inadequately by other techniques or not done at all. The principal advantages will be lower costs and greatly simplified design—payoffs from a ready supply of low-cost functional packages.

For most applications, semiconductor integrated circuits will predominate. Semiconductor devices are the only reasonable candidates presently in existence for the active elements of integrated circuits. Passive semiconductor elements look attractive too, because of their potential for low cost and high reliability, but they can be used only if precision is not a prime requisite.

Silicon is likely to remain the basic material, although others will be of use in specific applications. For example, gallium arsenide will be important in integrated microwave functions. But silicon will predominate at lower frequencies because of the technology which has already evolved around it and its oxide, and because it is an abundant and relatively inexpensive starting material.

Costs and curves
Reduced cost is one of the big attractions of integrated electronics, and the cost advantage continues to increase as the technology evolves toward the production of larger and larger circuit functions on a single semiconductor substrate. For simple circuits, the cost per component is nearly inversely proportional to the number of components, the result of the

equivalent piece of semiconductor in the equivalent package containing more components. But as components are added, decreased yields more than compensate for the improvement in complexity, tending to raise the cost per component.


At present, it is reached when 50 units are used per circuit. But the minimum is rising while the entire cost curve is falling (see graph below). A plot of costs suggests that minimum cost per component might be expected in circuits about 1,000 components per circuit (providing such functions can be produced in moderate quantities.) The manufacturing cost per component can be expected only a tenth of the present cost.

The complexity for minimum component cost increased at a rate of roughly a factor of two per year graph on next page). Certainly over the short term can be expected to continue, if not to increase. Over longer term, the rate of increase is a bit more uncertain though there is no reason to believe it will not remain constant for at least 10 years. That means by 1975, the number of components per integrated circuit for minimum will be 65,000.

I believe that such a large circuit can be built on a wafer.

Two-mil squares

With the dimensional tolerances already being achieved in integrated circuits, isolated high-performance units can be built on centers two thousandths of an inch apart.



Electronics, Volume 38, Number 8, April 19, 1965




a two-mil square can also contain several kilolms of resistance or a few diodes. This allows at least 500 components per linear inch or a quarter million per square inch. Thus, 65,000 components need occupy only about one-fourth a square inch.

On the silicon wafer currently used, usually an inch or more in diameter, there is ample room for such a structure if the components can be closely packed with no space wasted for interconnection patterns. This is realistic, since efforts to achieve a level of complexity above the presently available integrated circuits are already underway using multilayer metallization patterns separated by dielectric films. Such a density of components can be achieved by present optical techniques and does not require the more exotic techniques, such as electron beam operations, which are being studied to make even smaller structures.

Increasing the yield

There is no fundamental obstacle to achieving device yields of 100%. At present, packaging costs so far exceed the cost of the semiconductor structure itself that there is no incentive to improve yields, but they can be raised as high as



Electronics, Volume 38, Number 8, April 19, 1965

is economically justified. No barrier exists comparable to the thermodynamic equilibrium considerations that often limit yields in chemical reactions; it is not even necessary to do any fundamental research or to replace present processes. Only the engineering effort is needed.

In the early days of integrated circuitry, when yields were extremely low, there was such incentive. Today ordinary integrated circuits are made with yields comparable to those obtained for individual semiconductor devices. The same pattern will make larger arrays economical, if other considerations make such arrays desirable.

Heat problem

Will it be possible to remove the heat generated by tens of thousands of components in a single silicon chip?


If we could shrink the volume of a standard high-speed digital computer to that required for the components themselves, we would expect it to glow brightly with present power dissipation. But it won't happen with integrated circuits. Since integrated electronic structures are two-dimensional, they have a surface available for cooling close to each center of heat generation. In addition, power is needed primarily to drive the various lines and capacitances associated with the system. As long as a function is confined to a small area on a wafer, the amount of capacitance which must be driven is distinctly limited. In fact, shrinking dimensions on an integrated structure makes it possible to operate the structure at higher speed for the same power per unit area.

Day of reckoning

Clearly, we will be able to build such component-crammed equipment. Next, we ask under what circumstances we should do it. The total cost of making a particular system function must be minimized. To do so, we could amortize the engineering over several identical items, or evolve flexible techniques for the engineering of large functions so that no disproportionate expense need be borne by a particular array. Perhaps newly devised design automation procedures could translate from logic diagram to technological realization without any special engineering.

It may prove to be more economical to build large

Greg: 1965



Age 35

MIT Technology Review

NEWS & ANALYSIS · FEATURES · VIEWS · MULTIMEDIA · DISCUSSIONS · TOPICS · POPULAR: 50 SMARTEST COMPANIES 2015

TWELVE TOMORROWS Science Fiction Inspired by Today's New Technologies [Pre-Order](#)

Want to go ad free?


FEATURED STORY 1 COMMENT

The End of Moore's Law?

The current economic boom is likely due to increases in computing speed and decreases in price. Now there are some good reasons to think that the party may be ending.

By Charles C. Mann on May 1, 2000

Age 40




CLOUD CXO HARDWARE MICROSOFT STORAGE INNOVATION MOBILITY MORE NEWSLETTERS

JUST IN **MICROSOFT COMMITS TO 10-YEAR SUPPORT LIFECYCLE FOR WINDOWS 10**

Is Moore's Law dead at 40 or is this just a mid-life crisis?

Last week, Michael Kanellos published this FAQ on the 40th anniversary of Moore's law, which is famously known as the phenomenon that computer processing power will double every 18 months.? Actually, Gordon Moore only said that transistor count would double every 24 months and it was David House (a former executive of Intel) who extrapolated that performance would double every 18 months as a result of the increase in transistors.?

 By [George Ou](#) for [Real World IT](#) | April 5, 2005 -- 01:48 GMT (18:48 PDT) | Topic: [Processors](#)

The End of Moore's Law

Microchips are getting smaller—and that's the problem.



By Adam L. Penenberg

Until recently, **Moore's Law**, the observation that the number of transistors on a microchip doubles every 18 months to two years, seemed a self-fulfilling prophecy. When Intel co-founder Gordon Moore issued his famous prediction 40 years ago, a chip could hold a few dozen transistors. Today, Intel can cram almost **1 billion transistors**, each of which is less than 100 nanometers in size, on a single microchip. (One nanometer is 1 millionth of a millimeter—the equivalent of about 10 hydrogen atoms.) The transistors on Intel's chips are so tiny that they're not visible to the naked eye. *

2015: Age 50


<re/code> NEWS ▾ | REVIEWS | VIDEO | PODCASTS | VOICES | WRITERS

Moore's Law Hits 50, but It May Not See 60



Dr. Gordon Moore

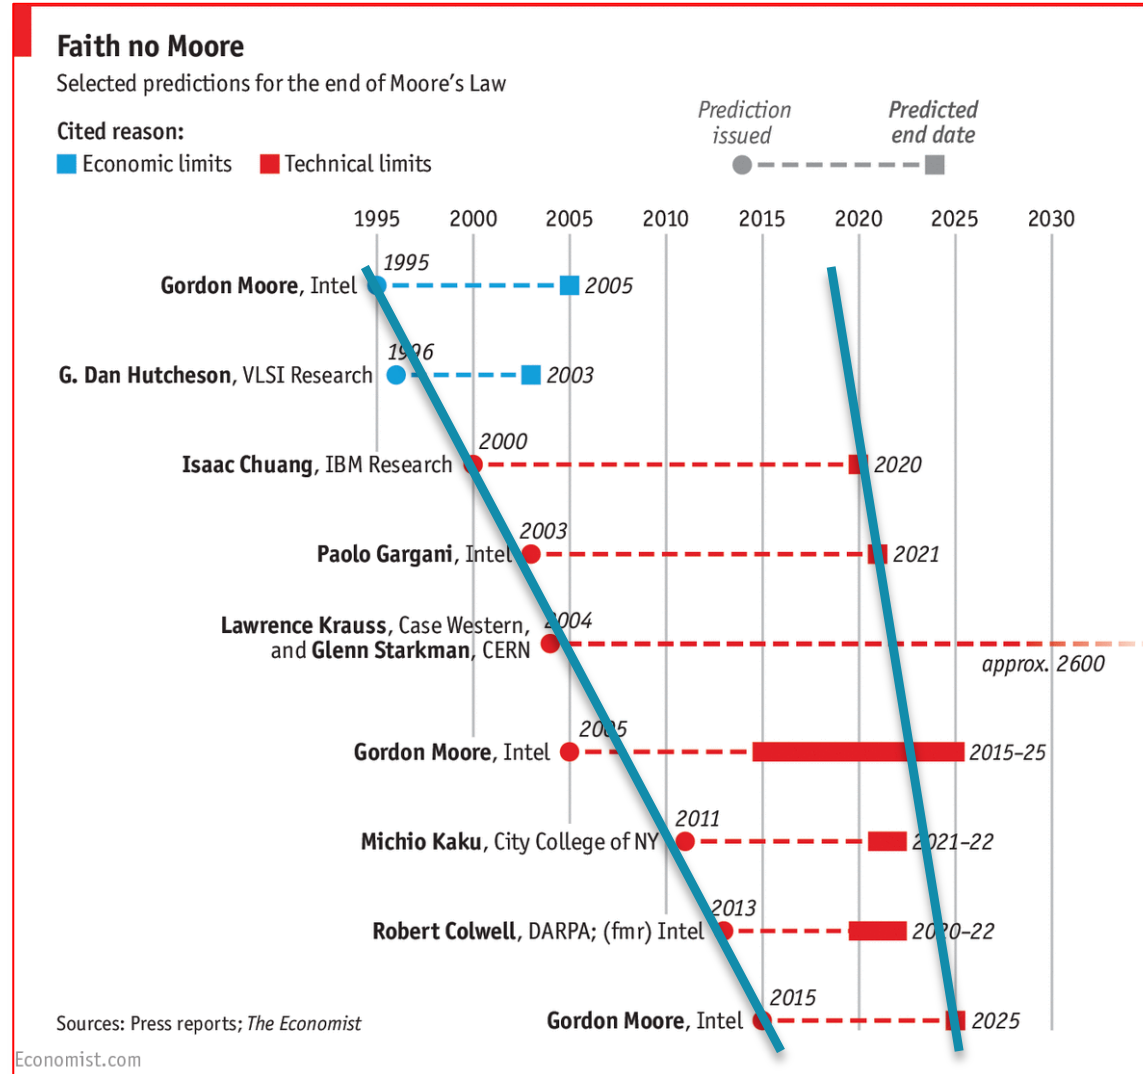
/ ENTERPRISE Intel

 **By Arik Hesseldahl** | [@ahess247](#) | EMAIL | ETHICS

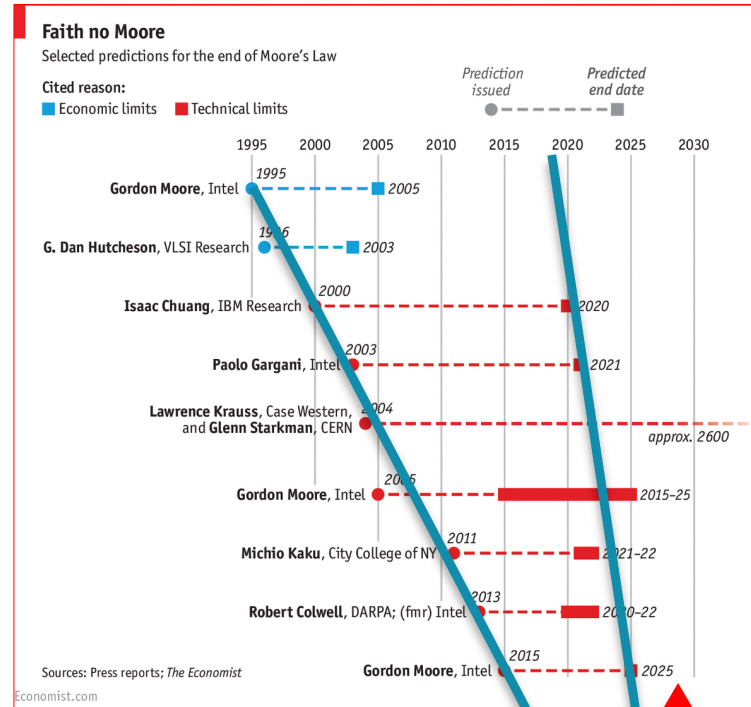
April 15, 2015, 3:11 PM PDT

Handy chart from The Economist:

The Economist




In the spirit of Moore's Law: Extrapolating a few points...



Moore's Law ends: 2029





Economist, April 2015

And here's the crunch: that minimum cost per transistor has been rising since 28nm chips hit the market five years or so ago. That is partly a result of decreasing yields, but also because of the escalating cost of the photolithography equipment needed to fabricate ever-smaller circuits. In short, the cost-effectiveness of chip manufacturing seems to have hit a sweet spot at about 28nm.



Data fueling the mantra

Technology	Area scale factor	Gates/mm ² (KU)	Gate utilization (%)	Used gates/mm ² (KU)	Parametric Yield Impact (Δ from Do yield)
90		637	86	548	97
65	0.57	1109	83	920	96
40	0.52	2129	78	1668	97



Sources: nVidia, ITPC, nov, 2011
Broadcom, IMEC, may 2012
GF, ISS, jan 2013

Data fueling the mantra

<http://www.ibs-inc.net/>

Technology	Area scale factor	Gates/mm ² (KU)	Gate utilization (%)	Used gates/mm ² (KU)	Parametric Yield Impact (Δ from Do yield)	Actual used gates/mm ² (KU)	Gates/wafer (MU)	Wafer price (\$)	Wafer price (Δ)	Cost per million gate (\$)
90		637	86	548	97	531	34009	1358		0.0399
65	0.57	1109	83	920	96	884	56554	1585	17%	0.0280
40	0.52	2139	78	1668	92	1535	98237	1899	20%	0.0193
28	0.54	3946	76	2999	87	2609	166982	2326	23%	0.0139
20	0.56	6992	65	4545	73	3318	212333	2981	28%	0.0140
16/14	0.56	12391	54	6691	61	4082	261222	4205	41%	0.0161

NRE

Alternate scenario:


Technology	Area scale factor	Gates/mm ² (KU)	Gate utilization (%)	used gates/mm ² (KU)	Parametric Yield Impact (Δ from Do yield)	Actual used gates/mm ² (KU)	Gates/wafer (MU)	Wafer cost (\$)	Wafer price (Δ)	Cost per million gate (\$)	Mask set cost	Design Cost
90		637	86	548	97	531	34009	1811		0.0533	800,000	24,000,000
65	0.57	1109	84	932	96	894	57235	2177	16.8	0.0380	1,400,000	40,000,000
40	0.52	2139	82	1754	95	1666	106642	2712	19.7	0.0254	2,000,000	50,000,000
28	0.54	3946	83	3275	94	3079	197035	3500	22.5	0.0178	2,500,000	80,000,000
16/14	0.63	6263	81	5073	92	4668	298723	4375	25.0	0.0146	5,000,000	150,000,000
10	0.63	9942	79	7854	91	7147	457430	5906	35.0	0.0129	7,000,000	202,500,000
7	0.63	15781	77	12151	90	10936	699920	7383	25.0	0.0105	10,000,000	273,375,000

193i steppers for 7nm will be 50% faster than for 28nm

- > 250 wafers per hour = < 15 seconds per wafer
- ~ 100 fields per wafer = < 0.15 seconds per die
- Chuck moves > 1m/second (develops >10G)
- 3 nm X and Y accuracy*


*drop England onto the earth,
align it to a precision of 12cm,
7 times per second.

Etch is half of LE, and etch is
improving also




Mask cost reduction in the pipeline

Shot count comparison



256k-APS (Aperture Plate System)

APS operated at -45kV



256k ($k=1024$) = 262,144 programmable beams of 20nm beam size and 50keV beam energy within 82µm x 82µm beam array field at 6" mask substrate

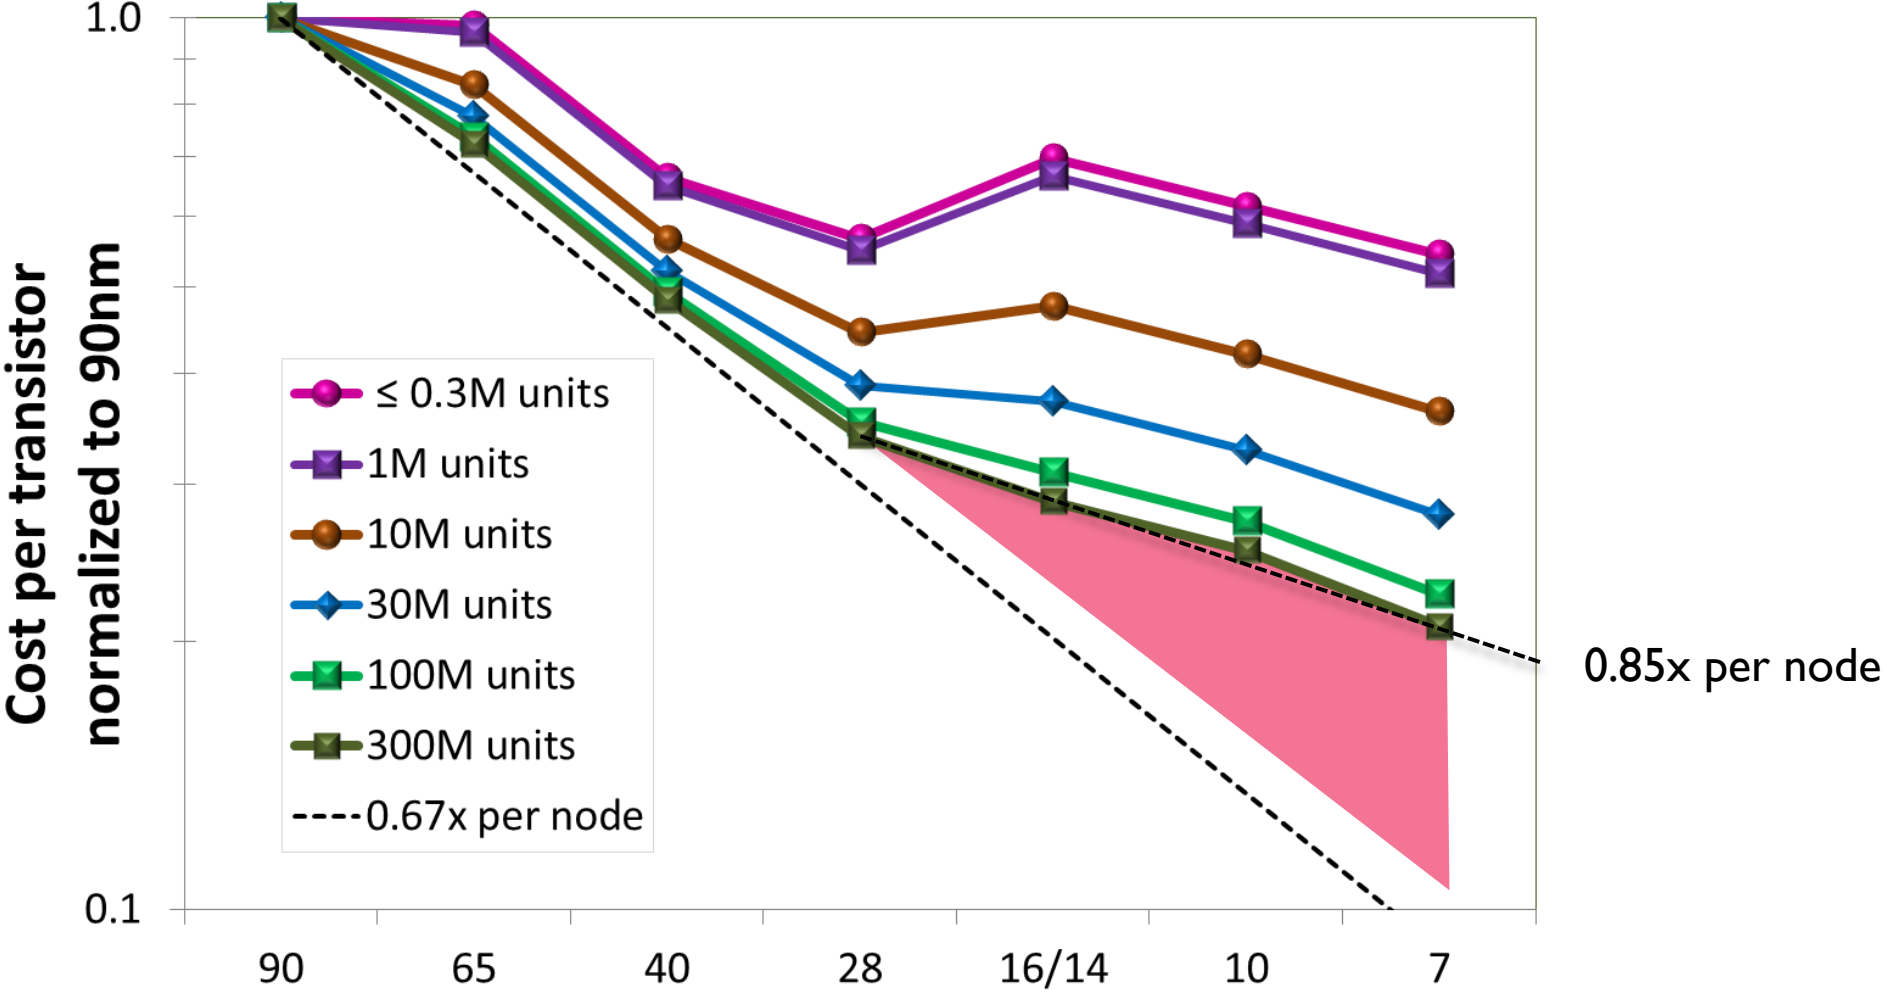


Tokyo Conference Center Shinagawa
June 26, 2012

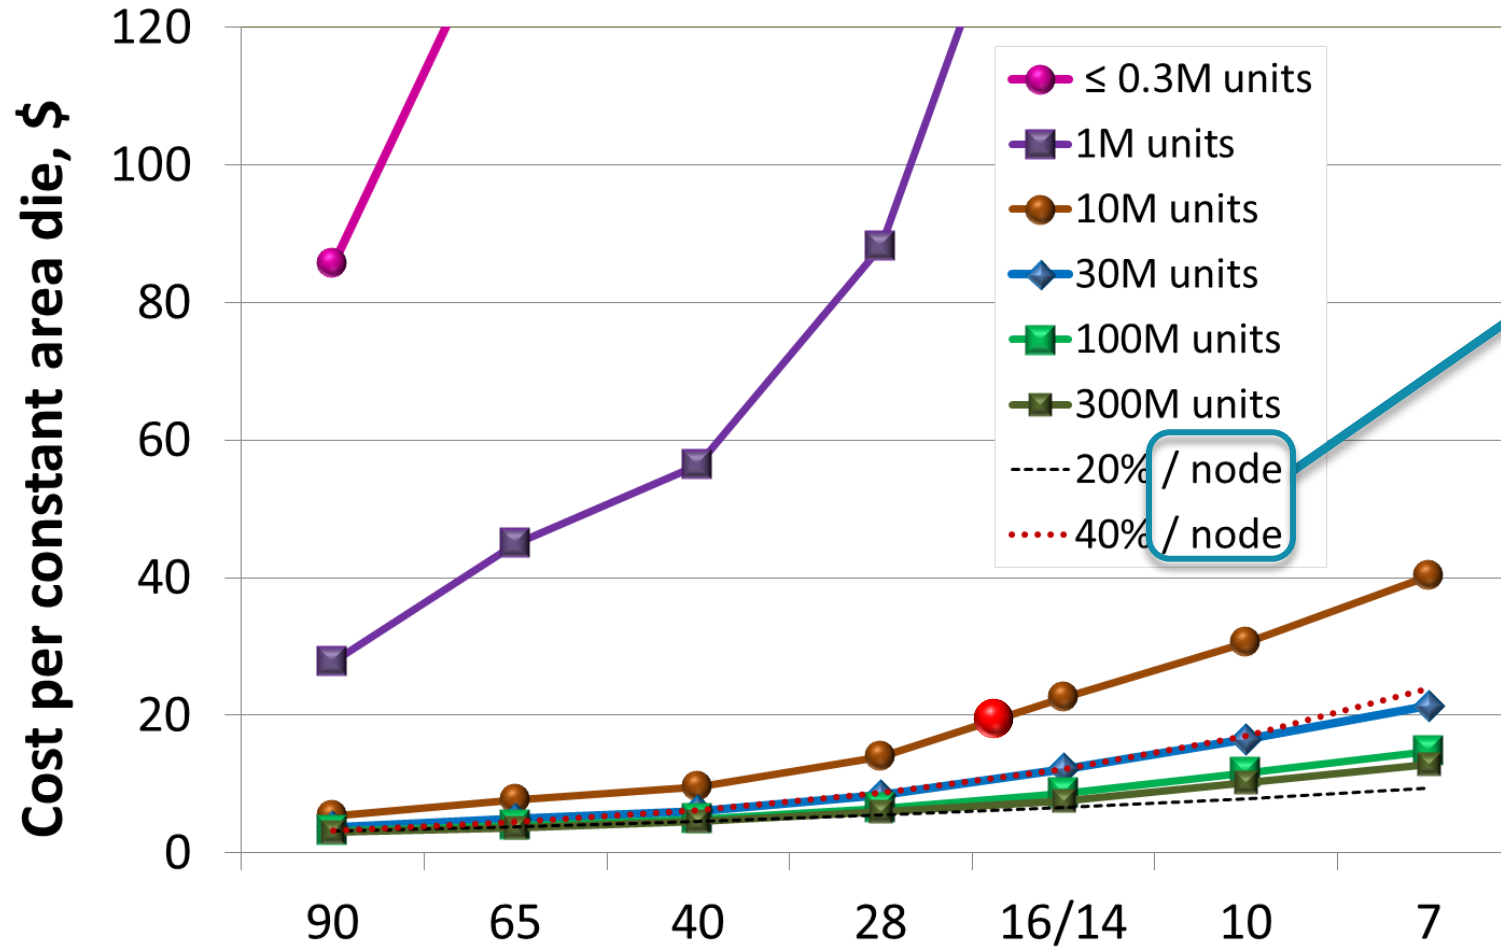
SEMATECH Symposium Japan 2012

Nodes / MDP	Metal 1 #1	Metal 1 #2	Metal 1 #3
20nm /Conventional	635,005	678,823	2,286,066
20nm /MB-MDP	301,057	345,613	1,456,902
Reduction Rate	52.59%	49.09%	36.27%

A cost per transistor scenario




Flipping to cost per die (100mm² die)




per node, not year
(typo in manuscript)

20nm / \$20
Break-even is
10M units


Cost pressure on gate count



The Question is digital




The Answer is analog



Can we do more with less (Moore)?

Today's systems: There is not just one problem



Can we do more with less (Moore)?

There are Three Problems

...and performance ...

in a power budget



Electricity Bill

Cooling Bill





Battery Capacity

Touch Temperature




Form Factor

Energy Harvesting



Cost



System scaling today:

1. Everything has a power budget,
2. But wants as much performance as possible

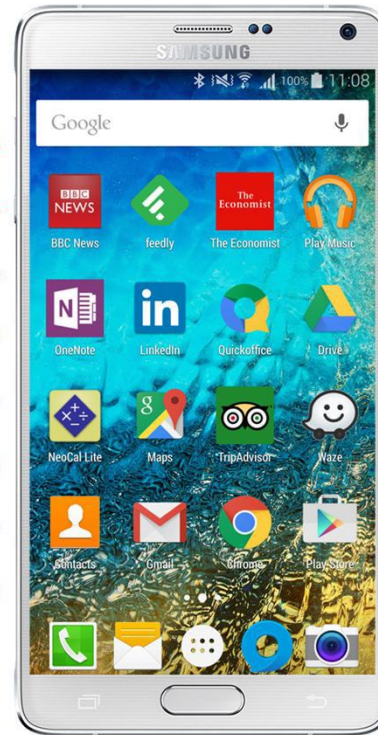
2009



Motorola Droid (2009)
By Creative Tools. Watermark removed by User:Atall [CC BY 2.0
(<http://creativecommons.org/licenses/by/2.0/>), via Wikimedia Commons

Display	5x
Camera	4x
Connectivity	20x
Sensors	3x
Video	34x
CPU	17x
GPU	40x
Memory Bandwidth	16x
Moore's Law	5-6x


2014



Samsung Galaxy S5 (2014)
By GalaxyOptimus (Own work) [CC BY-SA 3.0
(<http://creativecommons.org/licenses/by-sa/3.0/>), via Wikimedia Commons



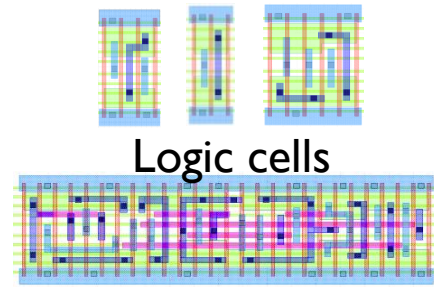

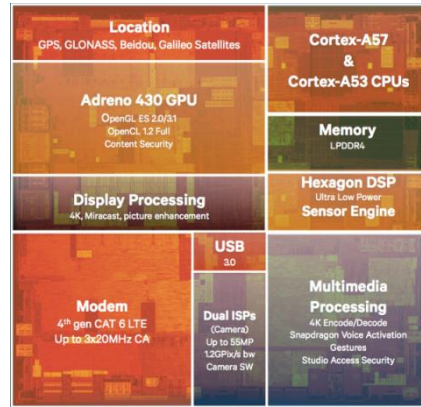

With more transistors,
improve performance within
an energy budget:

- Multi-core parallelism
- Out-of-Order
- Branch prediction
- Pre-fetching
- Cache hierarchy complexity
- Multi-Threading




R.Aitken, 2015 IEDM short course

Interdependent chain to make end product




Next Generation Transistors at IEDM: What you think they look like



2.1 First Demonstration of Ge Nanowire CMOS Circuits:


2.4 Germanium-based Transistors for Future High Performance and Low Power Logic Applications

Next Generation Transistors at IEDM: What they actually look like



Evening Panel: Emerging Devices – Will they solve the bottlenecks of CMOS?

The shrinking FET



Intrinsic FET capacitance is a minority!






Fig. 6. Breakdown of intrinsic & extrinsic switching capacitances for typical post-28nm simple logic gates (INV and NAND), post-layout parasitic extraction and simulated with SPICE. IMEC: IEDM 2013



HS Wong, et al., SISPAD 2009

The shrinking FET



- Electrostatics for gate length reduction:
- previously beneficial in reducing C
 - now beneficial in reducing R_C

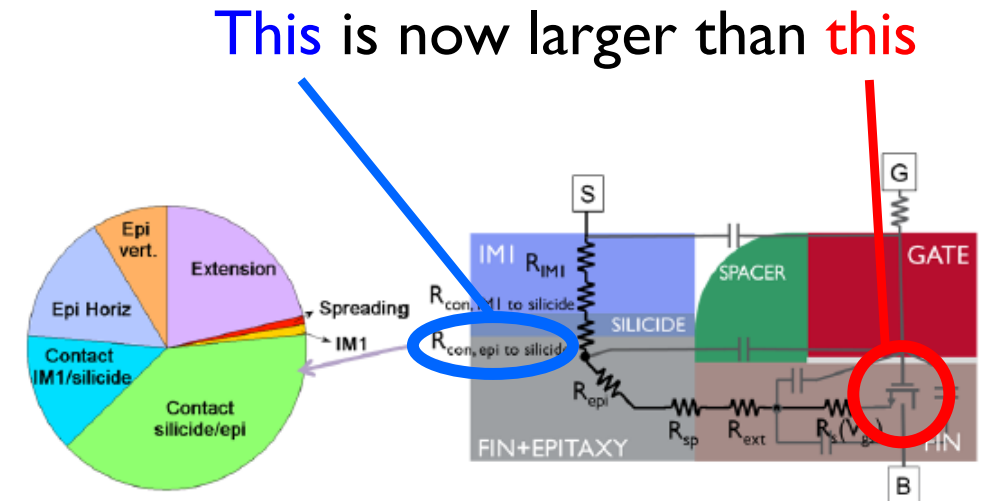





Fig. 12. Transconductance g_m of an n-type 10nm-Node FinFET as a function of gate bias evaluated with the processing assumptions from the

21.7 $1.5 \times 10^9 \Omega \text{cm}^2$ Contact Resistivity on Highly Doped Si:P Using Ge Pre-amorphization and Ti Silicidation.

Non-scaling parasitics




Session 22: Nano Device Technology – Steep Slope Transistors

22.2 First Foundry Platform of Complementary Tunnel-FETs in CMOS Baseline Technology for Ultralow-Power IoT Applications

22.6 Sub-60mV-Swing Negative-Capacitance FinFET without Hysteresis

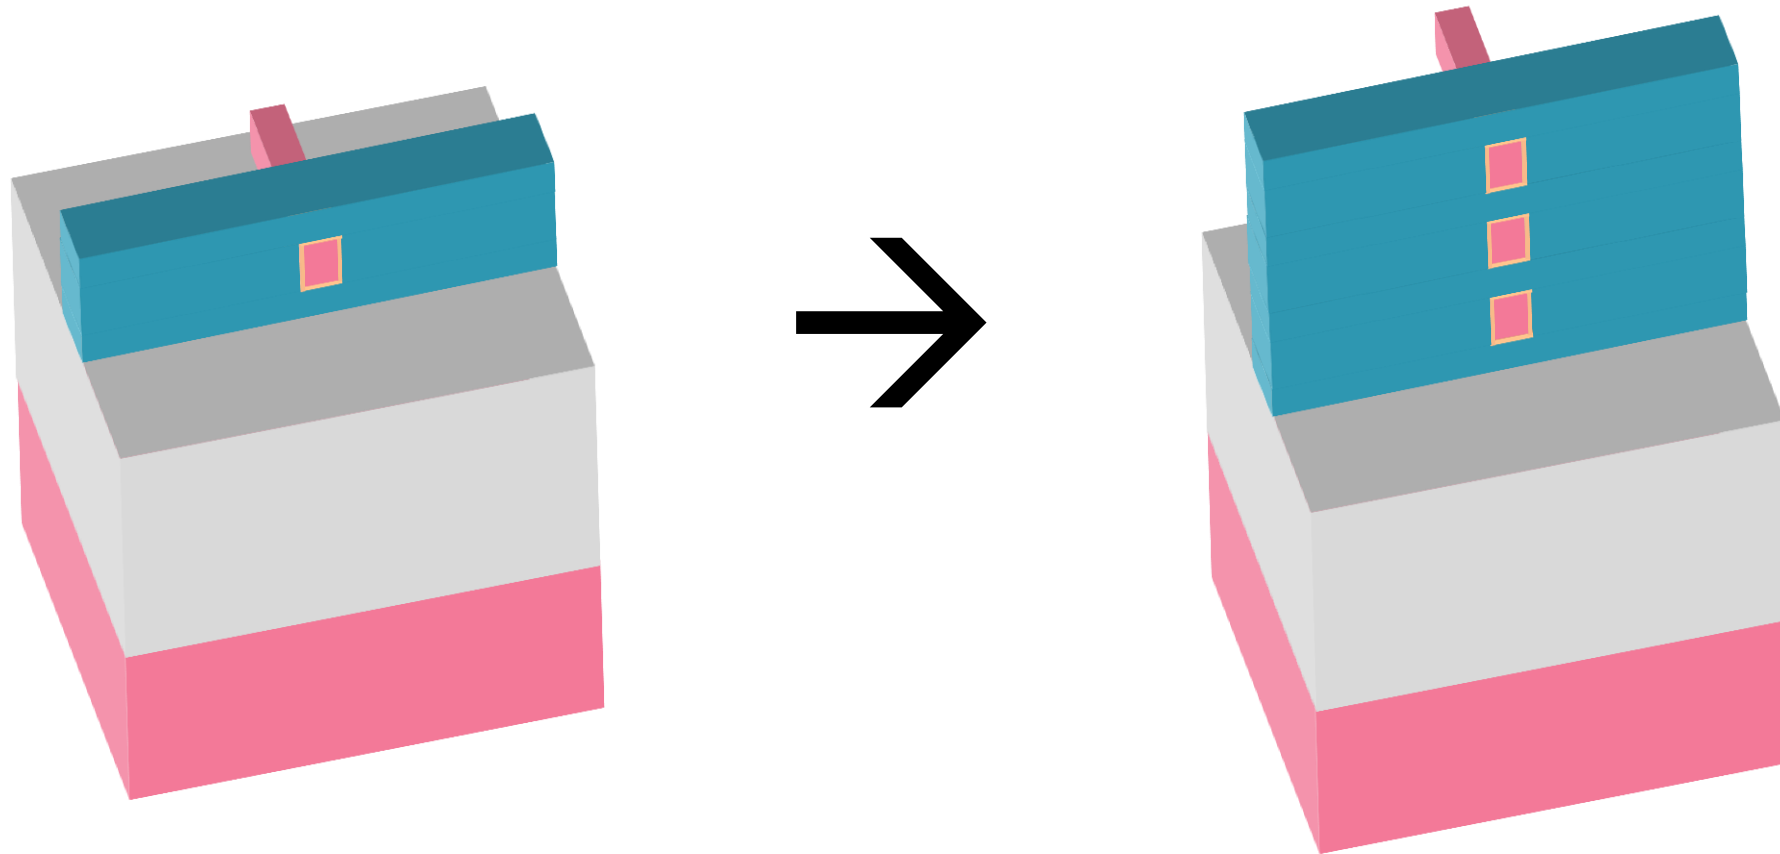
FinFET performance scaling secret: “3D factor”

Planar → 1st generation FinFET → 2nd generation FinFET:




- Part of performance is improved electrostatics
- But part is simply more folded width
- “fin depopulation”: Can’t design with one-fin devices
- More 3D = more parasitics

Gate All-Around Horizontal Nanowire



Chasing fin height, or back to 2D?



Yang, et al., GLSVLSI 2010,
pp 263-268

Meric, et al. IEDM 2010, pp. 556-559

Hunt, et al. *Science*, June 2013, v340, n6139

Liu, et al., *Nature Nanotechnology* 8,
2013, pp 119-124



12.1 Phonon-Limited Performance of Single-Layer, Single-Gate Black Phosphorus n- and p-type Field-Effect Transistors,

12.3 Designing Band-to-Band Tunneling Field-Effect Transistors with 2D Semiconductors for Next-Generation Low-Power VLSI,

12.4 How Good is Mono-Layer Transition-Metal Dichalcogenide Tunnel Field-Effect Transistors in sub-10 nm?

12.7 Understanding the Nature of Metal-Graphene Contacts: A Theoretical and Experimental Study,

Real V scaling: to the mV's:



32.5 Spintronic Majority Gates

32.6 Spintronic Logic Circuit and Device Prototypes

150°C?
Power networks?
Signal integrity?

Moore's Law and Dennard scaling

“shrinking dimensions on an integrated structure makes it possible to operate the structure at higher speed for the same power per unit area” [1]



Table 1
Scaling Results for Circuit Performance

Device or Circuit Parameter	Scaling Factor
Device dimension t_{ox} , L , W	$1/\kappa$
Doping concentration N_a	κ
Voltage V	$1/\kappa$
Current I	$1/\kappa$
Capacitance $\epsilon A/t$	$1/\kappa$
Delay time/circuit VC/I	$1/\kappa$
Power dissipation/circuit VI	$1/\kappa^2$
Power density VI/A	1


Also from The Economist: September 2015

Compactness is less important for another fast-growing branch of the information-technology business, **cloud computing**. In cloud-service providers' cavernous data centres, space is not at a premium, the way it is inside the latest iPhone. What increasingly matters most to cloud providers is **energy efficiency**: how much power their racks of servers consume, and how they can keep them sufficiently cool to ensure that their chips do not fry.


Fortunately, one of the corollaries of Moore's law is that the **energy efficiency of transistors follows the same exponential law, doubling around every two years**. And like the law itself, it's not quite dead yet.

Dark Silicon

We get more transistors, we just can't afford to turn them all on



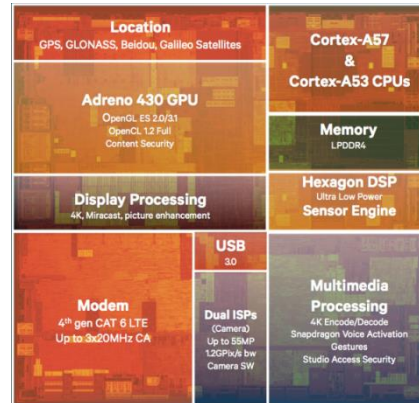
■ Sinha, Cline, Yeric, Chandra, Cao*, ISLPED 2012



Design effort reduces effect on products:

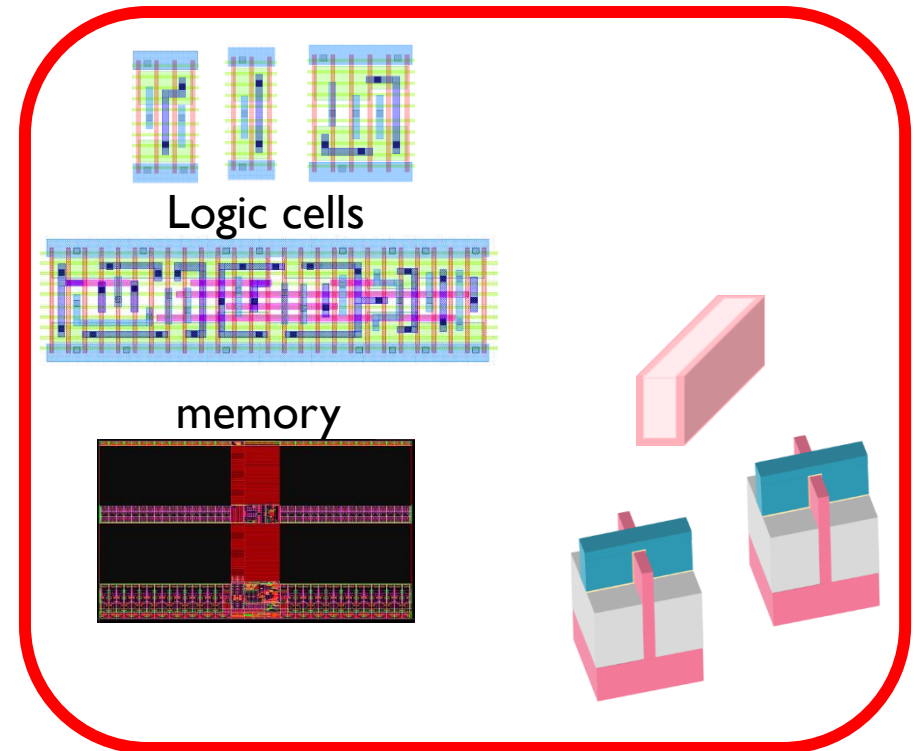
- Power and clock gating
 - Dynamic Voltage/Frequency Scaling (DVFS)
 - Multi-core
 - Memory assist
- } Spend transistors to buy power

Can we get more out of less (Moore)?



Example: Qualcomm 810

Design-Technology Co-Optimization (DTCO)



Gate pitch, logic density, and frequency

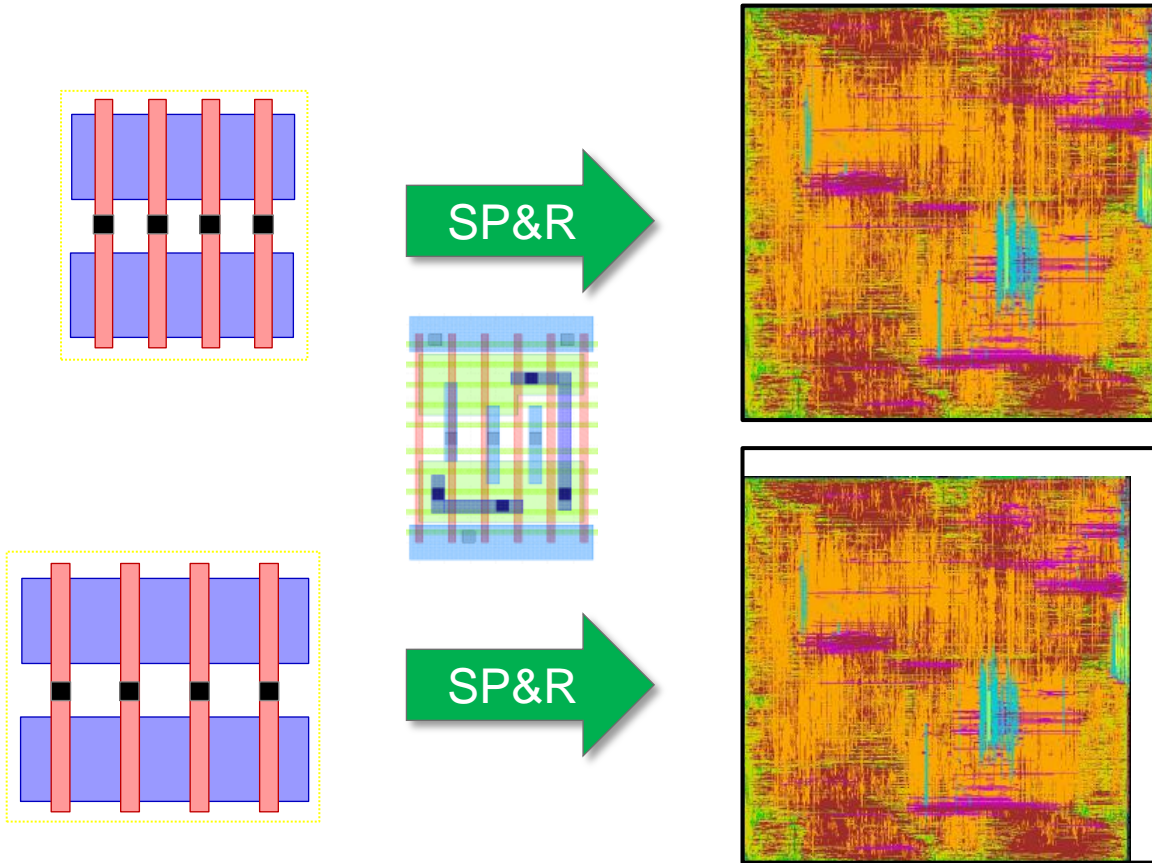
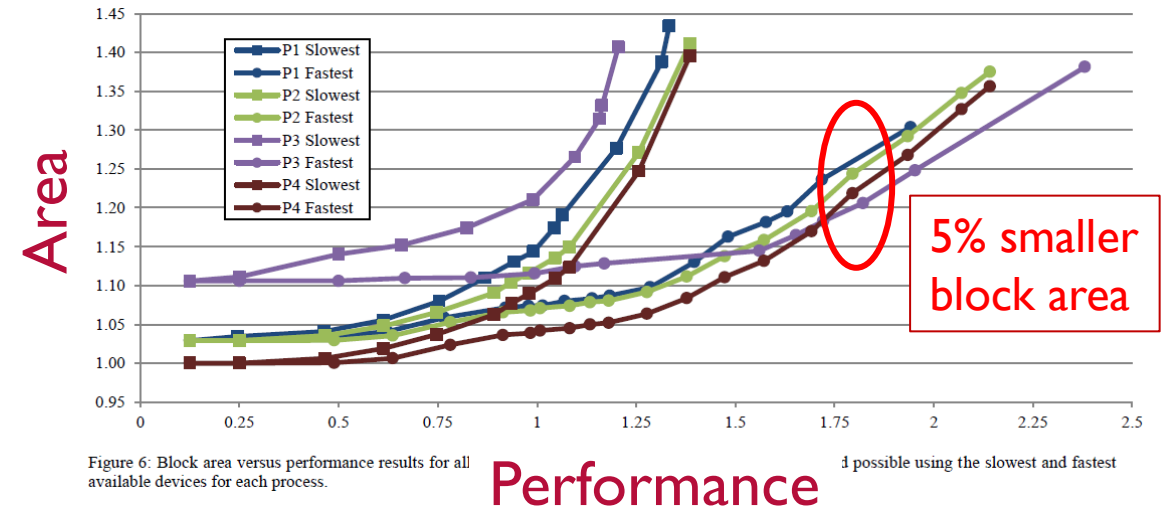


Table 1: Summary of Process Variations Benchmarked

Process Version	Relative Poly Pitch	Other Features
P1	1	-
P2	1.076	-
P3	1.153	15% larger gate pitch
P4	1.076	2 pitch smaller low drive flops

M. Frederick, IEDM 2014



Strain volume, contact-gate capacitance
(and complicated litho issues)

Gate pitch, logic density, and frequency

- Inverse Moore Scaling!
- Transistor performance can be traded for area reduction
- Perceived value of performance vs. area is dependent on application targets




Table 1: Summary of Process Variations Benchmarked

Process Version	Relative Poly Pitch	Other Features
P1	1	-
P2	1.076	-
P3	1.153	-
P4	1.076	2 pitch smaller low drive flops

15% larger gate pitch

M. Frederick, IEDM 2014





Figure 6: Block area versus performance results for all available devices for each process.

Strain volume, contact-gate capacitance
(and complicated litho issues)

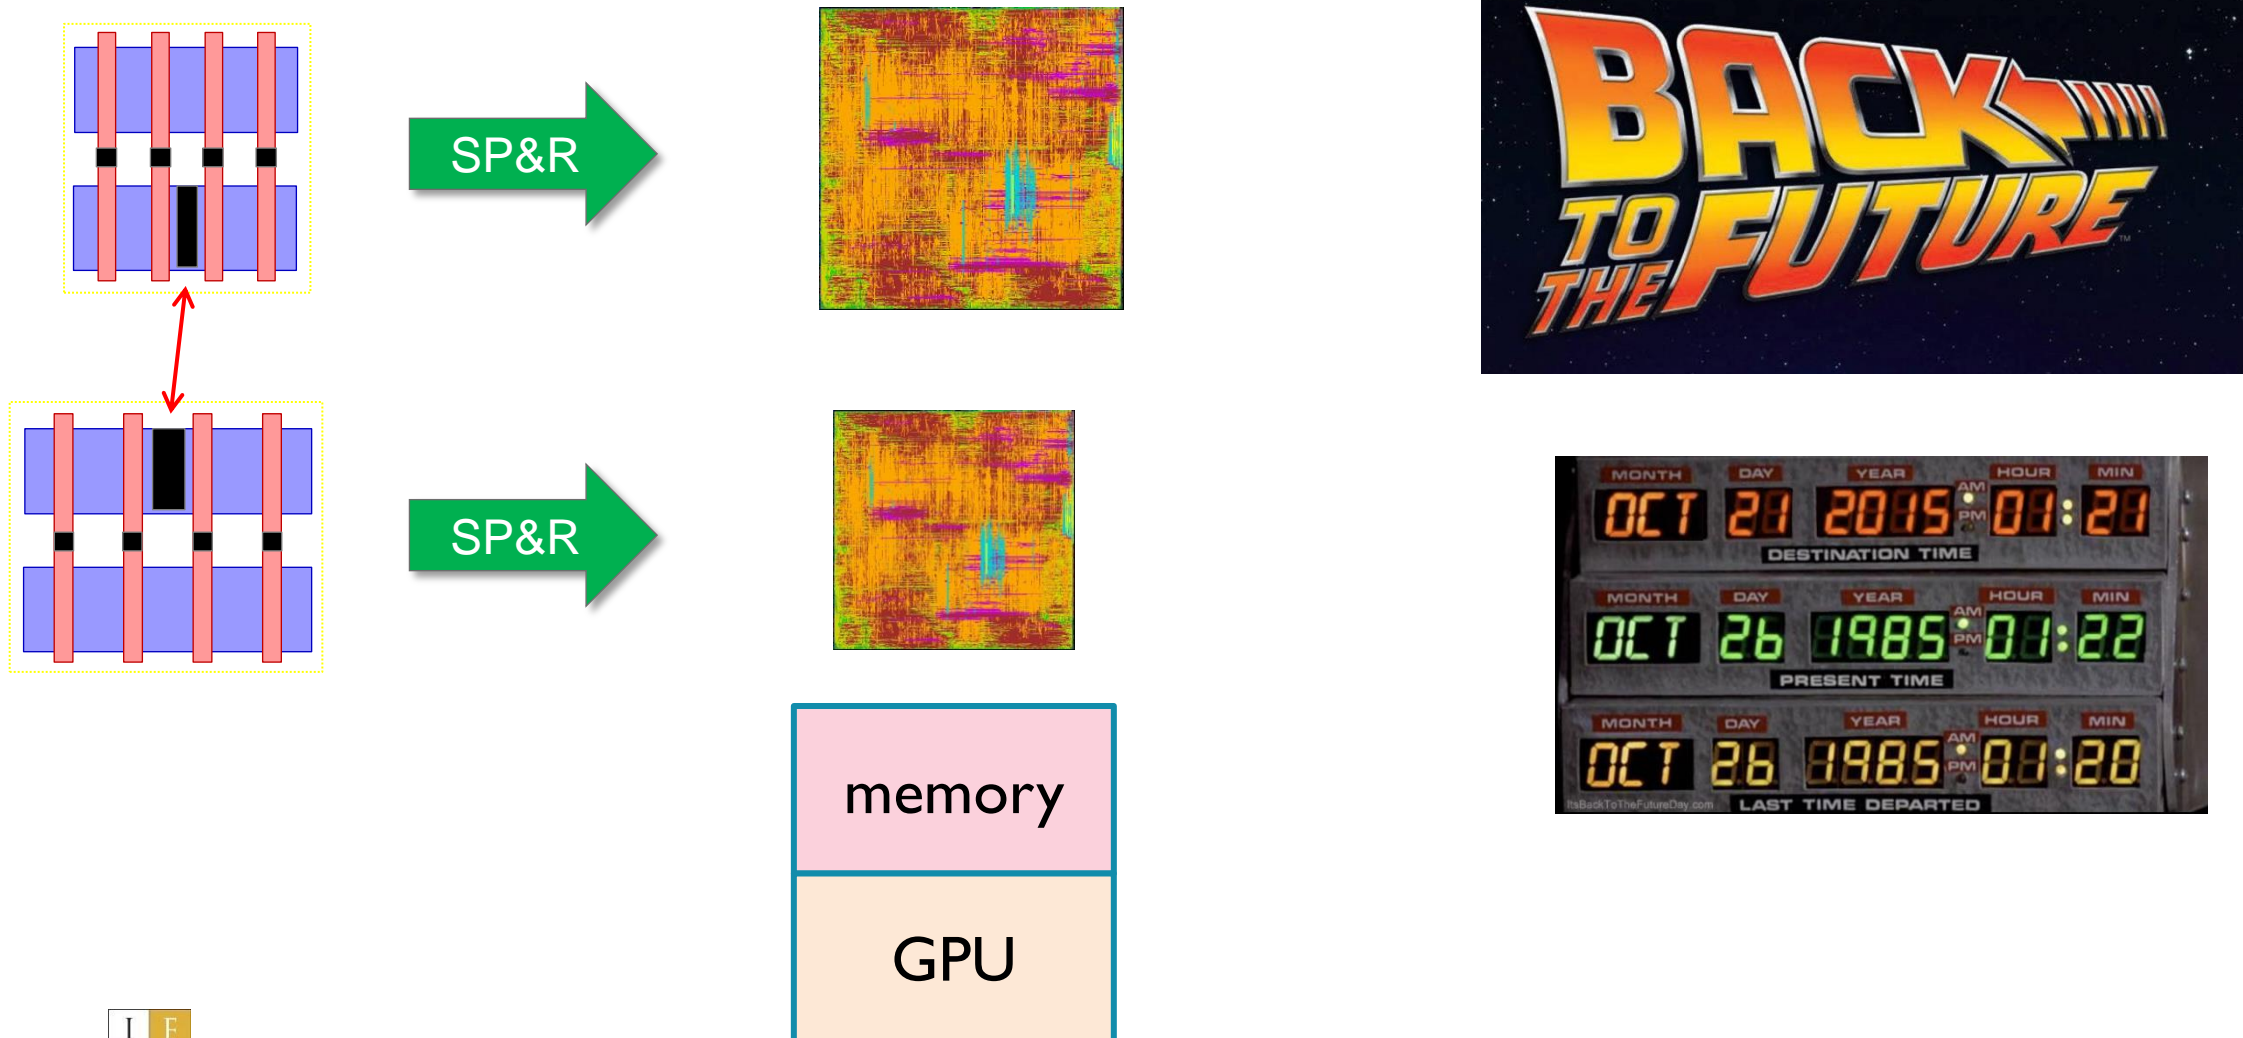
BEOL cost: Good news for FEOL (?)



Retirement planning
advice:

diversify your
investments


Can we support multiple gate pitches on one die?




The lost half node

Deep sub- λ lithography has forced regularity. This has cost density.

'Vintage' Sequential Logic




'State of the Art' Sequential Logic



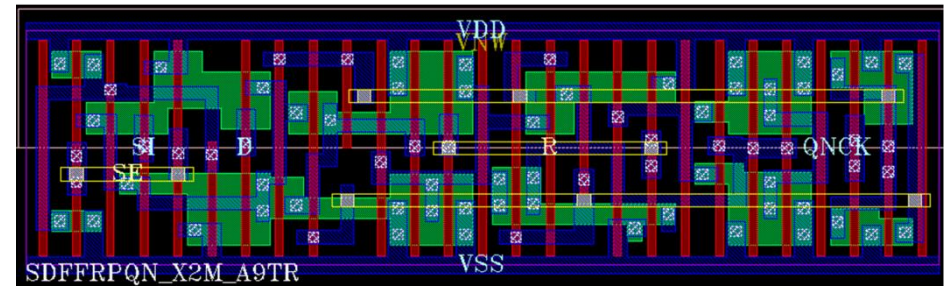
- Lost 'stacked devices'
- Forced 'diffusion tuck under'
- Forced 'double diffusion' break
- Lost 'tapered devices'
- Forced 'coarse granularity' in L and W

Cumulative density loss: ~20-40%

65nm





32nm



Liebmann, Pietromonaco, SPIE_8684_12


Middle Of Line (MOL)

- Cost bump: 28 → 16/14
- Adds RC
- Adds gate efficiency pressure




Middle Of Line (MOL)

- Anytime you can avoid a V0, you free up routing resources, and reduce chip size



-  1st pass gate
-  2nd pass gate (does not have to be ALD WF, just a metal)




Expectations: V_T options

~ 2 Moore's Law Nodes





Pleak (mW)


More V_T options = lower power products




Device options = reduced design cost



Performance vs. Leakage Power



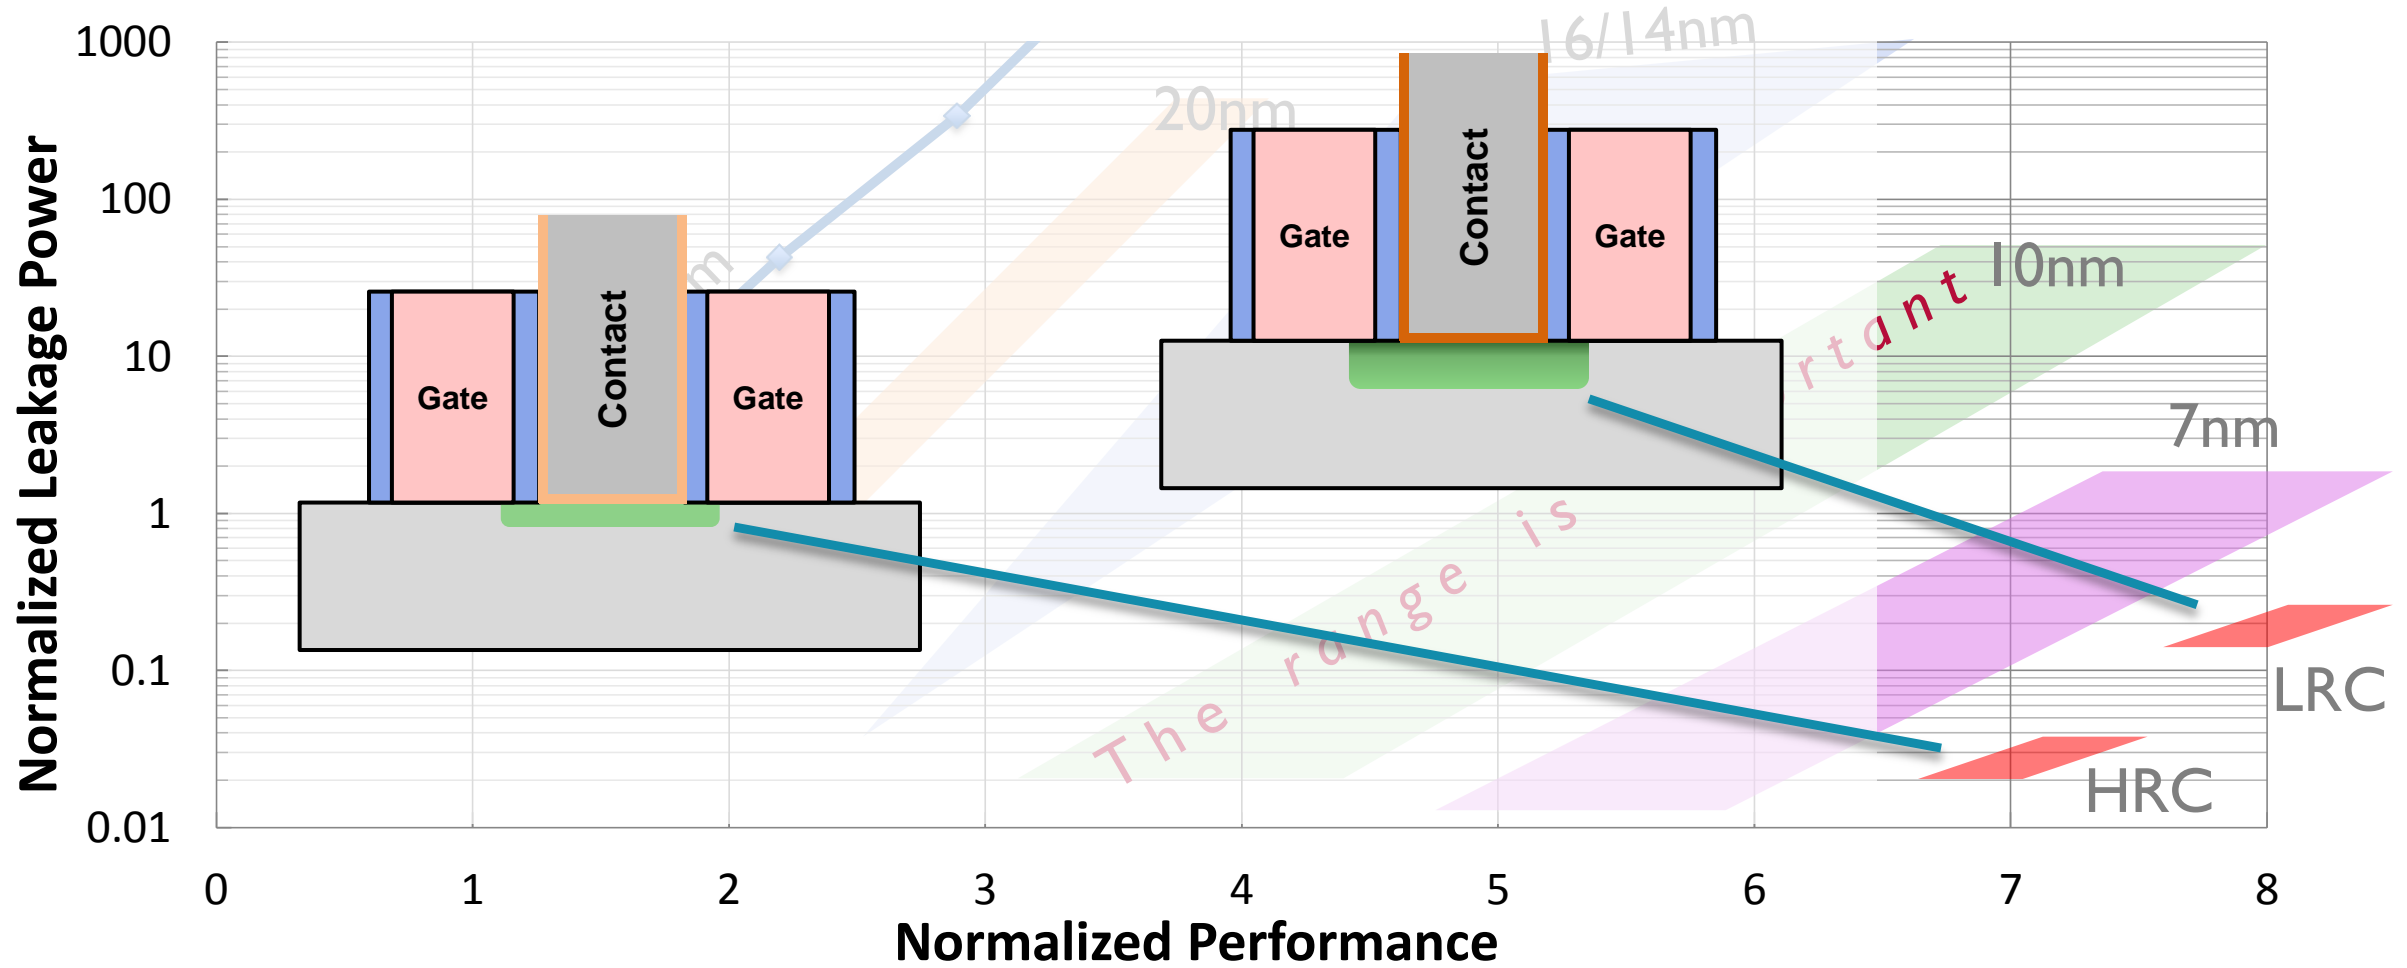
What FET heterogeneity can be cost effective?



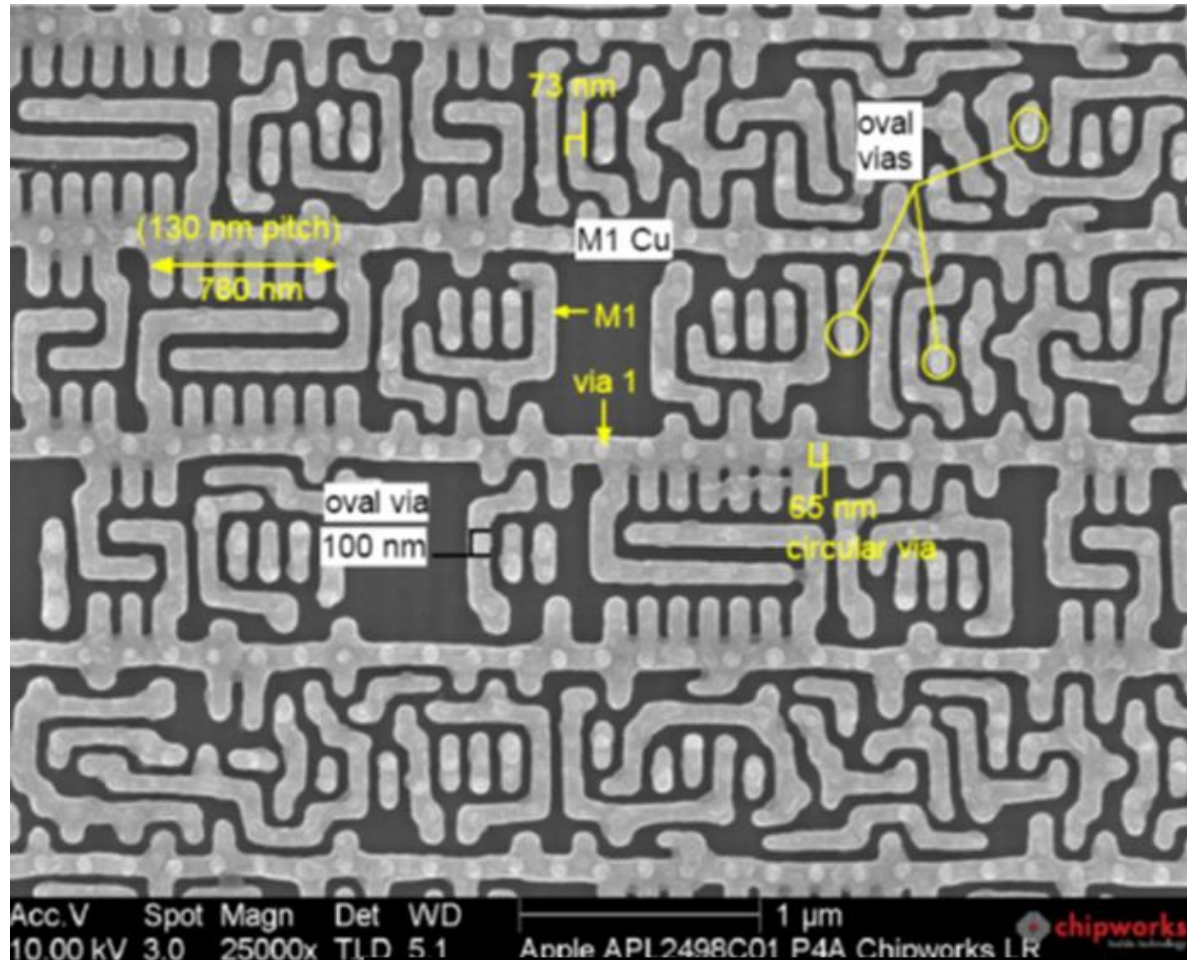
18.1 Energy-Delay Performance Optimization of NEM Logic Relay



An R_C equivalent to V_T ?



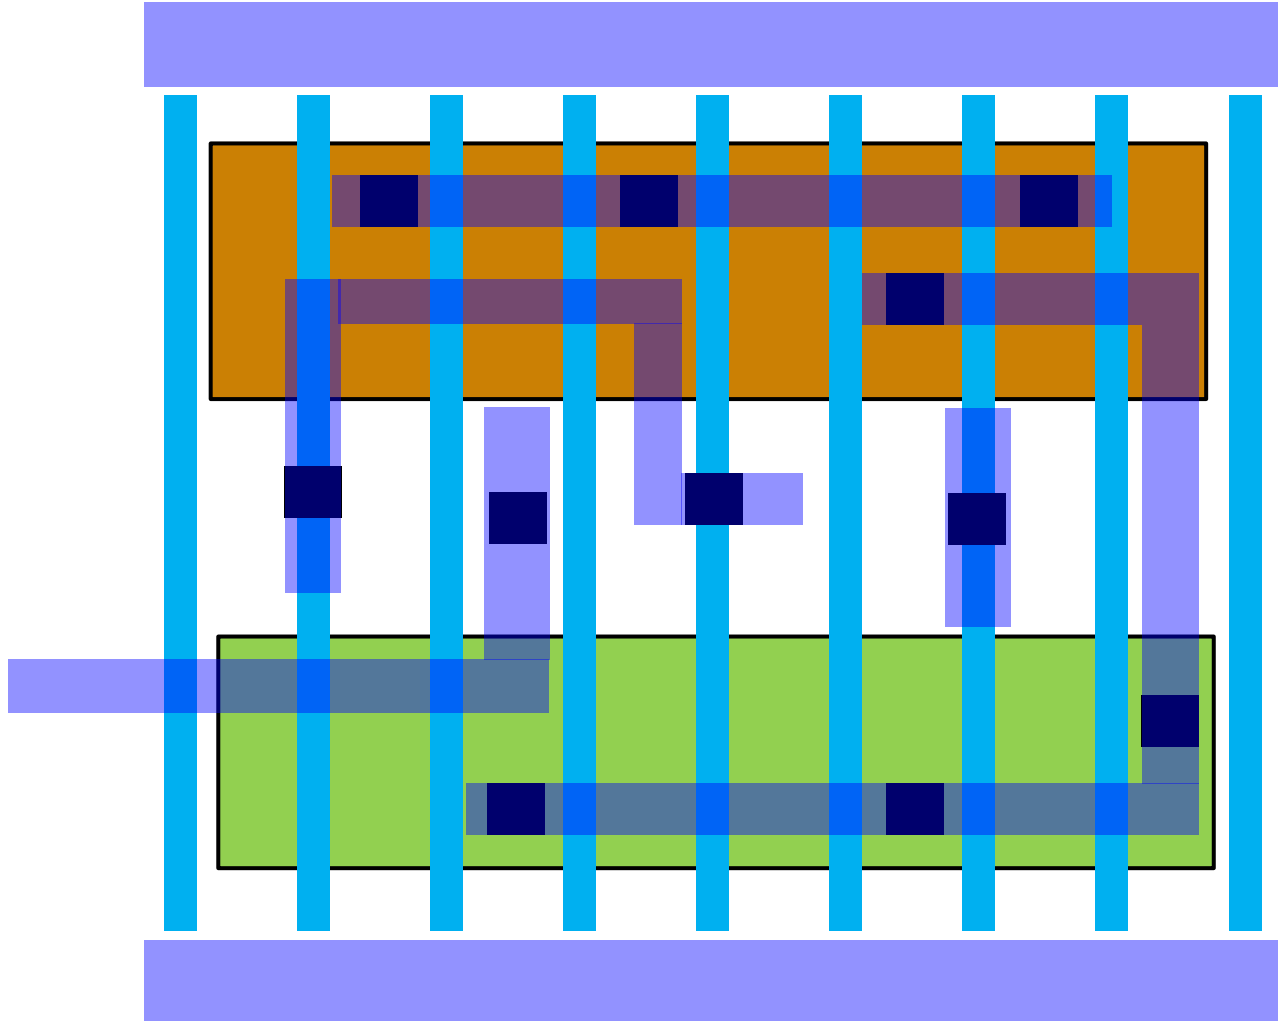


Lots of curvy M1 shapes: The bane of lithography.

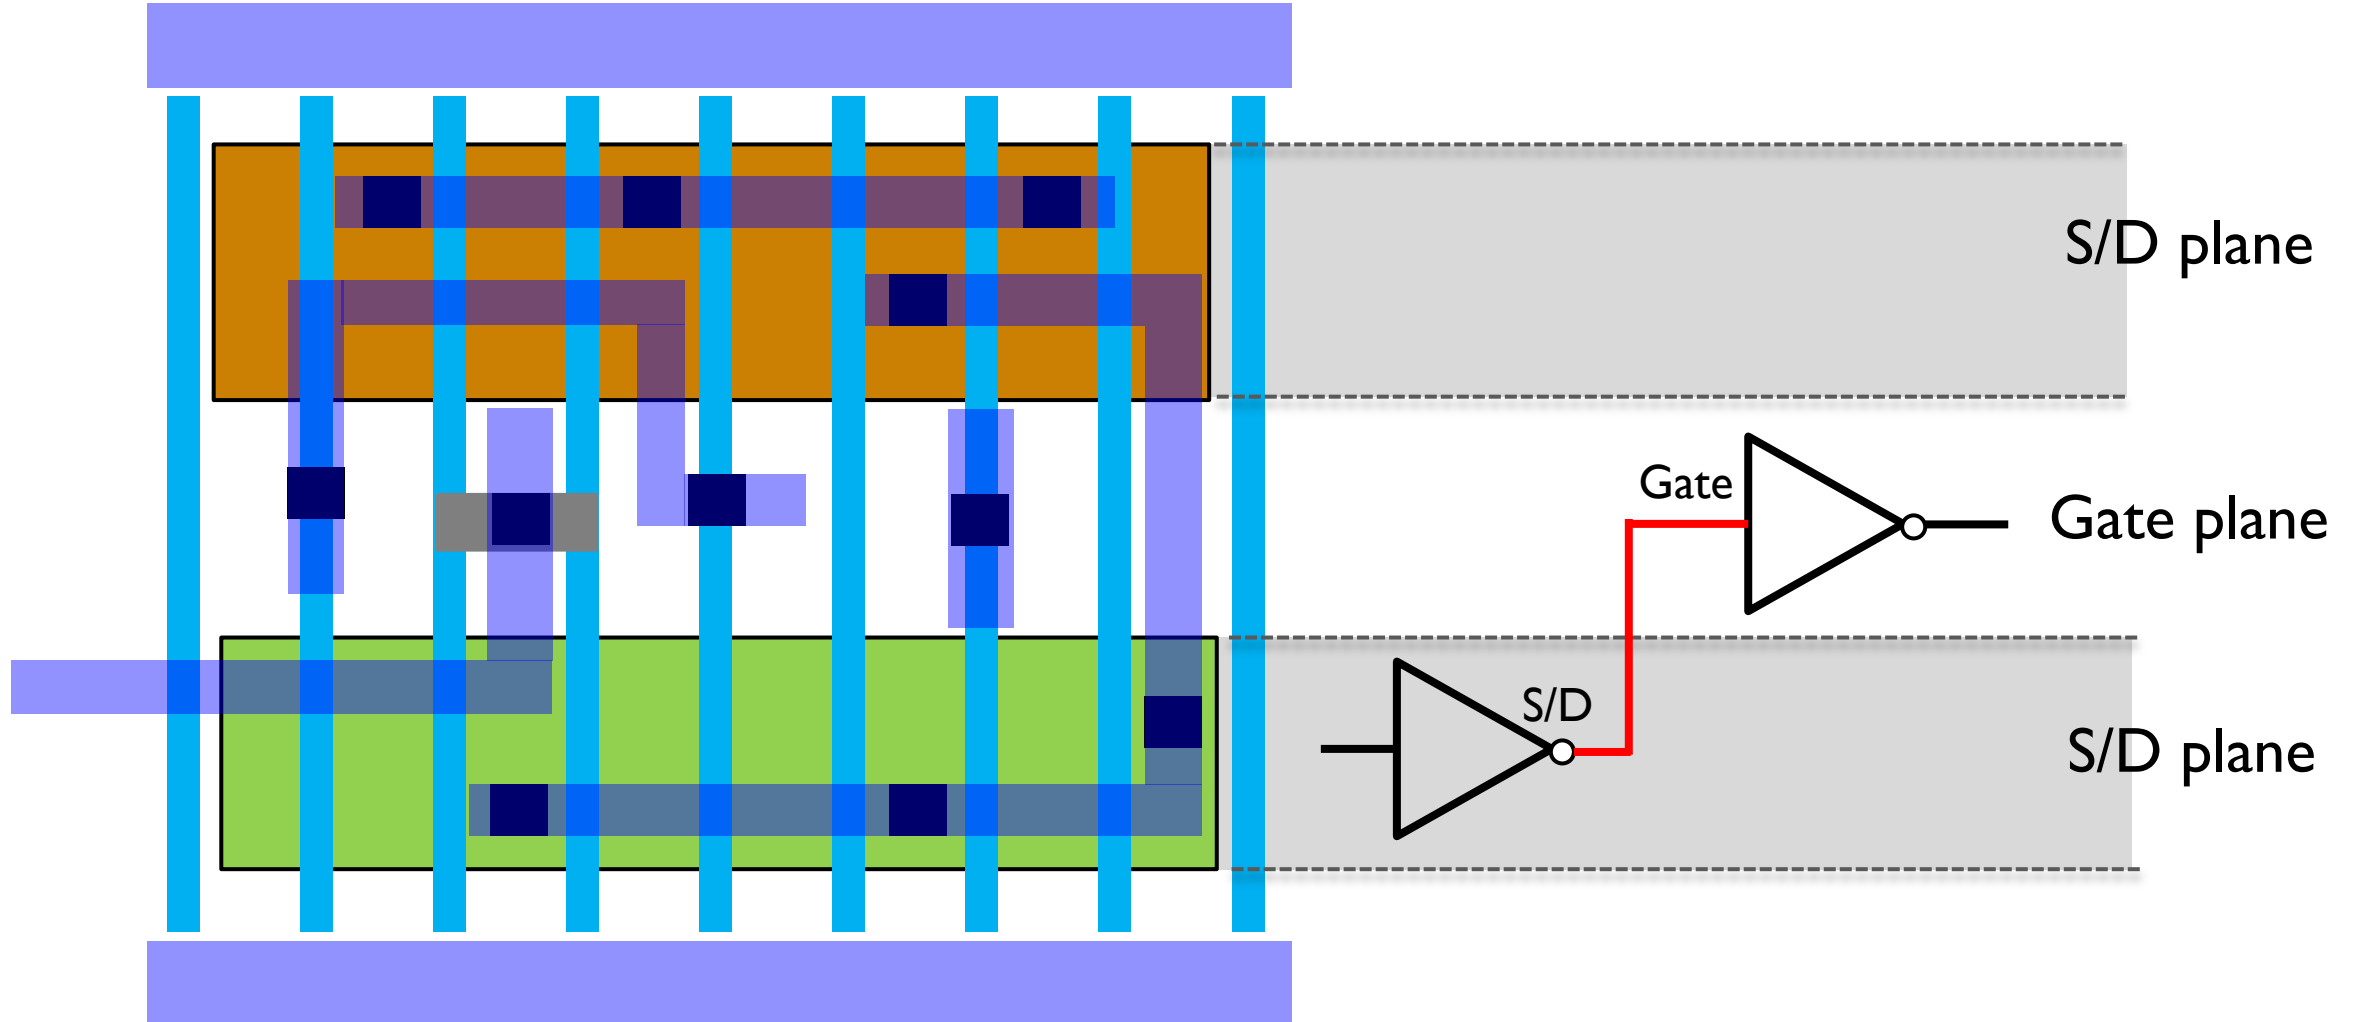


<http://www.chipworks.com/blog/technologyblog/2012/07/31/samsung-32-nm-technology-looking-at-the-layout/>

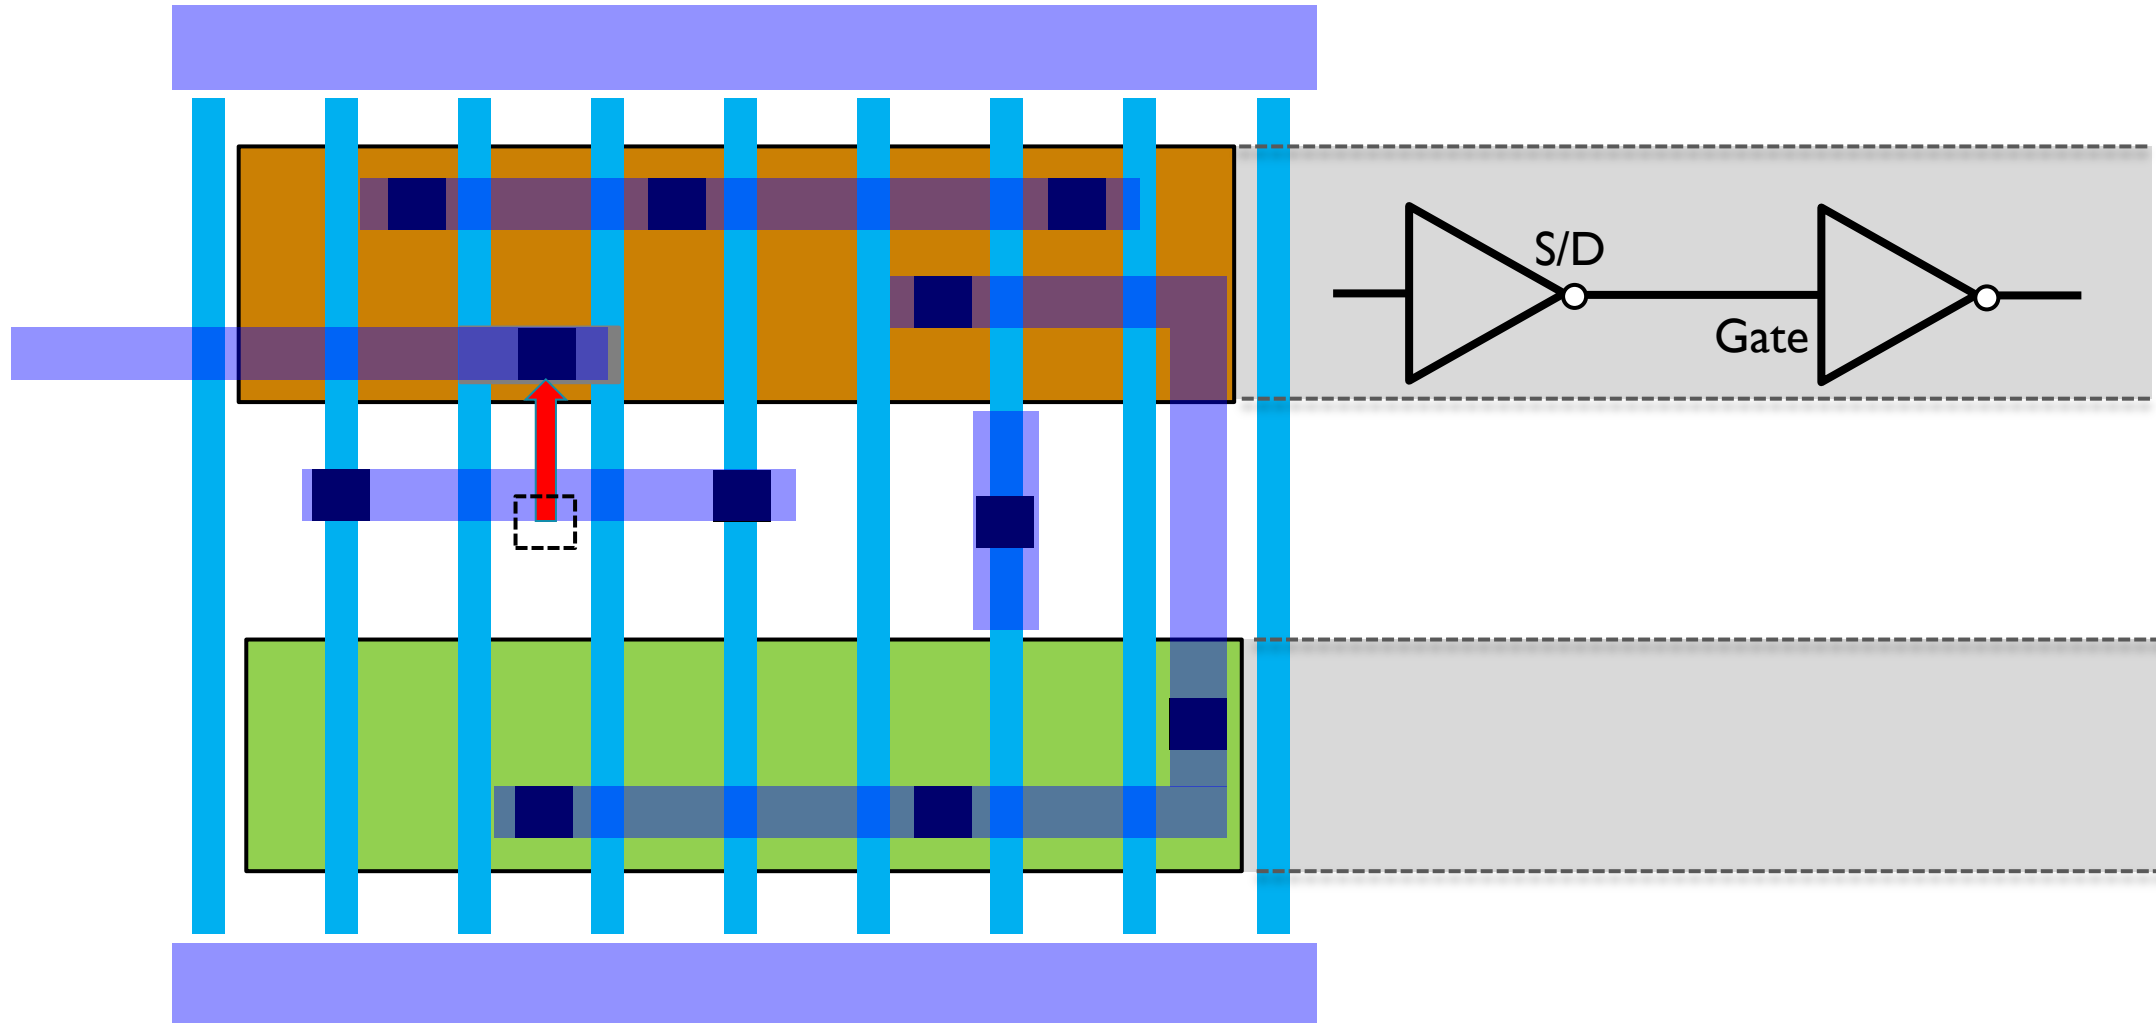
What does that have to do with transistors?



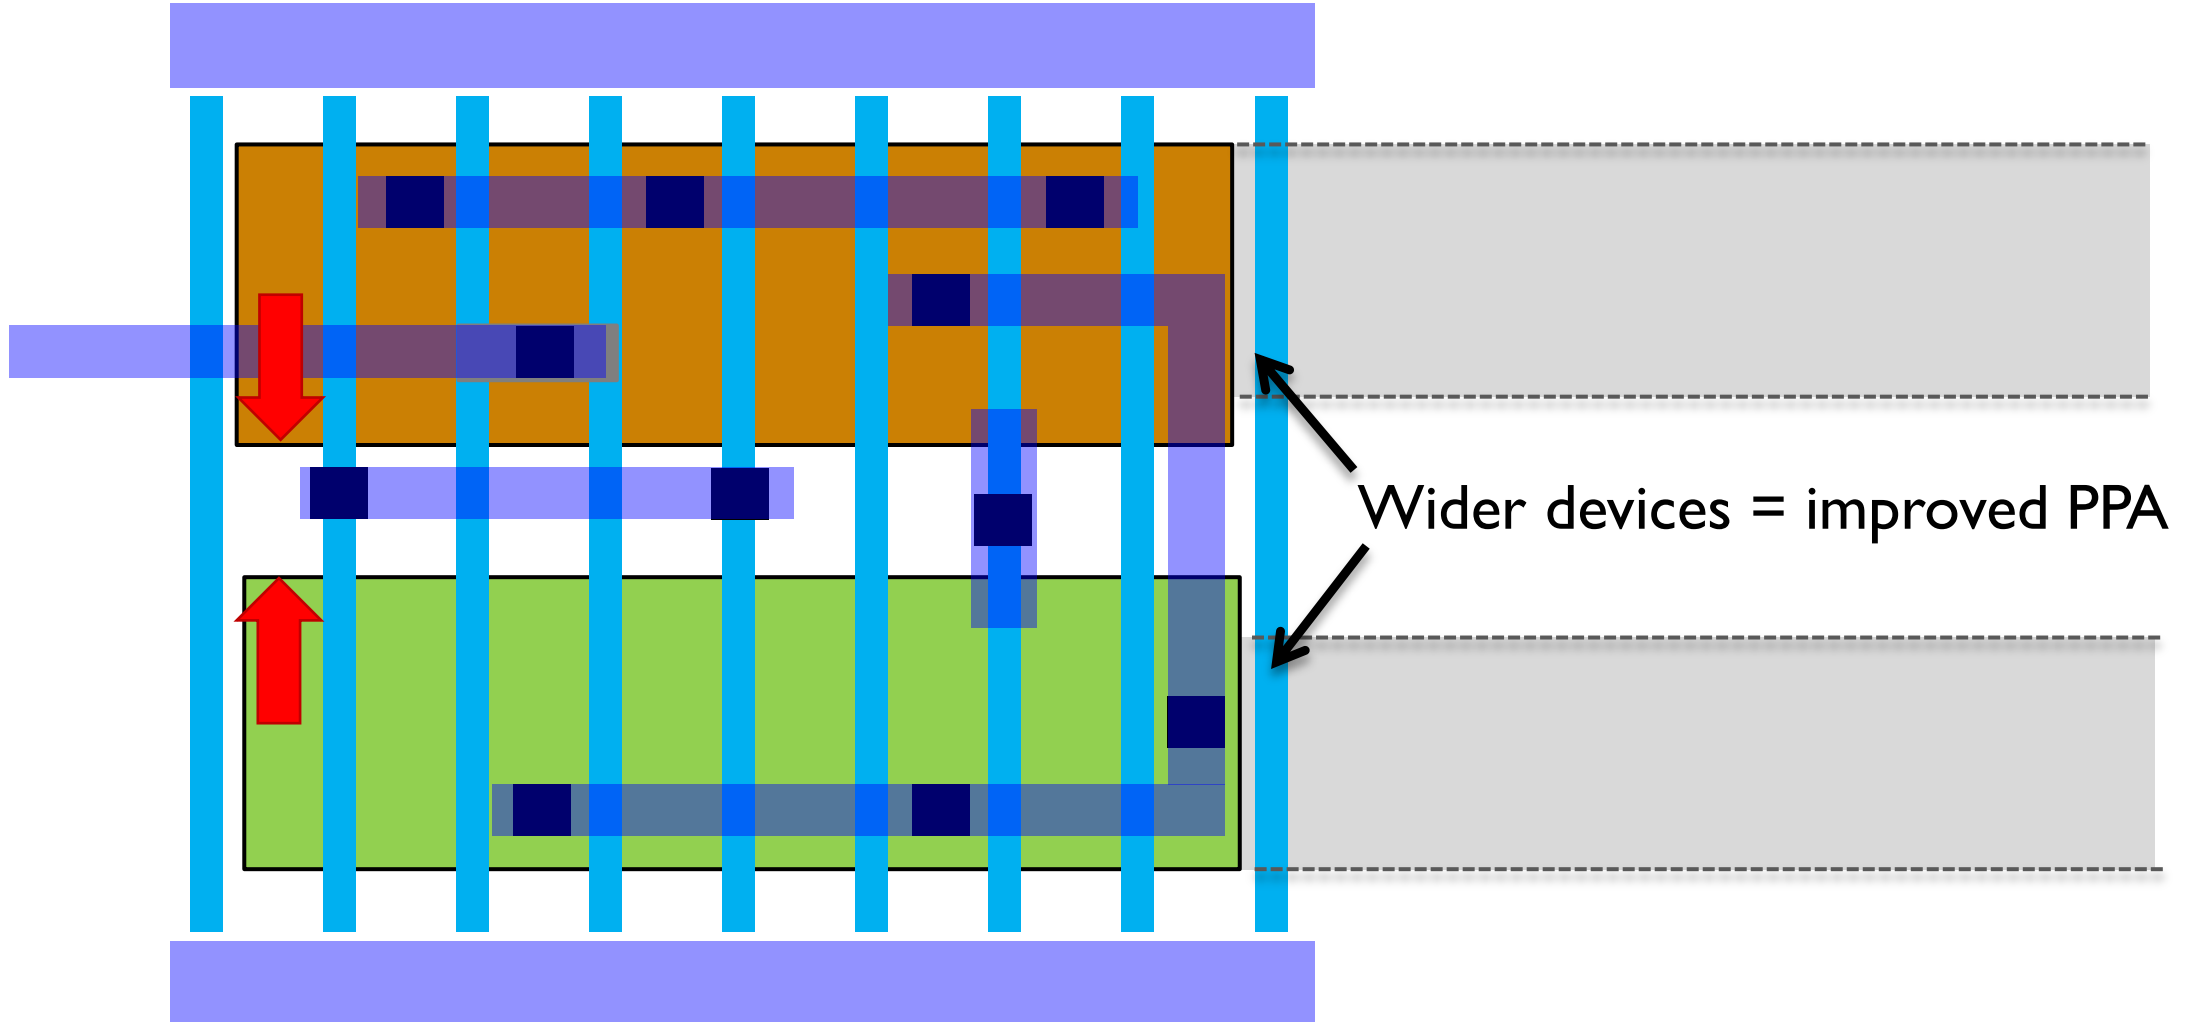
What does that have to do with transistors?




A potential improvement: gate contact over active



Plus performance improvement (or area shrink)



Another important scaling law:



$$\text{Variation} \propto \frac{A_{VT}}{\sqrt{W \cdot L}}$$



M. Pelgrom

Red process faster than blue process!!

Investing in reducing variability may be more helpful than mobility

IEDM 2015 and variation

Line Edge Roughness. Work Function variation. Random Dopant Fluctuation

Session 11: Circuit Device Interaction – **CMOS Scaling and Circuit/Device Variability**

15.3 Novel Wafer-Scale Uniform Layer-by-Layer Etching Technology for Line-Edge-Roughness Reduction and Surface-Flattening of 3-D Ge Channels,

21.2 Variation Improvement for Manufacturable FINFET Technology

Session 20: Characterization, Reliability and Yield – **Transistors Ageing, Variability and the Impact on Circuit Design**

20.5 Technology Scaling and Reliability: Challenges and Opportunities

20.8 Implications of Variability on Resilient Design (Invited), R. Aitken, ARM

The irony of reliability limits to Moore's Law

IEDM 2014

Will Reliability Limit Moore's Law?

Anthony S. Oates

TSMC Ltd., 168 Park Ave. 2, Hsinchu Science Park, Hsinchu, Taiwan 30075; aoates@tsmc.com

Abstract— Up to the present time reliability has not limited the rapid evolution of Si process technologies. However, the near future will bring a continual stream of innovations in transistor architecture and gate dielectric and interconnect materials. Maintaining historical high levels of reliability in this environment will be challenging. In this paper we discuss the reliability issues that have the potential to limit the future pace of technology progress.

minimal. PBTI in NMOS has been added as an issue to be monitored [3], while other transistor mechanisms have been readily optimized. Now to the FinFET transistor architecture is under indications pointing to successful implementation of the introduction of FinFETs to reliability appears to be relatively benign since the structure new mechanisms of degradation [4]. However, the confined fin geometry is exacerbated compared




Fig. 4: Use condition low-k failure time prediction for L=1000m interconnect line length at V=0.75Volt.

The answer is yes:

- Electromigration (already)
- Soft Errors, not just in memory (interleaving, ECC), but in logic, especially HPC
- RTN, BTI, etc.: the defects don't scale

Dennard warned us about electromigration

Table 2
Scaling Results for Interconnection Lines

Parameter	Scaling Factor
Line resistance, $R_L = \rho L/Wt$	κ
Normalized voltage drop IR_L/V	κ
Line response time $R_L C$	1
Line current density I/A	κ



...and Moore's Law is paying the price

Electromigration today (I6/I4) materially affects Moore's Law scaling entitlement




Maximum Current Limits \rightarrow fan out limits \rightarrow more buffering

We now spend significant transistors (and power) on EM

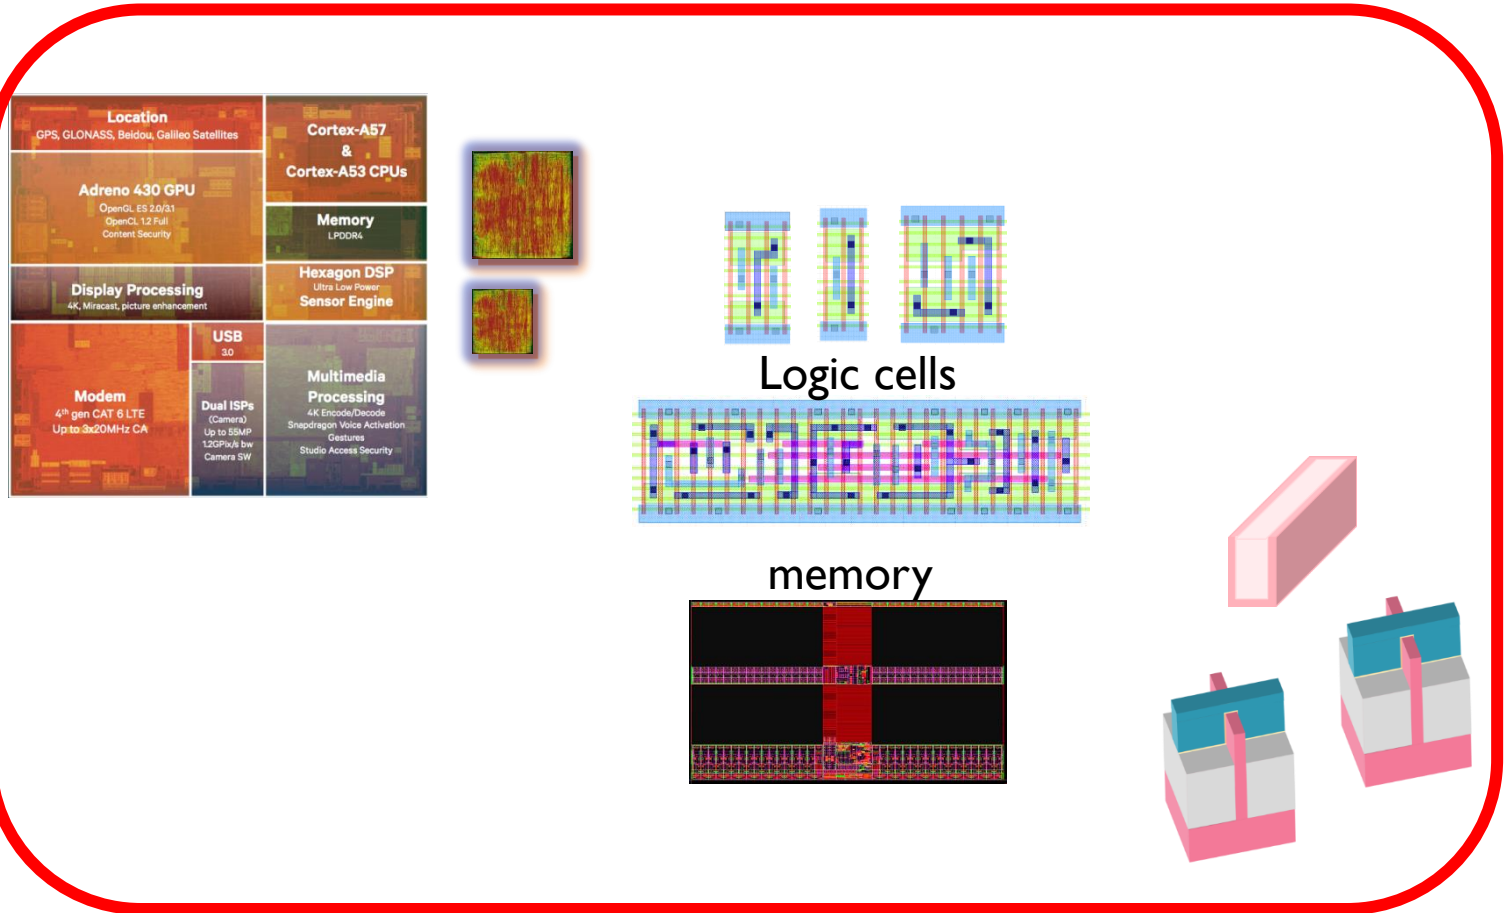
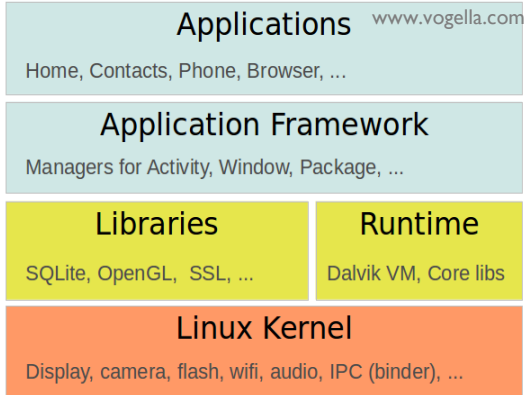
Device to circuit scaling: Summary

- Moore's Law scaling is slowing down
 - BEOL lithography cost is growing, with no near-term reset (EUV, DSA, etc.)
 - Cost per chip pressure, may include pressure on transistors/chip (chip area)
- Dennard scaling pressure is as bad or worse, and is often interchangeable
 - Within existing device electrostatics, invest in variation reduction
 - Contacts becoming a key limiter, and are a key focus of next generation transistors (via and wire R also increasing, adding R to C in general trend)
 - New devices that can offer reduced V without sacrificing I are ultimately needed
- DTCO to pull in the S-curve
 - As costs increase, and/or node timing slows, more radical changes become viable
 - Added device flexibility: heterogeneity, gate patterning flexibility, voltage swing, etc.
→ A Moore's Law Node may gained or lost
- Reliability and Yield: How much longer can we stay the course?

2015: Age 50

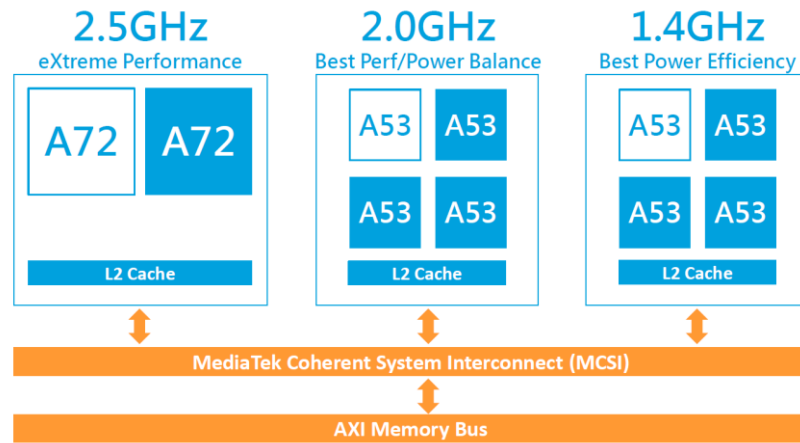


Architecture-Technology Interactions




Heterogeneous Multi Core in the Dark Silicon era

Deca/10-Core CPU Architecture



Source: anandtech

- World's 1st Integrated Cortex-M4**
 - Clock speed up to 364MHz
 - Dedicated SRAM size of 512KB
 - Isolated low power domain
 - Direct access to DRAM
- Open Platform – More Differentiation**
 - ARM-based, friendly to developers
 - Rich sample code and technical support
 - Complete tool chain



Power gain from Tri-cluster CPU architecture

	Dual-cluster power consumption	Tri-cluster power consumption	Improvement
FB launch	0.385W	0.318W	17%
FB Read	0.139W	0.084W	40%
FB Message	0.157W	0.101W	36%
FB scroll	0.217W	0.152W	30%
Beauty Plus	0.487W	0.378W	23%
Temple run launch	0.378W	0.316W	17%
Temple run play	0.303W	0.199W	34%
voice call	0.204W	0.121W	41%
Web Page loading	0.655W	0.627W	5%
Web Page Browsing	0.326W	0.273W	17%
Youtube HD	0.256W	0.156W	39%
Video Record	0.289W	0.197W	32%
Video Playback	0.113W	0.067W	41%
Homescreen idle	0.050W	0.026W	48%
Gmail	0.104W	0.061W	42%




Source: phonearena

13 | Confidential


Current trend:
Lower utilization for higher efficiency

Future trend:
More dedicated accelerators

Wire scaling modifies Moore's Law results



- Insert more transistors to maintain performance
- Add to power problems
- Problems are not just lateral, but more and more vertical



<http://www.zyvexlabs.com/EIPBNuG/2005MicroGraph.html>

15.5 Present Status and Future Prospects of Nano-Carbon Interconnect Technologies for LSIs

Evening Panel: Is there a potential for a revolution in on-chip interconnect?

3DIC: More than Moore (?)



“It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected” [1]

8.5 New Challenges and Opportunities for 3D Integrations

8.7 Enabling Low Power BEOL Compatible Monolithic 3D+ Nanoelectronics for IoTs ...

8.8 Advanced 3D Monolithic Hybrid CMOS with Sub-50 nm Gate Inverters ...

3DIC and Moore's Law




3DIC cost + 3 x Cost(●)

Cost(●)


(plus defect clustering advantage)

3D folding granularity


2D




Core / block folding



Gate-level folding




Monolithic 3D




Power/Performance:
~One Moore's Law Node

Chang, Acharya, Sinha, Cline, Yeric, Lim;
ISLPED 2015


3DIC early adopters



micron.com




Xilinx.com



sony.net


CMOS Image Sensor and 3DIC: Virtuous circles



Smart Phones



3ICs



Demand

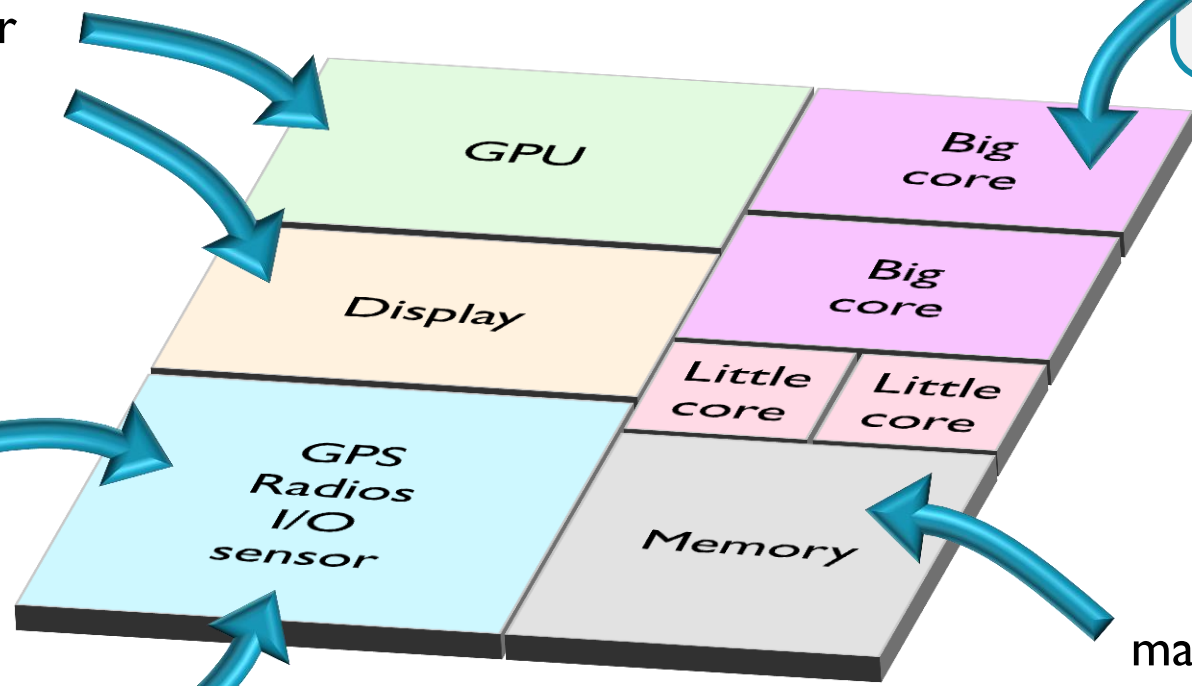


Heterogeneous SoC

Parallelizable functions:
Favors efficiency over
raw performance

High performance
CMOS

Analog devices not
optimized in CMOS,
RF may prefer FDSOI
Leading edge feature
density wasted?

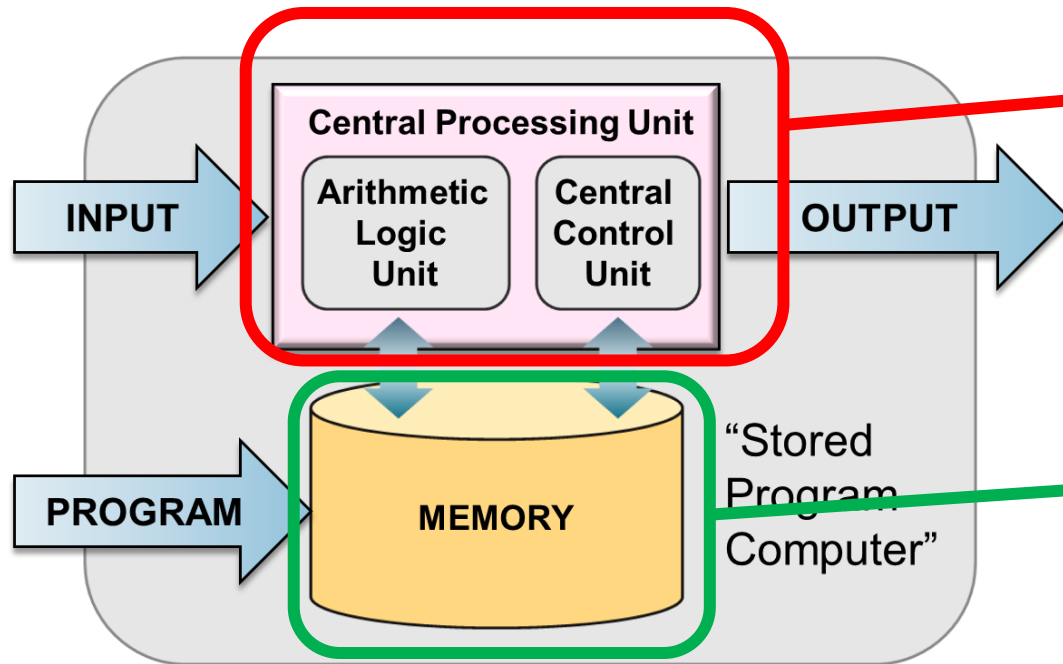


matching, leakage

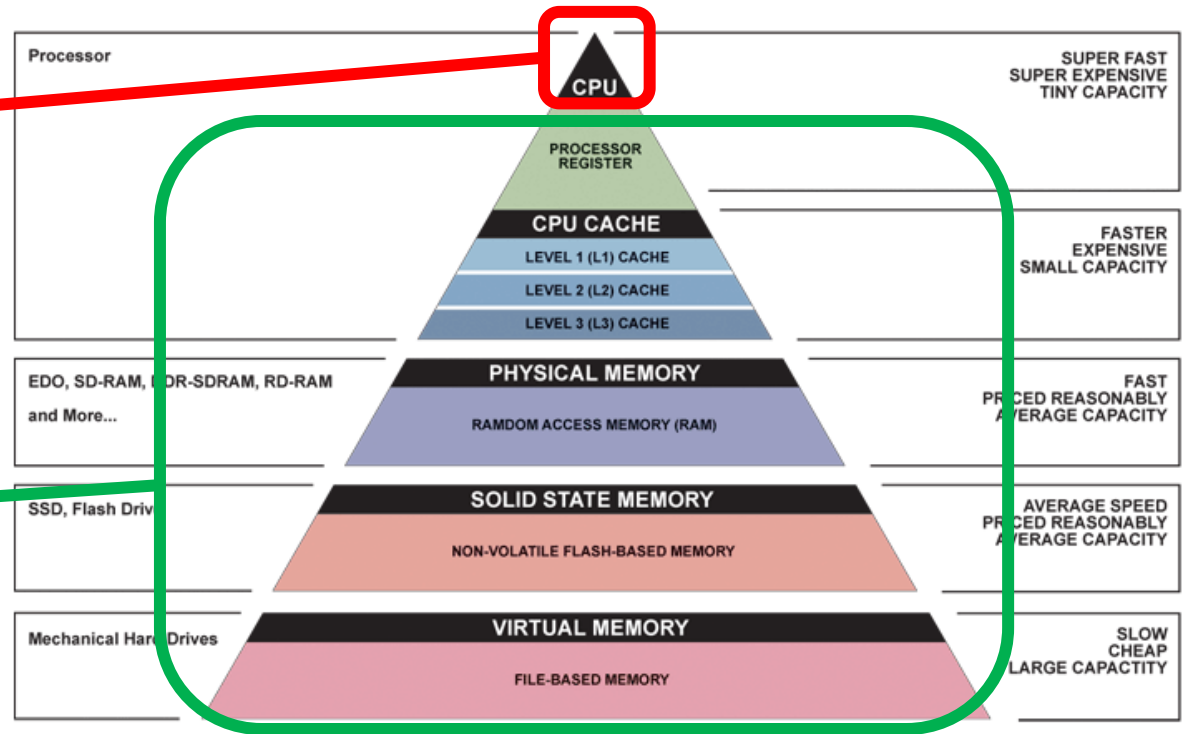
Always on sensor
Leakage / performance

Now.... think along product portfolios,
and not just one chip

Modern von Neumann computers



What software thinks memory looks like



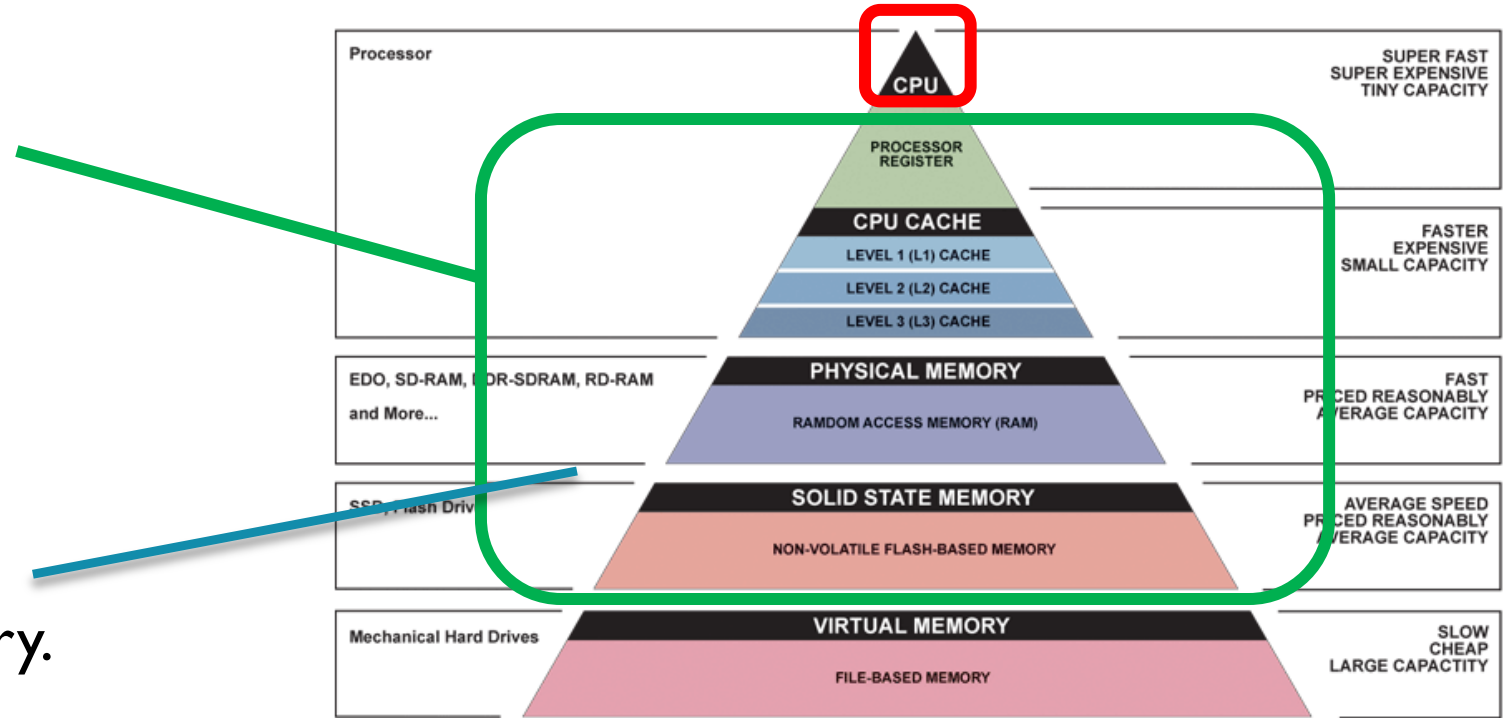
▲ Simplified Computer Memory Hierarchy
Illustration: Ryan J. Leng

<http://www.bit-tech.net>

Modern von Neumann computers

All semiconductor memories have all three problems

Application trends demand exponential commodity memory.



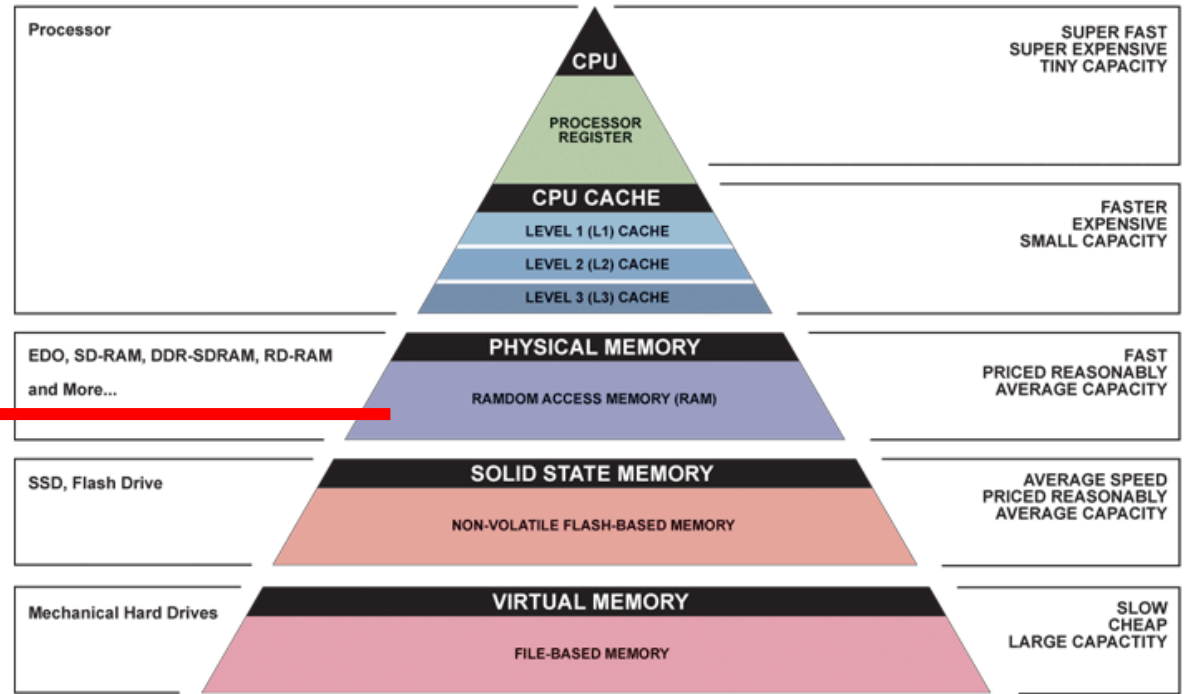

▲ Simplified Computer Memory Hierarchy
Illustration: Ryan J. Leng

<http://www.bit-tech.net>

DRAM scaling at the S curve

DRAM up to half of system power
Still consumes power when not doing anything

2-3 “known” nodes left




▲ Simplified Computer Memory Hierarchy
Illustration: Ryan J. Leng

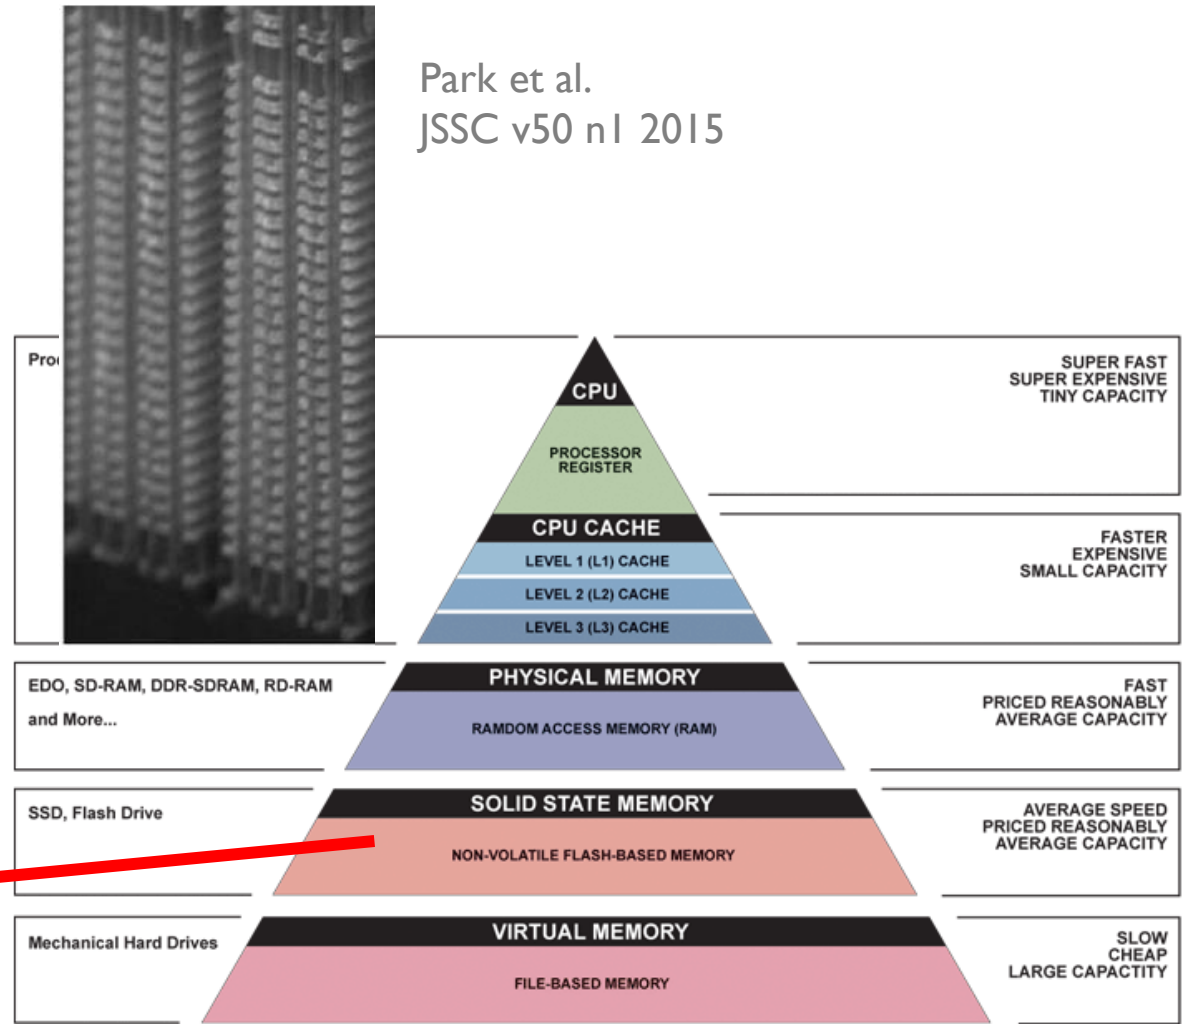
<http://www.bit-tech.net>

NAND: MLC and Vertical


...but for how much longer?



Jung H. Yoon, flashmemorysummit.com



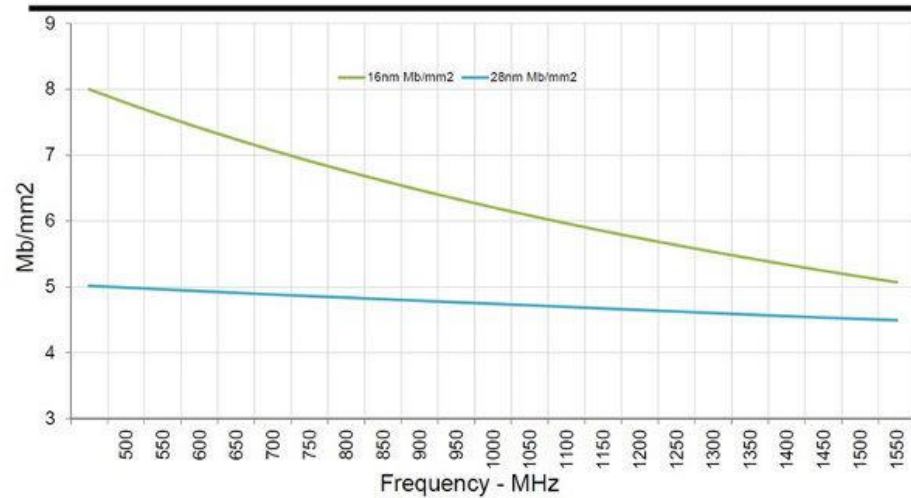
Park et al.
JSSC v50 n1 2015



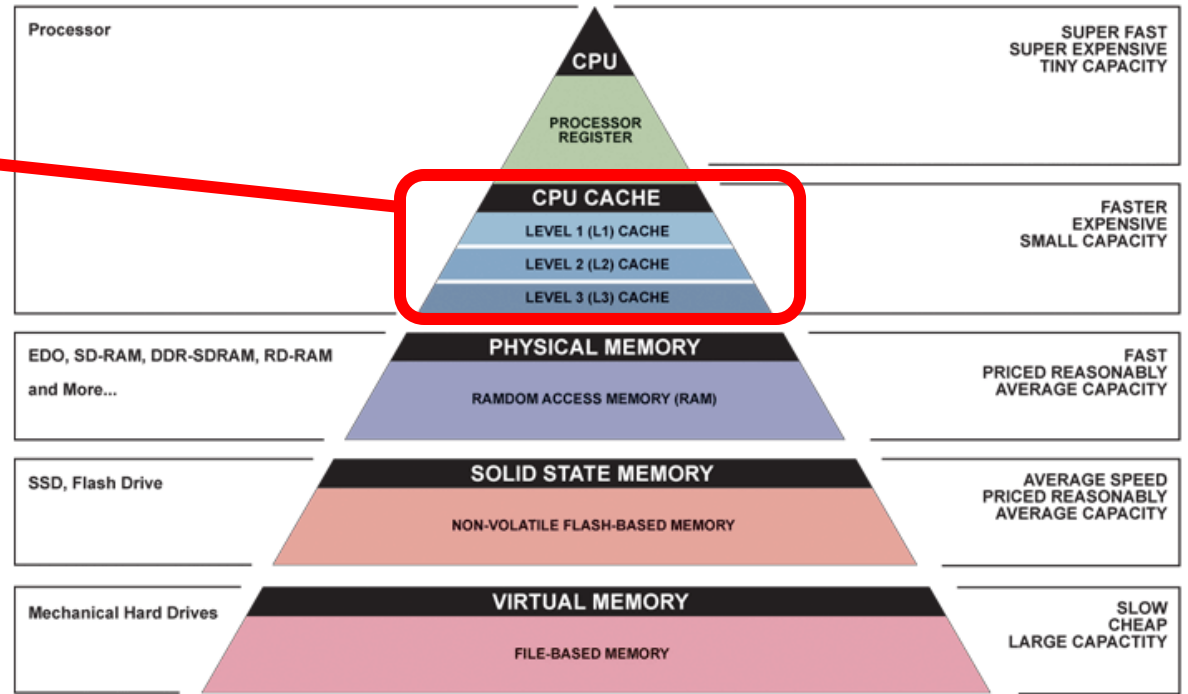
NAND Flash is a horrible NVM technology: Cost, Power, Speed, Endurance
Until you consider the alternatives

SRAM cache scaling

SRAM DENSITY - 16nm vs 28nm



- The trend is for smaller bank sizes and more overhead (<50% efficiency)
- Bigger transistors, more transistors
6T → 8T
- Often the voltage scaling bottleneck



▲ Simplified Computer Memory Hierarchy
Illustration: Ryan J. Leng

<http://www.bit-tech.net>

Wanted: Super Memory

	SRAM	DRAM	NAND	SuperMem
Area (F ²)	~120	4-6	<4 (eff)	<=6
Write Speed	<300ps	10ns	50ns+	<1ns
Read Speed	<300ps	10ns	10ns	<1ns
Leakage	High	Low	0	0
Active Power	“Low”	Low	High (write)	Low (1pJ/bit)
Nonvolatile	No	No	Yes	Yes
High Voltage	No	No	Yes	No
Logic process	Yes	No	No	Yes
Endurance	Infinite	Infinite	10 ⁵	>10 ¹⁵

- Super memory needs to be able to scale!
- Bringing denser, faster, and/or lower power memory to compute can greatly offset scaling problems
- A fast, high endurance NVM would enable new paradigms in persistent compute

Memory technology at IEDM 2015

3.3 A Floating Gate Based 3D NAND Technology with CMOS Under Array

3.6 Crystalline-as-Deposited ALD Phase Change Material Confined PCM Cell for High Density Storage Class Memory,

7.7 Distribution Projecting the Reliability for 40 nm ReRAM and beyond based on Stochastic Differential Equation,

10.6 Programming-Conditions Solutions Towards Suppression of Retention Tails of Scaled Oxide-Based RRAM,

10.1 Non Volatile Memory Evolution and Revolution

26.1 Fully Functional Perpendicular STT-MRAM Macro Embedded in 40 nm Logic for Energy-efficient IOT Applications


26.2 Systematic Optimization of 1 Gbit Perpendicular Magnetic Tunnel Junction Arrays for 28 nm Embedded STT- MRAM and Beyond,

26.4 Solving the Paradox of the Inconsistent Size Dependence of Thermal Stability at Device and Chip-level in Perpendicular STT-MRAM

26.7 A Novel Bi-stable 1-Transistor SRAM for High Density Embedded Applications,

The contenders

- Filamentary RRAM:
 - Stochastic variability
 - Endurance
 - Scalability (filaments don't scale)
- PCM:
 - Variability
 - Endurance
- MRAM:
 - Power/speed tradeoff (including read margin)
 - Disturb
 - Cost



Age 50: Children



Children of Moore's Law

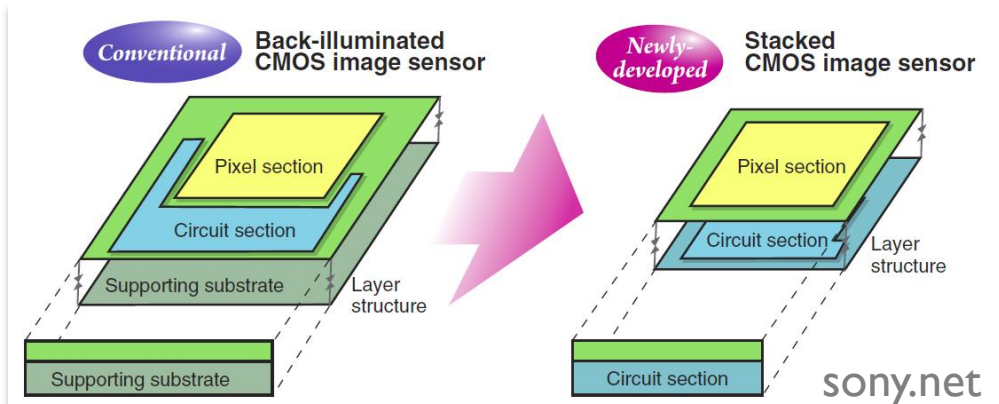
- Solar Panel
- Sensor
- Modem



\$5000
(or \$149/month over 5 years)



\$370



The city of Philadelphia reduced weekly trash collections from 17 to 3
The city of Barcelona estimates \$4B savings over 10 years

Child of Moore's Law: Photovoltaics

- \$0.5 on Alibaba:



Child of Moore's Law: MEMS sensors


- MEMs can further benefit from VLSI “ization” (ecosystem standards, etc.)





Chris Wasden at 2014 MEC, via semiwiki

1T SENSORS IN 10 YEARS						
Year	Unit Price	Units Sold	Industry Revenues	Developed Population	MEMS Rev/ Person	MEMS Unit/ Person
2005	30.000	46,666,667	5,000,000,000	4,000,000,000	1.25	0.01
2010	15.000	466,666,667	7,000,000,000	4,000,000,000	1.75	0.12
2015	1.800	8,333,333,333	15,000,000,000	4,000,000,000	3.75	2.08
2020	0.216	138,888,888,889	30,000,000,000	4,000,000,000	7.50	34.72
2025	0.026	1,388,888,888,889	60,000,000,000	4,000,000,000	15.00	347.22

2007 the average cost of an accelerometer sensor was \$3.
In 2014, the average was 54 cents



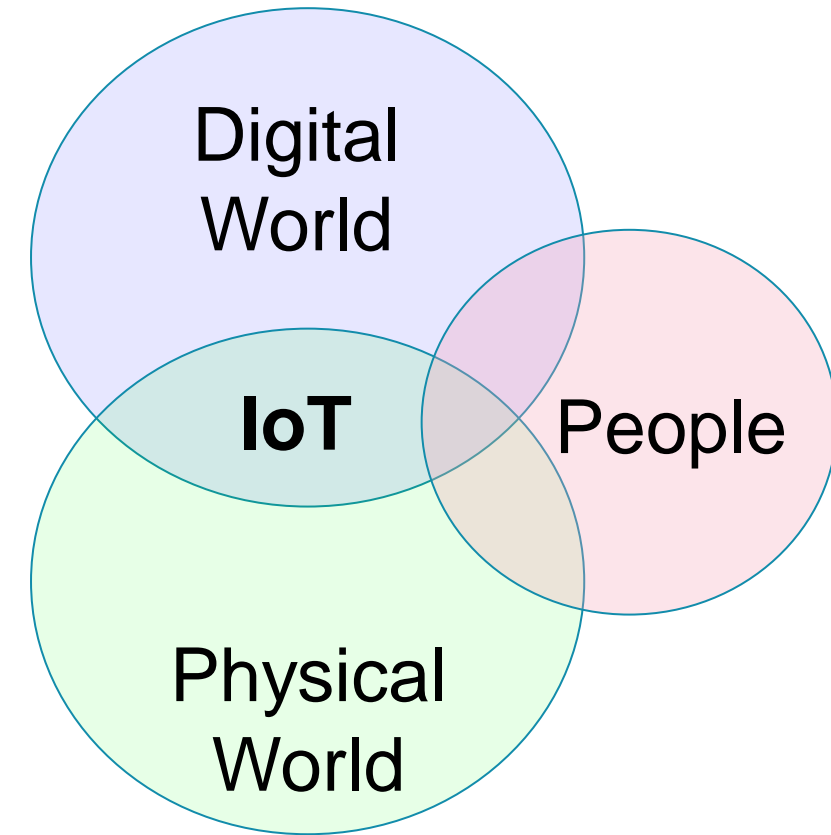
Future sensors: Also children of Moore's Law progress



Adamant technologies, e.g.

Moore's Law progress created the IoT

1. Vanishing edge node compute
2. Ubiquity of smart phones
 - A thing-to-person hub
3. Ubiquity of Internet, and plummeting cost to connect to it
4. Plummeting cost of (and richer set of) sensing
5. Reduced cost of energy harvesting




reference 31


Example IoT: Michigan Micro Mote

- 3DIC for form factor
- Energy Harvesting →
- Near Threshold operation
 - Variation 10x more important

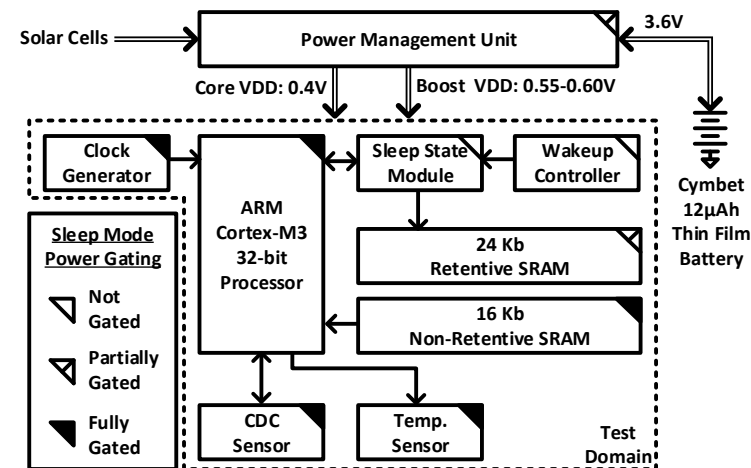

Battery Processor Solar Cells



G. Chen et al., ISSCC, 2010.





<http://www.eecs.umich.edu/eecs/about/articles/2015/Worlds-Smallest-Computer-Michigan-Micro-Mote.html>



Energy Efficiency: Things you can do with 100pJ

- Run a Cortex[®]-M0 for 10 cycles
- Write one bit of flash
- Write ~300 bits of DRAM or SRAM
- Send ~5 bits across LPDDR4
- Transmit 2 bits of UWB data
- Transmit 0.02 bits over Bluetooth LE
- Drive an electric car 100fm (@1MJ/km) ~0.05% of the distance across Si atom

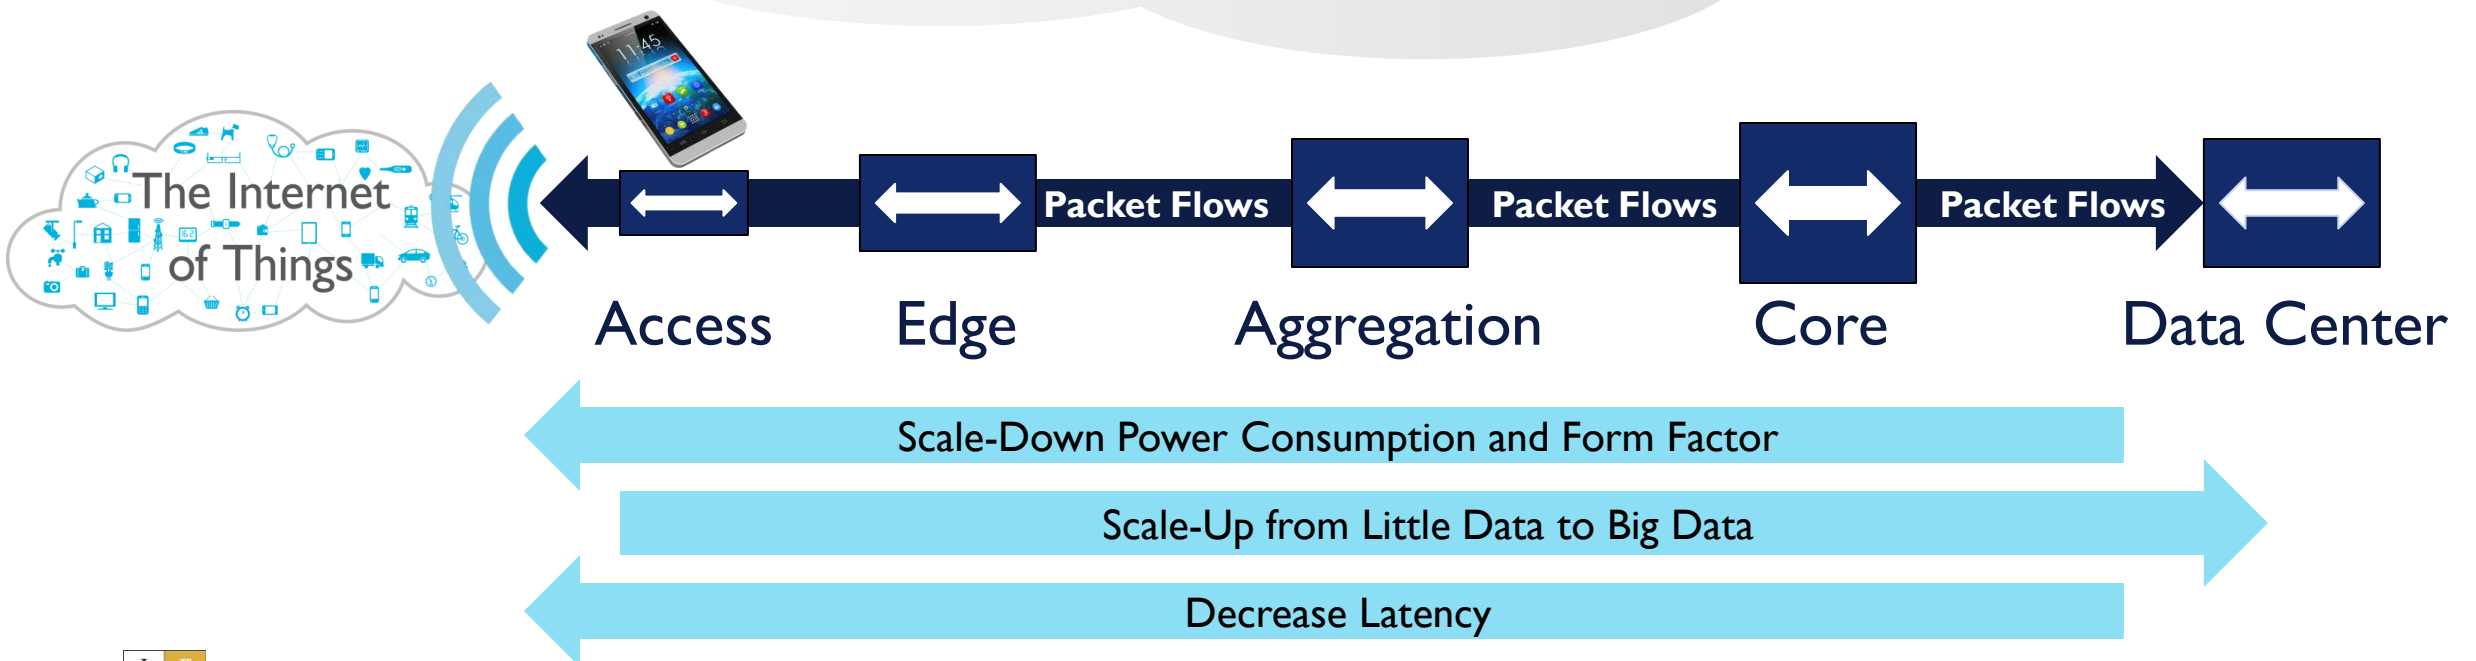
The IoT is an NVM problem



Energy costs to transmit, compute, and store data will define the shape of the IoT
VSLI Technology advancements will re-write the boundary conditions

Intelligent Flexible Cloud

- Applications run where the data is, independent of the network node
- Heterogeneous compute is distributed into the network
- **ALL OF THIS COMMUNICATION MUST BE SECURE**



Internet of Things at IEDM 2015

13.1 Ultra Low Power Sensor Platforms for Personal Health and Personal Environmental Monitoring

19.4 Free Form CMOS Electronics: Physically Flexible and Stretchable

13.5 Precision Mass Measurements in Solution Reveal Properties of Single Cells and Bioparticles

13.6 Fabrication and Analysis of SiN Nanopores for Direct DNA Sequencing

18.7 Output Enhancement of Triboelectric Energy Harvester by Micro-Porous Triboelectric Layer

19.1 Flexible Electronics Manufacturing: Flexible Digital x-ray to Flexible Hybrid Electronics

19.5 Large Area Sensing Surfaces: Flexible Organic Printed Interfacing Circuits and Sensors

19.8 Flexible 2D FETs using hBN Dielectrics

Session 25: Circuit Device Interaction – **More than Moore – Value Added Technologies**

25.4 Low-Cost and TSV-free Monolithic 3D-IC with Heterogeneous Integration of Logic, Memory and Sensor Analogy Circuitry for Internet of Things,

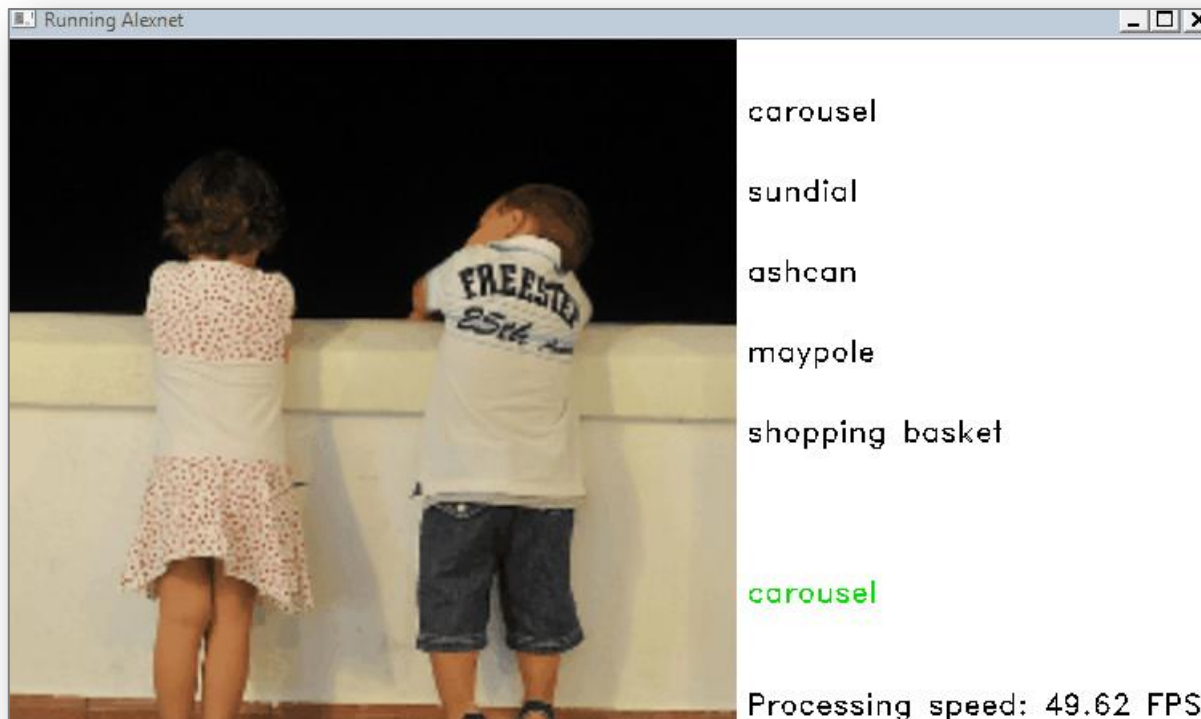
25.5 New Devices for Internet of Things: A Circuit Level Perspective

25.7 An Integrated Silicon Photonics Technology for O-band Datacom,

Child of Moore's Law: Image Recognition

- Everyone carries a camera → storage is “free” → Internet of Images
- Hardware scaling allows semi-affordable Machine Learning (ML):

(>900M operations/image)



Machine Learning:


Use transistors
instead of expert
developers

“this is a face”

“this is a QR code”

Putting it all together: Achieving the 100x

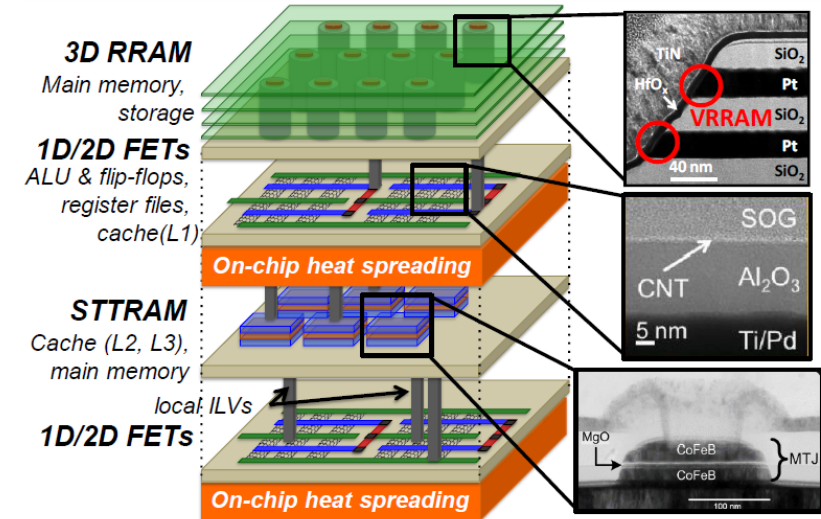
- Optimized layer processing
- New bandwidth paradigms
- New form factors



Shen et al., IEDM 2014

Example: Fine-Grained Interleaving of Memory with Logic for Machine Learning

100X Energy Delay Product benefit
vs. multi-core 2D ICs with DDR3 for Google PageRank



- New application drives new devices and heterogeneous integration

Stanford | SystemX Alliance

Neuromorphic computing at IEDM 2015

4.5 Memristive Based Device Arrays Combined with Spike Based Coding Can Enable Efficient Implementations of Embedded Neuromorphic Circuits

4.1 Brain-inspired Computing with Emerging Memories

4.4 Large-Scale Neural Networks Implemented with Non-volatile Memory as The Synaptic Weight Element: Comparative Performance Analysis

4.6 A Mixed-Signal Universal Neuromorphic Computing System

4.9 DARPA Neurocomputing


4.7 Oxide Based Nanoscale Analog Synapse Device For Neural Signal Recognition System

17.2 Investigation of the Potentialities of Vertical Resistive RAM (VRRAM) for Neuromorphic Applications,


17.7 Optimized Learning Scheme for Grayscale Image Recognition in a RRAM Based Analog Neuromorphic System

1.2 Quantum Computing in Si,




Past transistor pipeline



Future Transistor pipeline




What can you do with a future transistor?

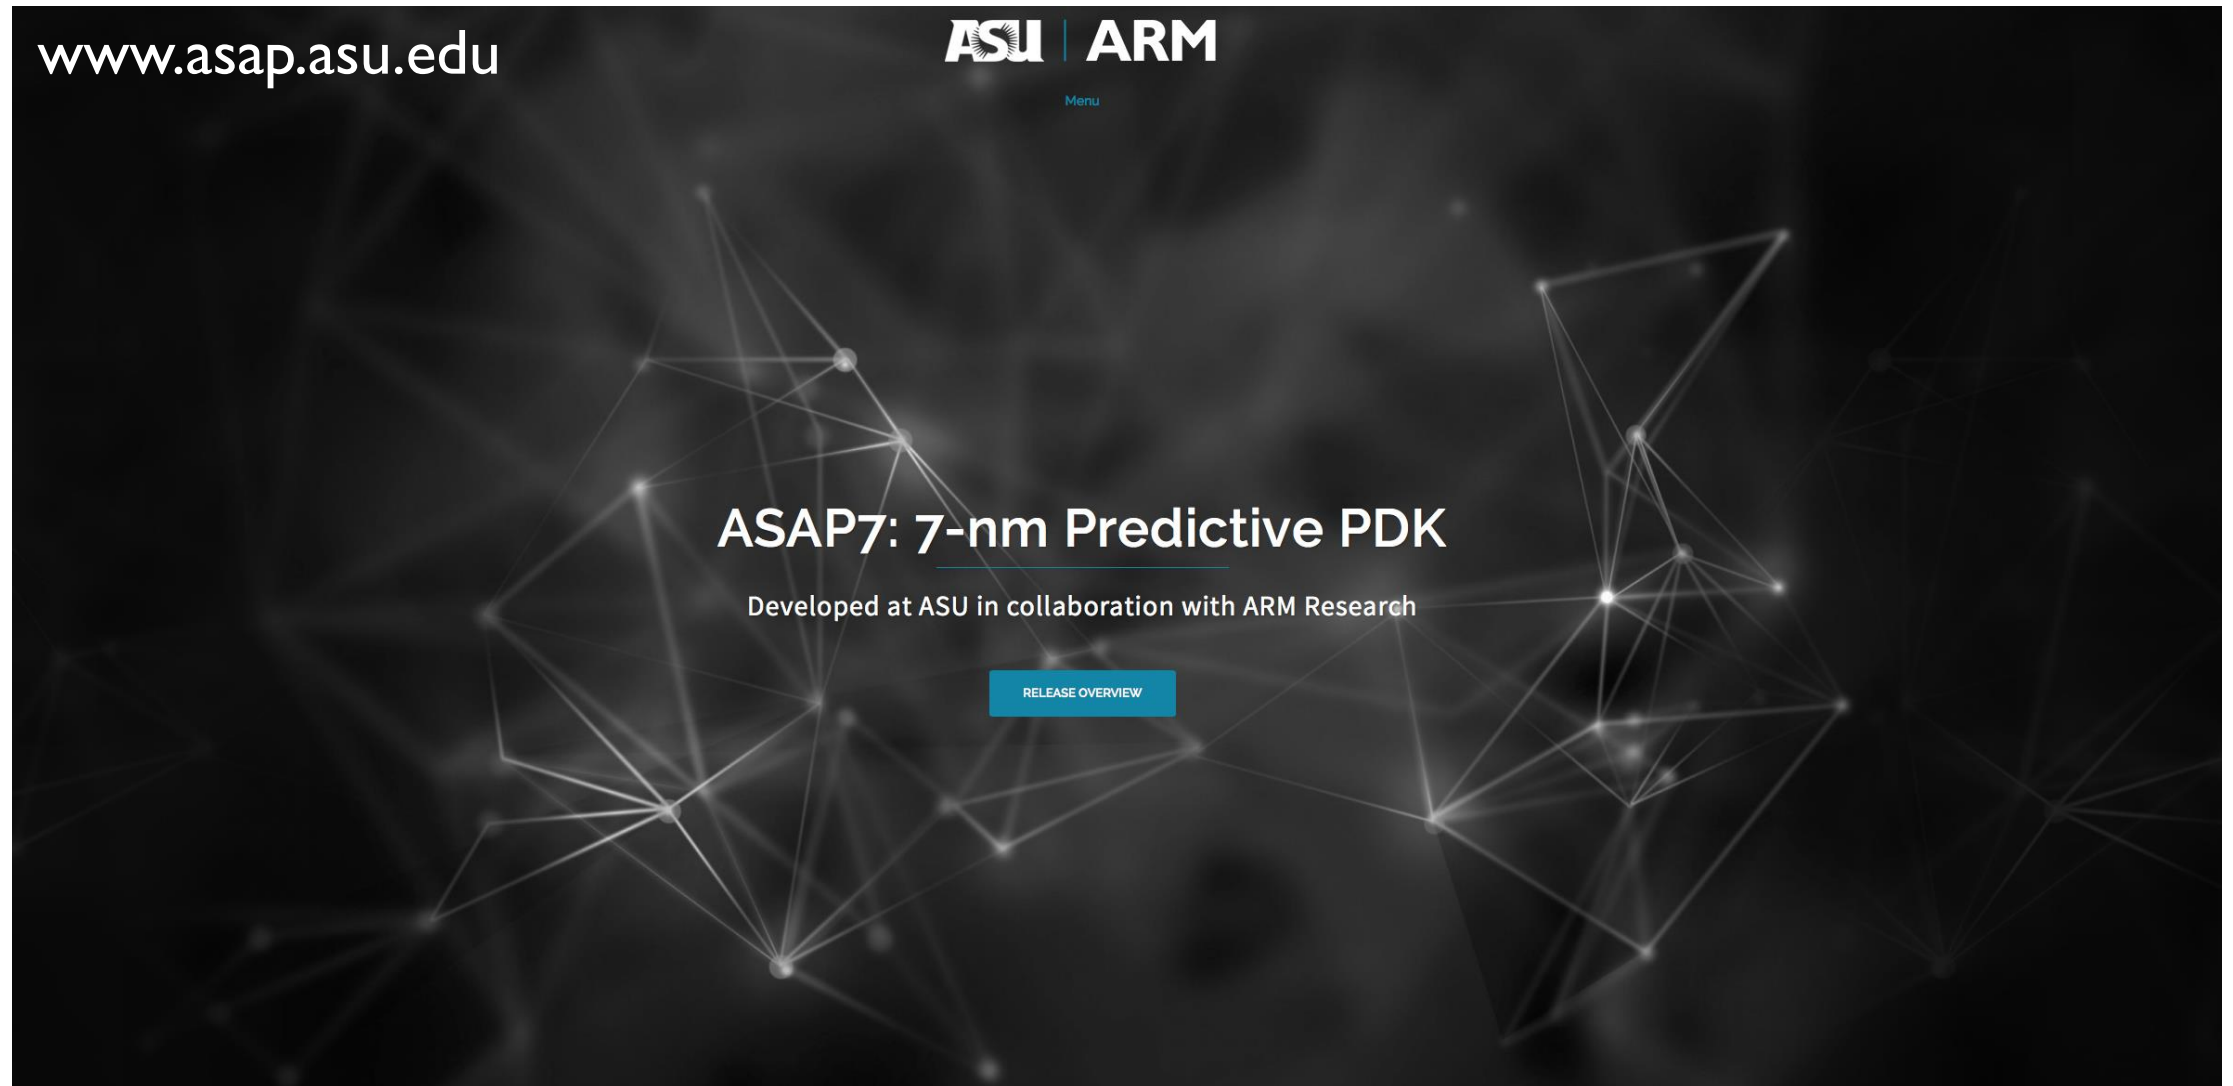


Needed:
technology
Benchmarking

(SRC ref #20)



ASU ASAP 7nm Predictive PDK



PDK enabling more accurate university research

www.asap.asu.edu

The ASAP 7nm Predictive PDK was developed at ASU in collaboration with ARM. The ASAP 7nm PDK contains SPICE-compatible FinFET device models (BSIM-CMG), Extraction Deck for the 7nm technology node. For more details regarding the technical information this process design kit is provided as an academic resource.

What's in a PDK?

1. Transistor models:
BSIM-CMG FinFET transistor models:
3 VTs and corners
2. Technology files for schematic capture/layout
3. DRC
4. LVS
5. Extraction

Predictive Technology Model

- Introduction
- Latest Models**
- Nano-CMOS
- Post-Silicon
- Interconnect
- Reliability

LATEST MODELS

Typical SPICE model files for each future generation are available here.

Attention: By using a **PTM** file, you agree to acknowledge both the URL of **PTM**: <http://www.asap.asu.edu> in all documents and publications involving its usage.

New!
June 01, 2012:
PTM releases a new set of models for multi-gate transistors (**PTM-MG**), for both HP and LSTP **BSIM-CMG**, a dedicated model for multi-gate devices.
Acknowledgement: PTM-MG is developed in collaboration with ARM.

Please start from [models](#) and [param.inc](#).

- 7nm PTM-MG [HP NMOS](#), [HP PMOS](#), [LSTP NMOS](#), [LSTP PMOS](#)
- 10nm PTM-MG [HP NMOS](#), [HP PMOS](#), [LSTP NMOS](#), [LSTP PMOS](#)
- 14nm PTM-MG [HP NMOS](#), [HP PMOS](#), [LSTP NMOS](#), [LSTP PMOS](#)
- 16nm PTM-MG [HP NMOS](#), [HP PMOS](#), [LSTP NMOS](#), [LSTP PMOS](#)
- 20nm PTM-MG [HP NMOS](#), [HP PMOS](#), [LSTP NMOS](#), [LSTP PMOS](#)

The entire package is also available here: [PTM-MG](#)

Saurabh Sinha, Lucian Shifren, Brian Cline, Greg Yeric.

Other students of the EEE598 Special Topics course:

Alan Sam, Nalim Gupta, Rohit Musalay, Akash Thakare, Srividhya Jambunathan, Ramana Rao Pandeshwar, Jayesh Sohanlal, Chandrakanth Puttaswamygowda, Adesh Namekumar, Varun Kaushik, Sanyogita Singh.

Cells:

Sai Chaitanya Reddy, Punit Shah, Anant Mithal.

Acknowledgements

Mentor Graphics: Tarek Ramadan (for excellent DRC and LVS training at ASU).

PDK enabling more accurate university research

www.asap.asu.edu

Release Overview

The ASAP 7nm Predictive PDK was developed at ASU in collaboration with ARM Research. The PDK is available for download from www.asap.asu.edu. The ASAP 7nm PDK contains SPICE-compatible FinFET device models (BSIM-CMG), Technology files for Cadence Virtuoso, Design Compiler, and the resulting design files. For more details regarding the technical specifications of the PDK, please refer the PDK Release Overview and the resulting design files. The PDK is designed to address the need for Mentor Graphics.

Contributors

Lawrence T. Clark, Vinayak
Sinha, Lucian Shifren
Students of the EEE598
Nam, Nalim Gupta, Rohit
Kandeshwar, Jayesh Soha
Kaushik, Sanyogita Sing

Sai Chaitanya Reddy, Punit Sh

Comes with the kit at download.

Related Publications








TBD

Acknowledgements

Mentor Graphics: Tarek Ramadan (for excellent DRC and LVS training at ASU).

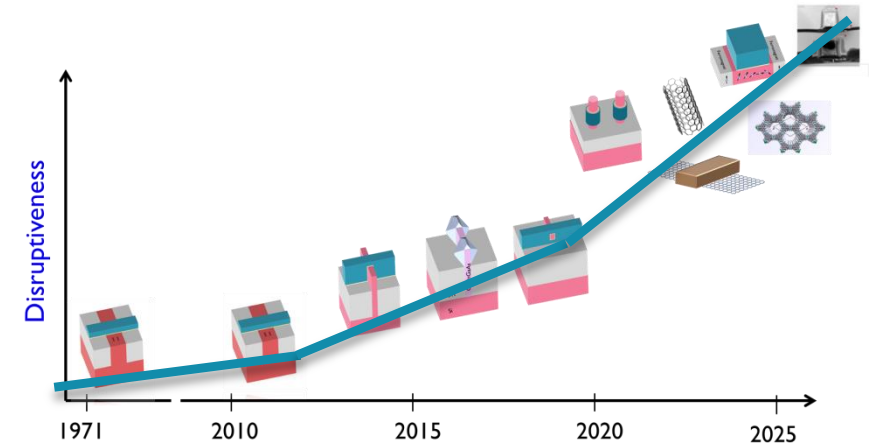
Summary

- Moore's Law and IEDM: "not free" on lorry to free in magazine
 - BUT, cost scaling is slowing down. There are no near term magic bullets.
 - Task/opportunity for technology investment: Pull in the curves
 - Fab / Tools / Process Integration / Device / Circuit (DTCO)
 - Dennard scaling is as challenged, and as important (FETs and wires)
- SoC: Continued informed innovation. Leverage 3DIC and novel memory
 - Energy efficiency vs. utilization → heterogeneity
 - Complex future technology choices need transistor-to-system benchmarking
- Breadth of future systems strains simple Moore's Law answers
 - Slowing fundamental scaling increases opportunity for radical change
 - Exciting new systems leverage the children of Moore's Law: Sensors, MEMS, ML
 - Systems must understand (and guide) underlying technology options



Moore's Law at 50, IEDM, and you: Summary

- You live in interesting times
 - New device physics in logic and memory
 - New levels of design efficiency
 - Tighter pipeline to new technologies



- We
- ~~You~~ are in the right place
 - Circuits and Systems help define device choices

Neuromorphic computing at IEDM 2015

- 4.5 Memristive Based Device Arrays Combined with Spike Based Coding Can Enable
- 4.1 Brain-Inspired Computing with Emerging Memories
- 4.4 Large-Scale Neural Networks Implemented with Non-volatile Memory as The Synapse
- 4.6 A Mixed-Signal Universal Neuromorphic Computing System
- 4.9 DARPA Neurocomputing
- 4.7 Oxide Based Nanoscale Analog Synapse Device For Neural Signal Recognition Systems
- 1.2 Quantum Computing in Si
- 17.2 Investigation of the Potentialities of Vertical Resistive RAM (VRRAM) for Neuromorphic Computing
- 17.7 Optimized Learning Scheme for Grayscale Image Recognition in a BRRAM Based Architecture

Internet of Things at IEDM 2015

- 13.1 Ultra Low Power Sensor Platforms for Personal Health and Personal Environmental Monitoring
- 19.4 Free Form CMOS Electronics: Physically Flexible and Stretchable
- 13.5 Precision Mass Measurements in Solution Reveal Properties of Single Cells and Bioparticles
- 13.6 Fabrication and Analysis of SiV Nanopores for Direct DNA Sequencing
- 18.7 Output Enhancement of Triboelectric Energy Harvester by Micro-Porous Triboelectric Layer
- 19.1 Flexible Electronics Manufacturing: Flexible Digital x-ray to Flexible Hybrid Electronics
- 19.5 Large Area Sensing Surfaces: Flexible Organic Printed Interfacing Circuits and Sensors
- 10.8 Flexible 2D FETs using hBN Dielectric

Session 25: Circuit Device Interaction – More than Moore – Value Added Technologies

- 25.4 Low-Cost and TSV-free Monolithic 3D-IC with Heterogeneous Integration of Logic, Memory and Sensor Analogy Circuitry for Internet of Things
- 25.5 New Devices for Internet of Things: A Circuit Level Perspective
- 25.7 An Integrated Silicon Photonics Technology for D-band Datacom

Together, at the system level, I expect Moore's Law level progress well past my own retirement