



Moore's Law: Where are we and which way are we going?

Greg Yeric
greg.yeric@arm.com

Fellow

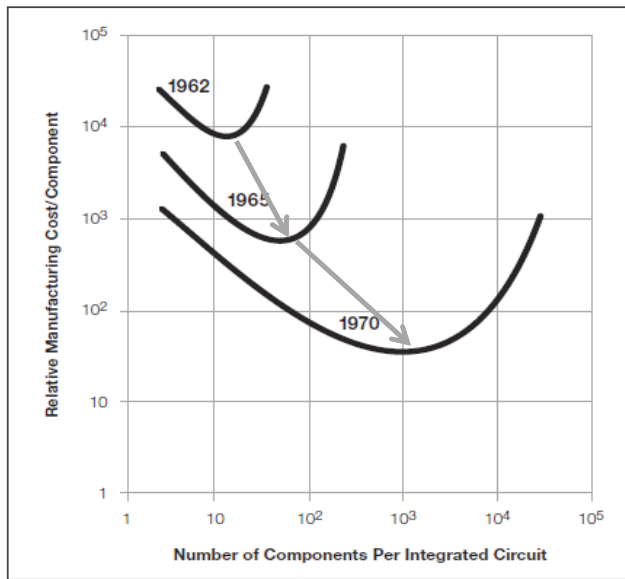
ARM Research

Moore's Law at ARM


ARM Research Future Silicon Technology Group

- Predictive technology modeling
- Disruptive technologies:
Assessment, Incubation
- Future hardware resilience


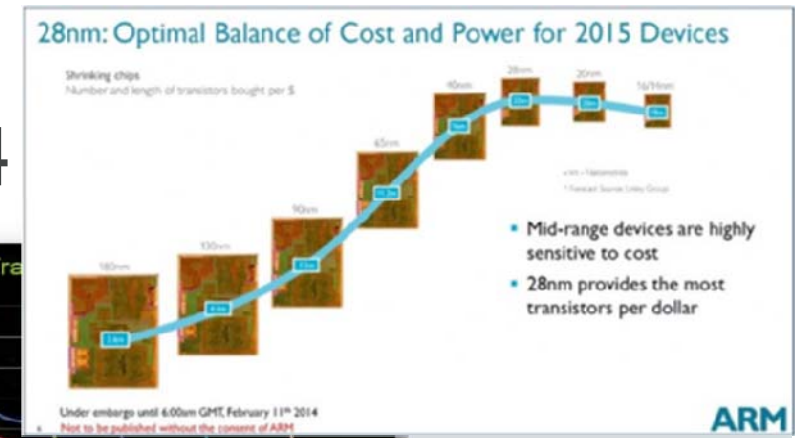
Has cost per transistor scaling reversed?



1965



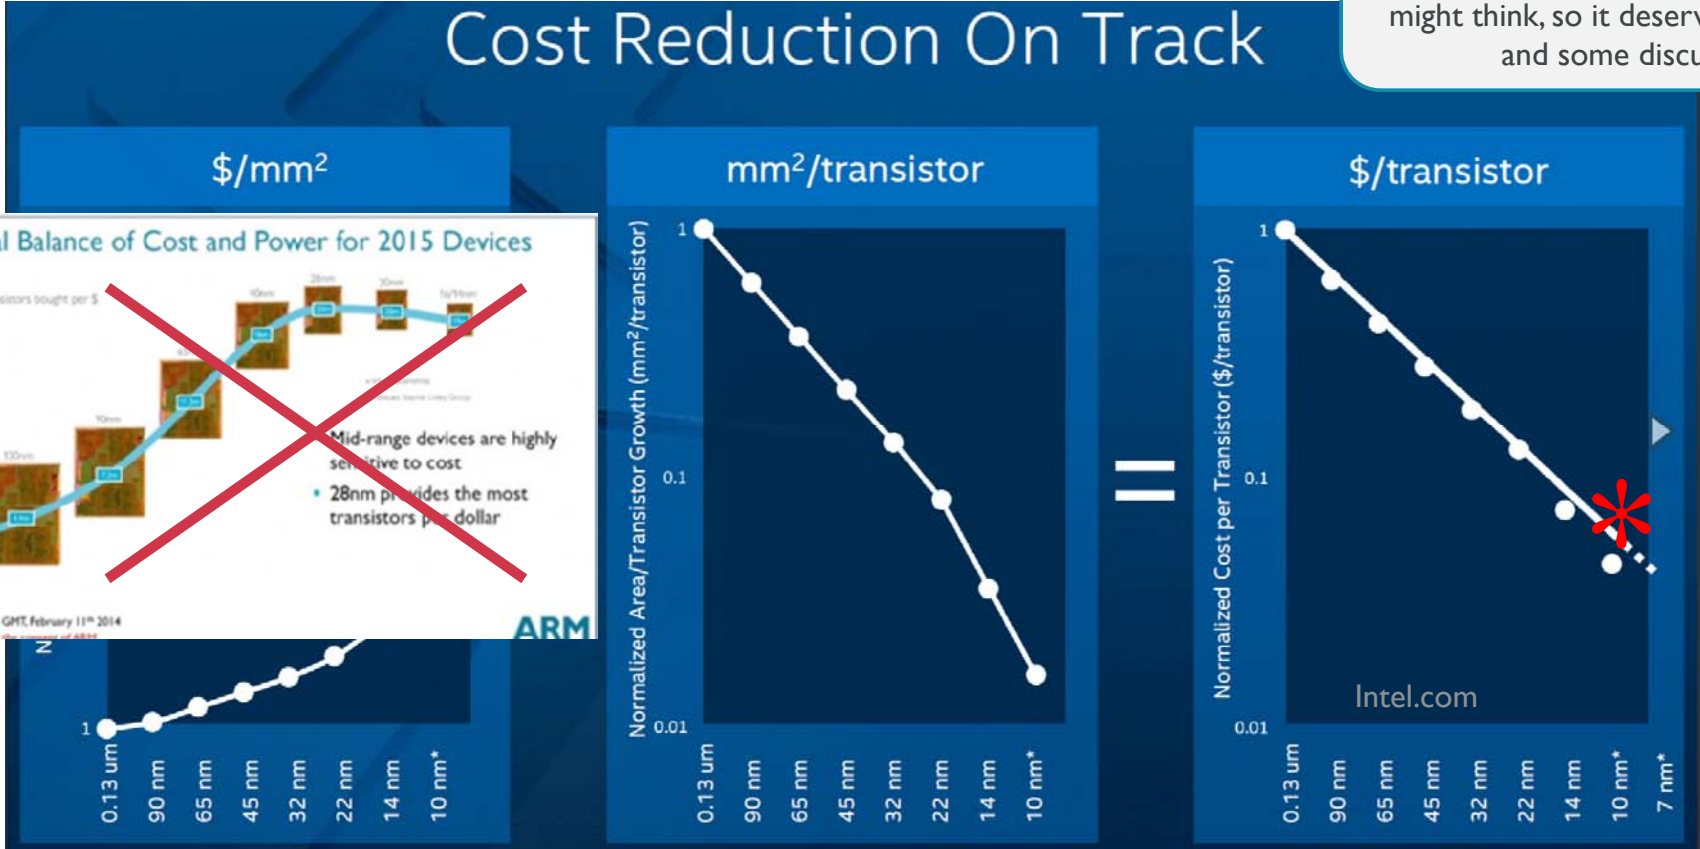
2014



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

Or is everything just fine?

Answer:
Cost/transistor still going down post-28nm, but not as straightforward as you might think, so it deserves an asterisk and some discussion



Standard Cell scaling...proper shrink


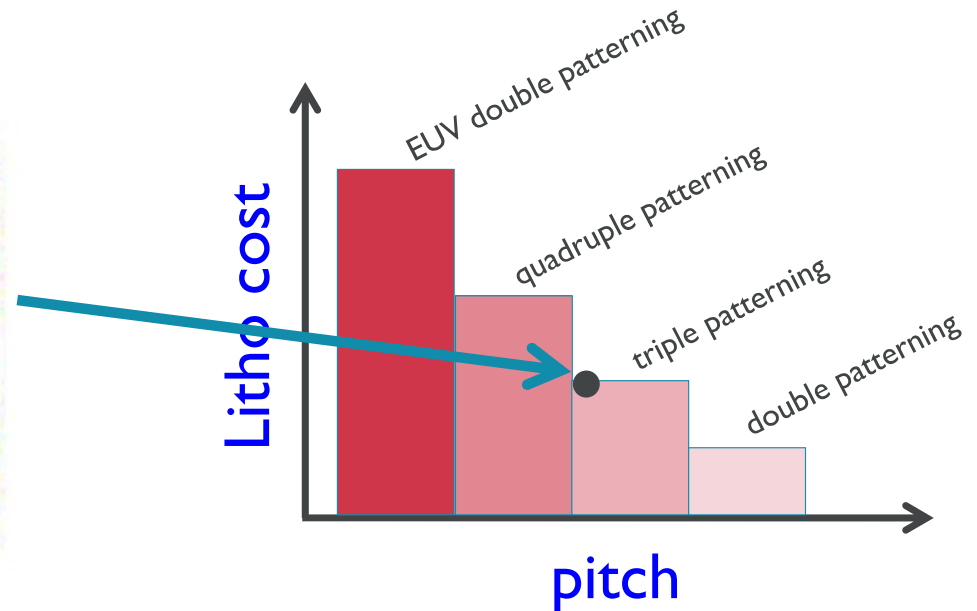
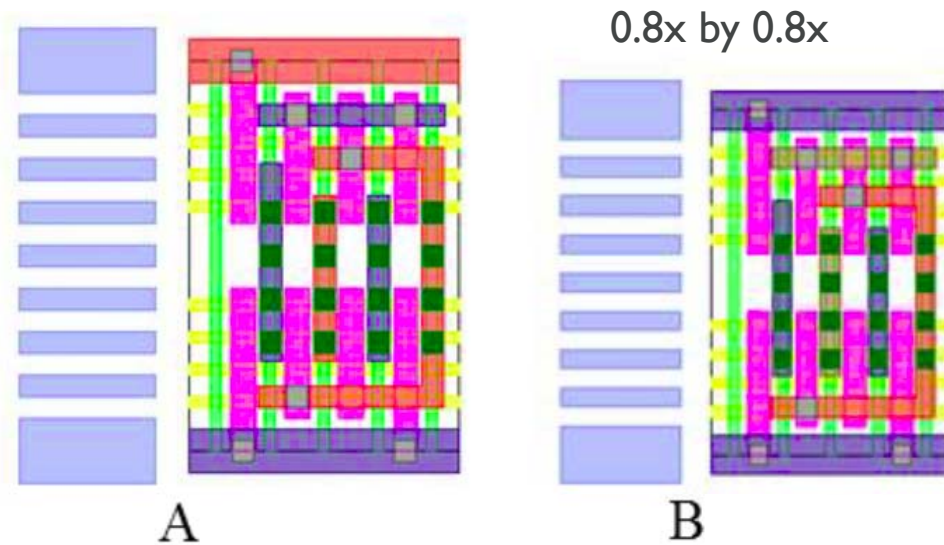


Fig. 2 A: 9T N14 AOI showing pin access points in green, B: 9T AOI in N10 showing triple-patterned M1 in blue, red, brown, C: 7.5T N10 AOI showing that adequate pin access can be preserved.

L. Liebmann et al, VLSI, p. 112 (2016).

A: node "N"
B: node "N+1"
Proper scaling 0.7x linear,
But if that hits a cliff in litho costs,
might want to try something smarter

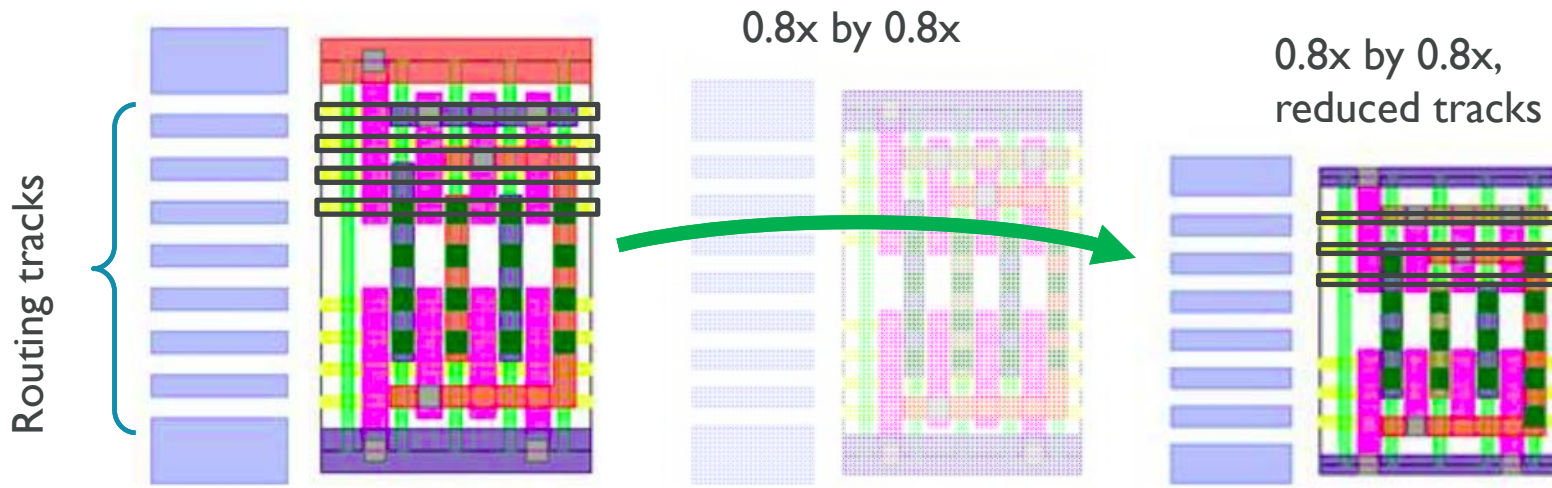
Standard Cell scaling: Affordable shrink...



Scale dimensions just to the next big cliff


...and we'll make it up somewhere else

Standard Cell scaling... with track reduction a.k.a. fin depopulation




Can get to desired transistor density at less than 0.7x linear shrink, if we can take tracks (height) out of the standard cells. This usually comes at the cost of few fins per cell (black)

Standard Cell scaling... 4, 3, 2 fins per transistor



This is a general trend in the industry. To hit performance targets, we need the FinFETs to deliver more current, which means the fins must grow taller. Not trivial.



OK, so now what?

Fin depopulation strategy quickly comes to an abrupt end. Must look for other tactics.

4, 3, 2 ... 1 fin?




Fig. 5 A: grating-and-cut compliant 6T N7 cell, B: N5 cell maintains unidirectional layout style even with the arrival of EUV lithography, C: N5 with over-scaled M0 as proposed by IMEC.

L. Liebmann et al, VLSI, p.112 (2016).

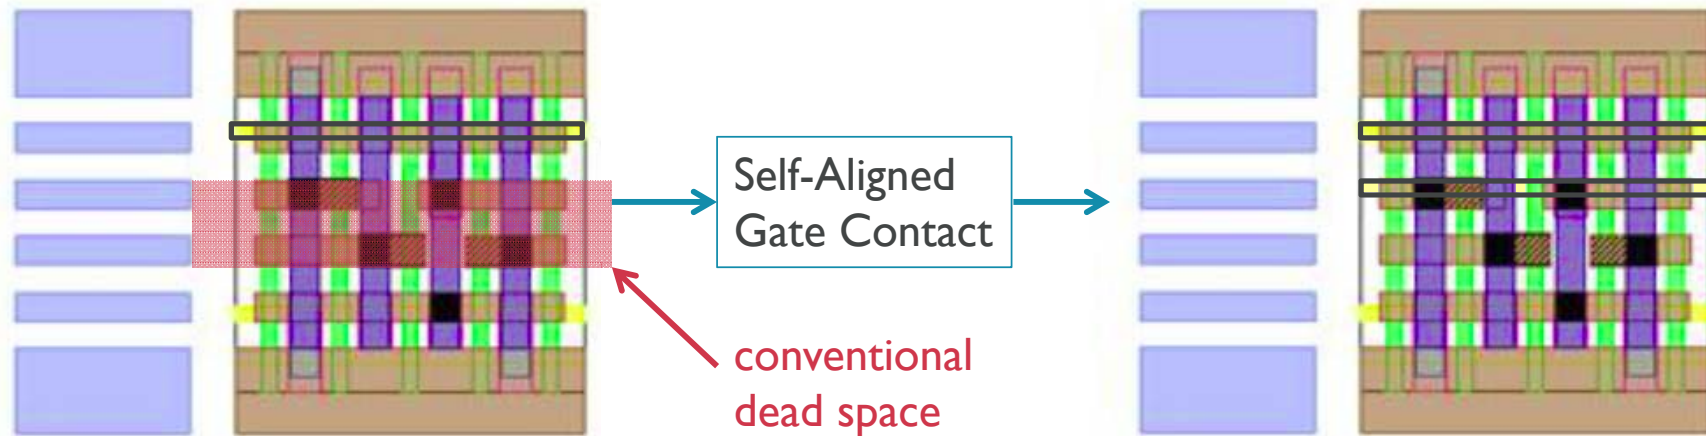
N7 to N5 Scaling

N5 is emerging to be a bitter-sweet node: the delay in the much anticipated transition to EUV forces layouts to remain unidirectional, Fig. 5B. Additionally, the critical pitches bump up against hard ‘non lithography’ limits for poly and fin pitch scaling as was shown in Table 1. Compensating for modest pitch scaling through further cell height reduction requires **substantial and risky process innovation such as the ability to contact poly-gates over active fins**. While these solutions are being actively explored, alternatives such as over-scaling only the lowest level of metal, Fig. 5C, as proposed by researchers at IMEC [2], could be called upon to maintain an overall 0.56x area scale factor for one more finFET node.

As traditional pitch-based Moore’s Law advancement wanes, process technologists are pursue “scaling boosters” to make up the slack. Analogous to design accelerators.

Example “scaling booster”: Self-Aligned gate contact

One potential density improvement : Enabling gate contacts (in red highlighted area) to extend over active fins without contacting the S/D



The majority of wafer cost increase is in the patterning of conventional wires
Previously **laughable** ideas for transistor integration might become rational options
Other scaling booster options: single diffusion breaks, fully self-aligned vias, super vias
→ Requires design-technology co-optimization (DTCO). Routability, e.g.

A word of caution on cost prediction math

7nm steppers will be 50% faster than 28nm steppers


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

*drop England onto the globe,
align it to a precision one brick,
5 times per second

are your circuits worthy?



Continuous improvement from all suppliers:
resist, light source, mask, etch....



Cost per transistor

- Lithographic pitch scaling as easy path for Moore's Law has passed*
 - Heroic efforts throughout the ecosystem are mitigating the rate of slowing
 - *5nm with a good EUV would be a welcome exception. (3nm may require EUV double patterning)
- “scaling boosters” required to make up the slack (ex. fin depopulation)
 - Future scaling boosters **could** get 5nm to desired cost, **might** get us to 3nm**
 - Self-aligned gate contact, fully self-aligned via, super via, single diffusion break
 - **Without breaking physical design. Design-breaking boosters exist, and may be required.
 - We move from long-lived lithographic scaling paradigm to two-node and even one-time boosts (quicker, bigger steps)
 - Managing this complexity will be key to getting to 3nm on time and on budget
 - Requires thorough benchmarking (DTCO) to down-select and find the best path

Moore's Law and Dennard Scaling

From Moore's 1965 paper

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



Table 1

Scaling Results for

Device or Circuit		
Device dimension		
Doping concentration		
Voltage V		
Current I		
Capacitance ϵA		
Delay time/circuit	VC/I	$1/\kappa$
Power dissipation/circuit	VI	$1/\kappa^2$
Power density	VI/A	1

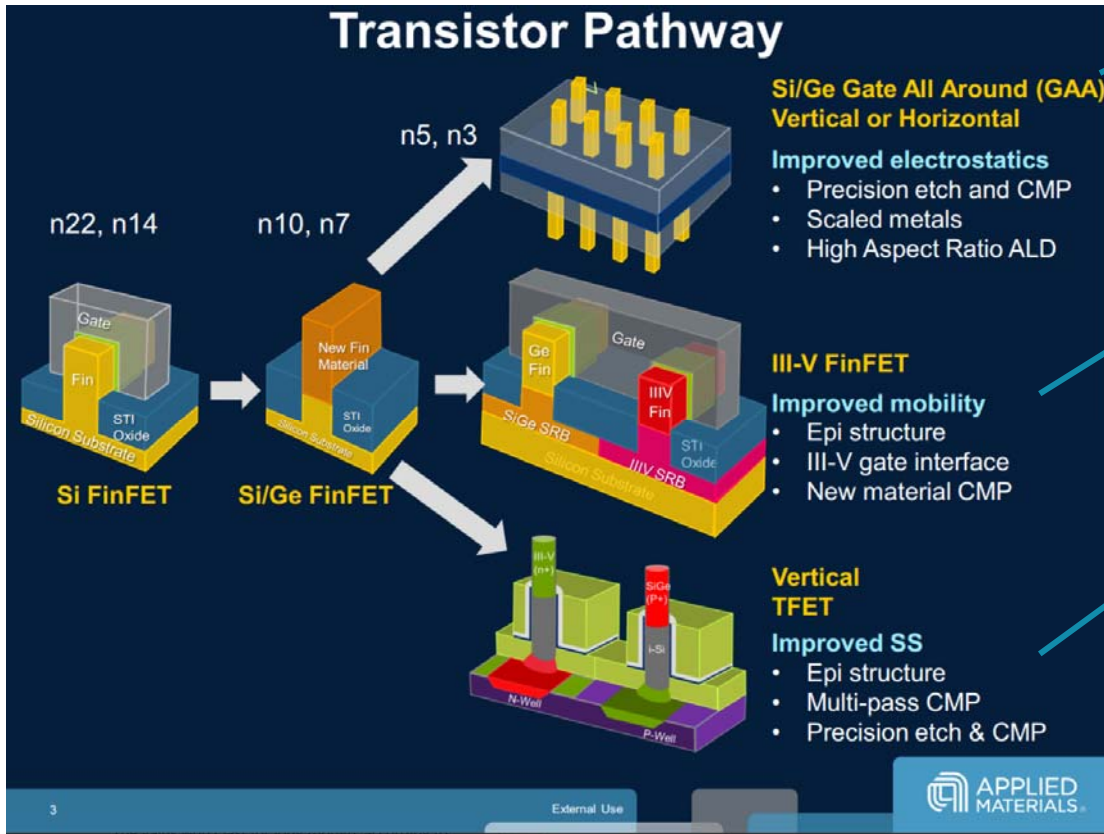
This doesn't happen automatically anymore.

From 28nm, we are often in the Reverse Dennard Scaling era: Smaller gate pitch can often mean slower (and larger!) chips

PPA improvement? Sure...pick any 2.

Pressure to bifurcate processes becomes stronger

After 7nm: lots of options, no clear path



Very disruptive to design

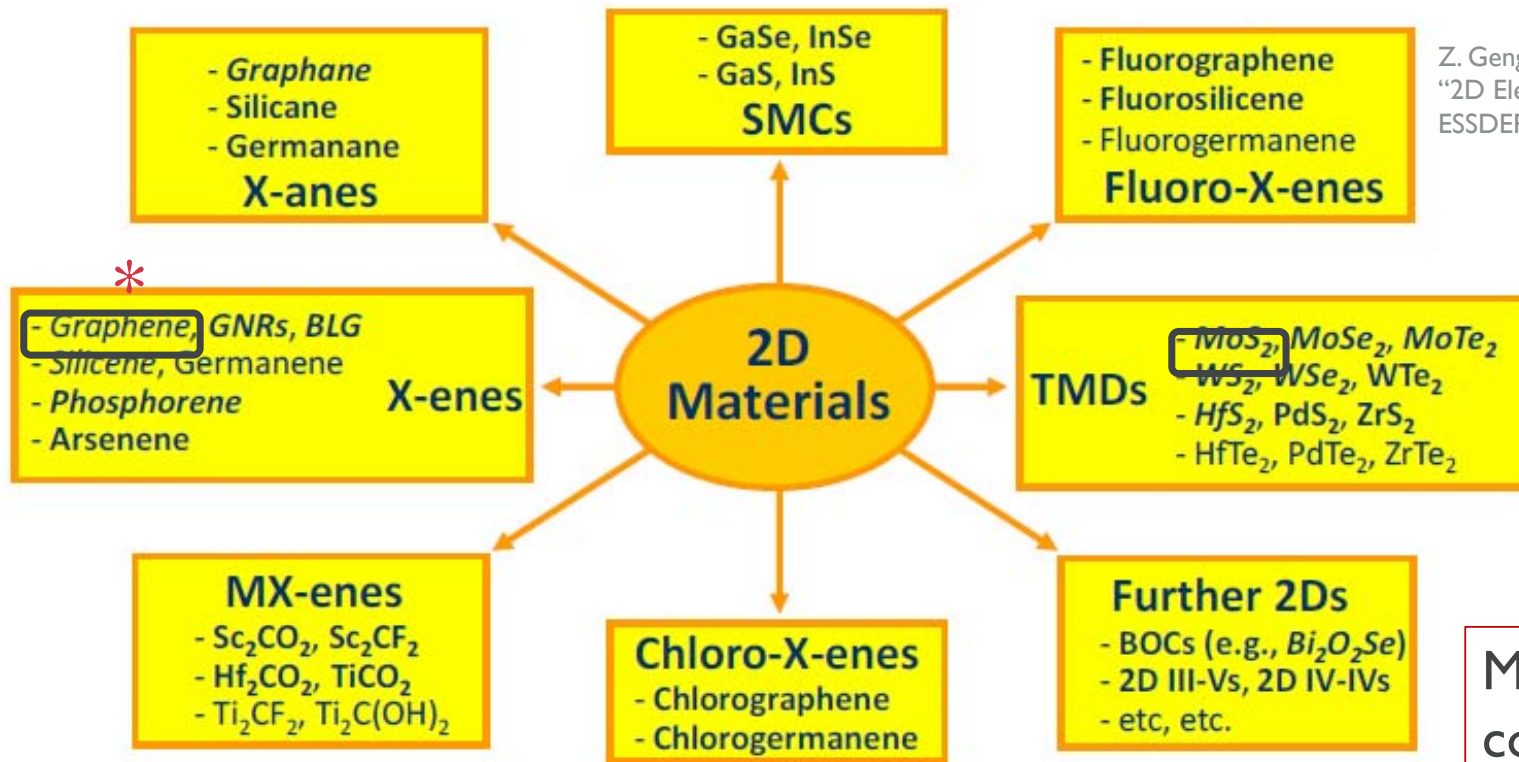
Most high μ devices suffer from Density of States and excessive leakage.
 Maybe OK for HPC, maybe not for mobile
 → CMOS as material #1 + material #2
Beginning of heterogeneous “CMOS”

TFETs are asymmetric: disruptive to design
 Existing TFETs are too slow

Dimensional scaling augmented by new materials and new physics

Challenge: finding the right amount of disruption

Example of the materials revolution: 2D materials



Z. Geng et al.,
 "2D Electronics - Opportunities and Limitations"
 ESSDERC 2016

* Metallic graphene could possibly be a future wire material


Many 2D materials could operate as SpinFETs

These are just the raw materials. Then you can dope, give them edge treatments, etc.

Graphene opened the Pandora's box of 2D materials.

BEOL Scaling summary:

Half of your performance and power is going here



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

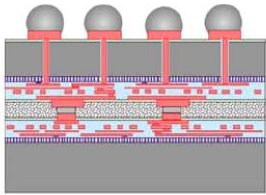
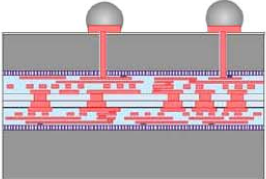
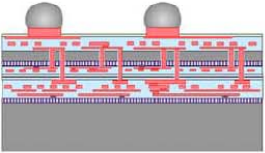
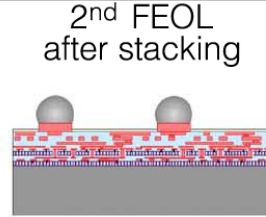
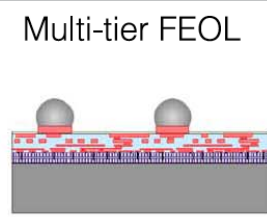
- Future transistors will struggle with severe parasitic load
→ we're already in reverse-Dennard scaling
- Low-k (airgap), low-R (cobalt, ruthenium): mitigate, but don't solve
- Problems are not just lateral, but more and more vertical (Via R)
- **Power/Performance scaling may be more challenged than cost scaling**

Beyond transistors and wires

Two technologies that could really help the system roadmaps:

- Find a better memory
 - PCM, RRAM, MRAM: Nothing yet is a “super” memory.
(from an on-die cache perspective. Storage class memory is great for systems, but another topic)
 - Endurance will be a key system factor (i.e., can we work with a memory less robust than SRAM)
 - Non-volatility is a big wild card, both in traditional cache hierarchy and power states
- Utilize 3DIC
 - Look at additional partitioning options to optimize PPAC.
 - There is a wide variety of present and future technologies to choose from.
 - Future technologies need design tools (chicken and egg problem)

The 3DIC Roadmap: homogenous SiP to transistors

	3D-SiC	3D-SOC	3D-SOC	3D-IC	
3D-Wiring level	Global	Semi-global	Intermediate	Local	FEOL
Partitioning	Die	blocks of standard cells		Standard cells	Transistors
3D Technology	Die stacking Die-to-Wafer stacks Die-to-Si-interposer	Parallel FEOL wafer processing Wafer-to-Wafer bonding		Sequential FEOL processing Active layer bonding or deposition	
2-tier stack schematic					
Characteristic	Known Good Die 3D stacks or Si-interposer stacking	BEOL between 2 FEOL layers <i>Overlay 2nd tier defined by W2W alignment/bonding</i>		FEOL/FEOL stack <i>Overlay 2nd tier defined by litho scanner alignment</i>	
Contact Pitch	40 ⇒ 20 ⇒ 10µm ⇒ 5µm	5 ⇒ 1µm	2 µm ⇒ 0.5 µm	200 ⇒ 100 nm	< 100 nm
Relative density:	1/16 ⇒ 1/4 ⇒ 1 ⇒ 4	4 ⇒ 100	50 ⇒ 400	5000 ⇒ 10000	> 10000


E. Beyne, IEEE Design & Test, May/June 2016

Homogeneous SiP

Today's technologies already helping high end system scaling.
(FPGA is also a good example of what Moore envisioned)


“It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected”
- Gordon Moore, 1965

	3D-SiC	
3D-Wiring level	Global	Sem
Partitioning	Die	blo
3D Technology	Die stacking Die-to-Wafer stacks Die-to-Si-interposer	Par
		Wafer-to-Wafer bonding
		Active layer bonding or deposition



4-High HBM2 Stack with Base Die

Bumps



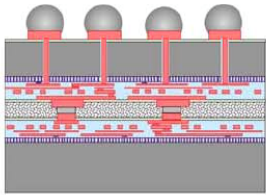
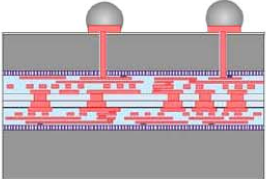
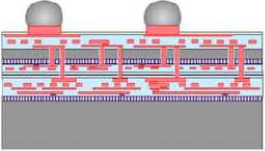
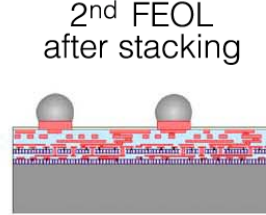
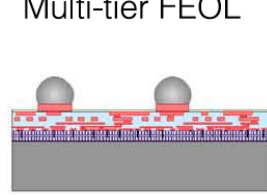
GPU

Silicon Carrier

Substrate

Nvidia Pascal claim: >3x GBps/watt improvement.

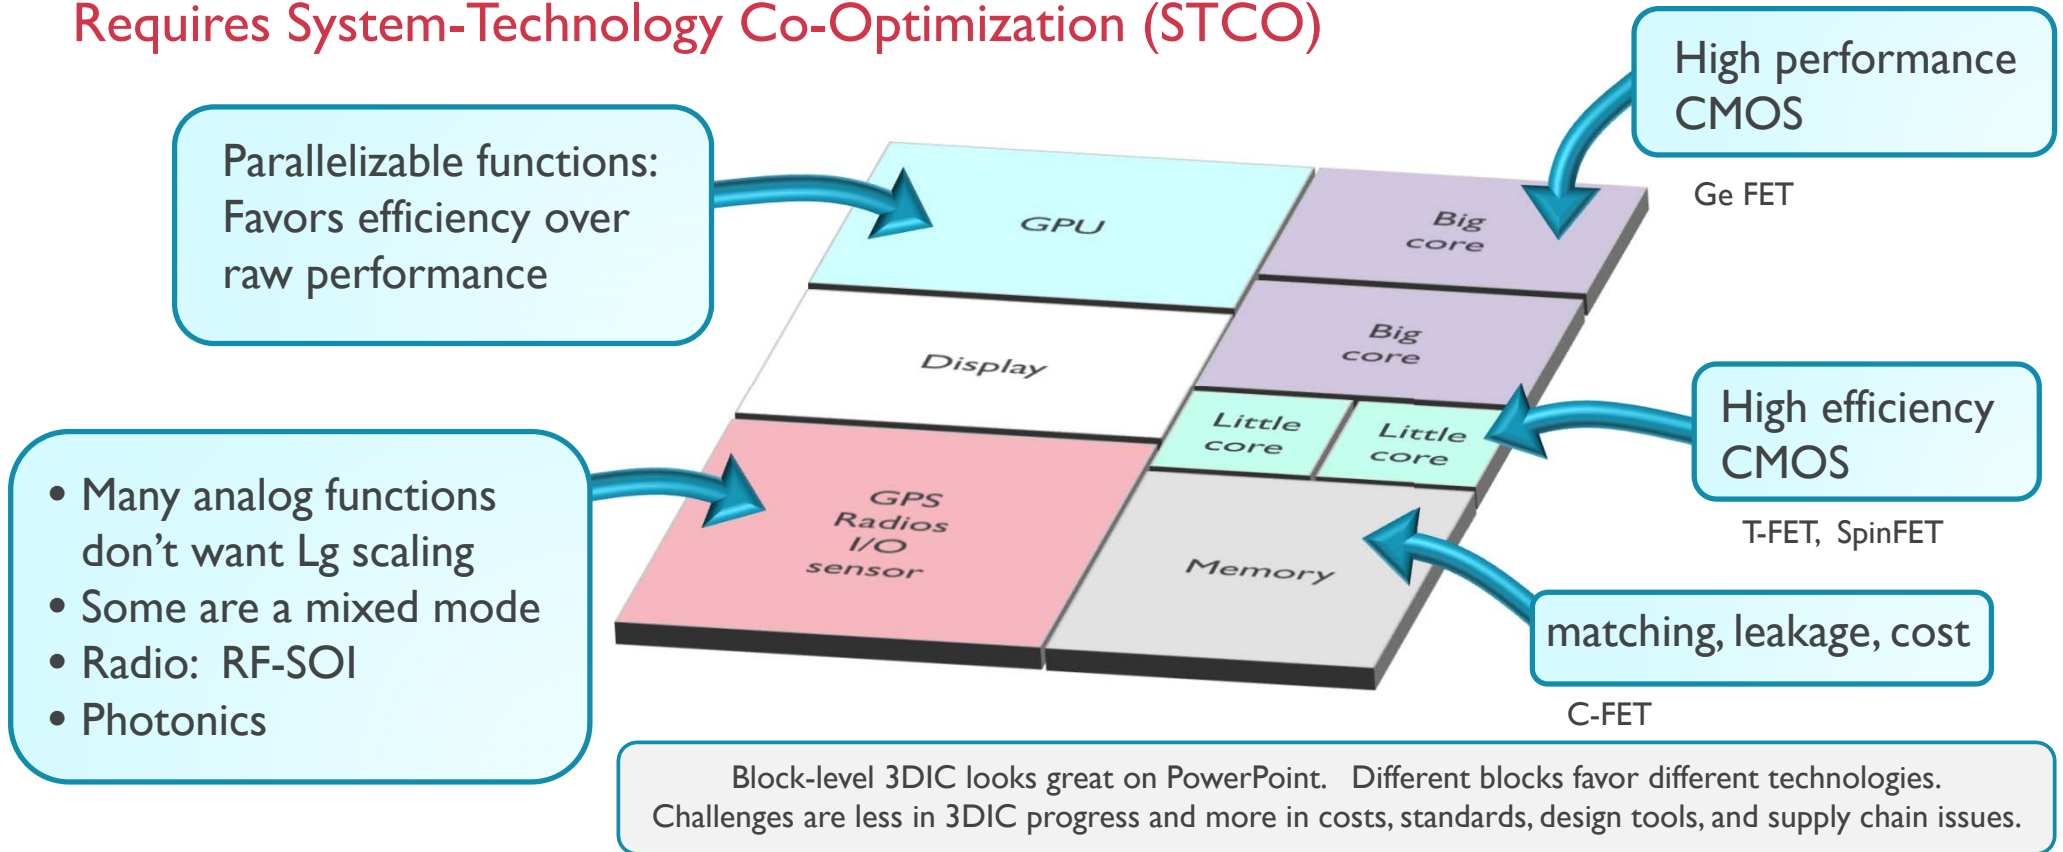
heterogeneous SiP: “3D-SOC” → partitioning

	3D-SiC	3D-SOC	3D-SOC	3D-IC	3D-IC
3D-Wiring level	Global	Semi-global	Intermediate	Local	FEOL
Partitioning	Die	blocks of standard cells	Standard cells	Transistors	
3D Technology	Die stacking Die-to-Wafer stacks Die-to-Si-interposer	Parallel FEOL wafer processing Wafer-to-Wafer bonding		Sequential FEOL processing Active layer bonding or deposition	
2-tier stack schematic					
Characteristic	Known Good Die 3D stacks or Si-interposer stacking	BEOL between 2 FEOL layers <i>Overlay 2nd tier defined by W2W alignment/bonding</i>	BEOL between 2 FEOL layers <i>Overlay 2nd tier defined by litho scanner alignment</i>	FEOL/FEOL stack	
Contact Pitch	40 ⇒ 20 ⇒ 10µm ⇒ 5µm	5 ⇒ 1µm	2 µm ⇒ 0.5 µm	200 ⇒ 100 nm	< 100 nm
Relative density:	1/16 ⇒ 1/4 ⇒ 1 ⇒ 4	4 ⇒ 100	50 ⇒ 400	5000 ⇒ 10000	> 10000

E. Beyne, IEEE Design & Test, May/June 2016

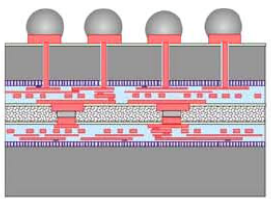
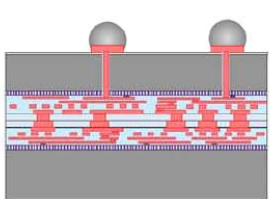
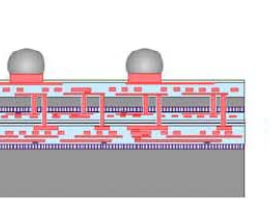
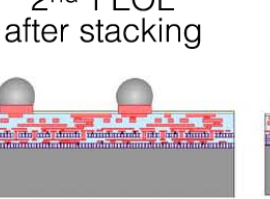
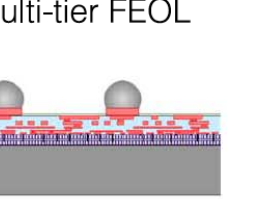
Block-level partitioning: 3D-SoC optimizes PPAC

Requires System-Technology Co-Optimization (STCO)



High density 3DIC

Tools that can bond wafers at this density are in the pipeline.
This would allow folding inside the cores themselves

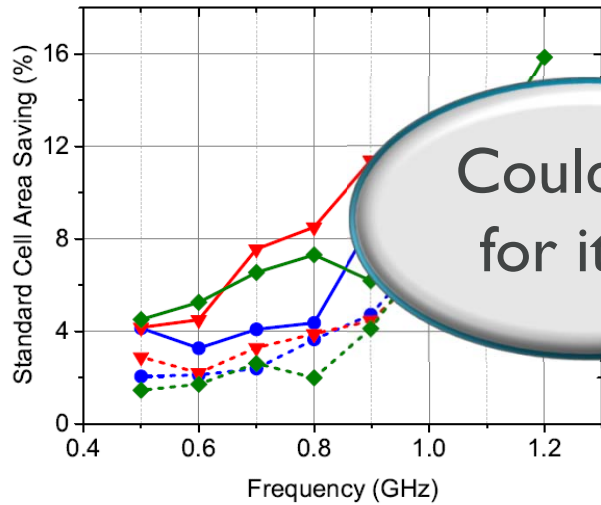
	3D-SIC	3D-SOC	3D-SOC	3D-IC	
3D-Wiring level	Global	Semi-global	Intermediate	Local	FEOL
Partitioning	Die	blocks of standard cells	Standard cells	Transistors	
3D Technology	Die stacking Die-to-Wafer stacks Die-to-Si-interposer	Parallel FEOL wafer processing Wafer-to-Wafer bonding	Sequential FEOL processing Active layer bonding or deposition		
2-tier stack schematic					
Characteristic	Known Good Die 3D stacks or Si-interposer stacking	BEOL between 2 FEOL layers <i>Overlay 2nd tier defined by W2W alignment/bonding</i>	BEOL between 2 FEOL layers <i>Overlay 2nd tier defined by litho scanner alignment</i>	FEOL/FEOL stack	
Contact Pitch	40 ⇒ 20 ⇒ 10µm ⇒ 5µm	5 ⇒ 1µm	2 µm ⇒ 0.5 µm	200 ⇒ 100 nm	< 100 nm
Relative density:	1/16 ⇒ 1/4 ⇒ 1 ⇒ 4	4 ⇒ 100	50 ⇒ 400	5000 ⇒ 10000	> 10000

E. Beyne, IEEE Design & Test, May/June 2016

Gate-level partitioning: 3D-IC

ARM / Georgia Tech study to be published at ICCAD.
There's a benefit, but is it a large enough one time bump to justify?

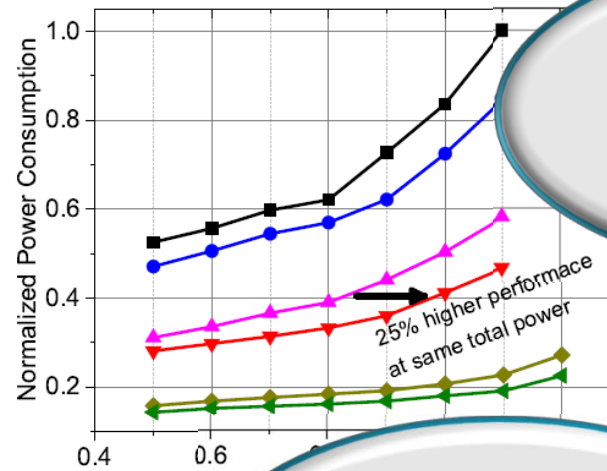
M3DIC vs. 2D: Area



Could pay for itself

- 28nm Cascade2D ● 28nm Shrunken2D
- ▲ 14/16nm Cascade2D ▲ 14/16nm Shrunken2D
- ◆ 7nm Cascade2D ◆ 7nm Shrunken2D

M3DIC vs. 2D: Power/Performance



While delivering power and/or performance uplift

Enough to justify a one-time bump?

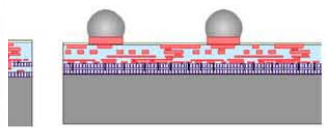
Aforementioned chicken and egg problem


More details to be presented at ICCAD 2016:

K. Chang, S. Sinha, B. Cline, R. Southerland, M. Doherty, G. Yeric and S. K. Lim,

“Cascade2D: A Design-Aware Partitioning Approach to Monolithic 3D IC with 2D Commercial Tools”

High density 3DIC at the transistor level

	3D-SIC	3D-SOC	3D-IC
3D-Wiring level	Global	Die	FEOL
Partitioning	Die	Die	Transistors
			FEOL processing
			Bonding or deposition
			Multi-tier FEOL
			
			FEOL/FEOL stack
			defined by litho scanner
			< 100 nm
			> 10000




Options include:

- Memory over logic
- Efficient CMOS over Fast CMOS
- PMOS over NMOS (CFET)

	3D stacks or Si-interposer stacking	Overlaid alignment
Contact Pitch	40 \Rightarrow 20 \Rightarrow 10 μm \Rightarrow 5 μm	5 \Rightarrow 1
Relative density:	1/16 \Rightarrow 1/4 \Rightarrow 1 \Rightarrow 4	4 \Rightarrow 1

E. Beyne, IEEE Design & Test, May/June 2016

Pressure on R&D costs in the industry



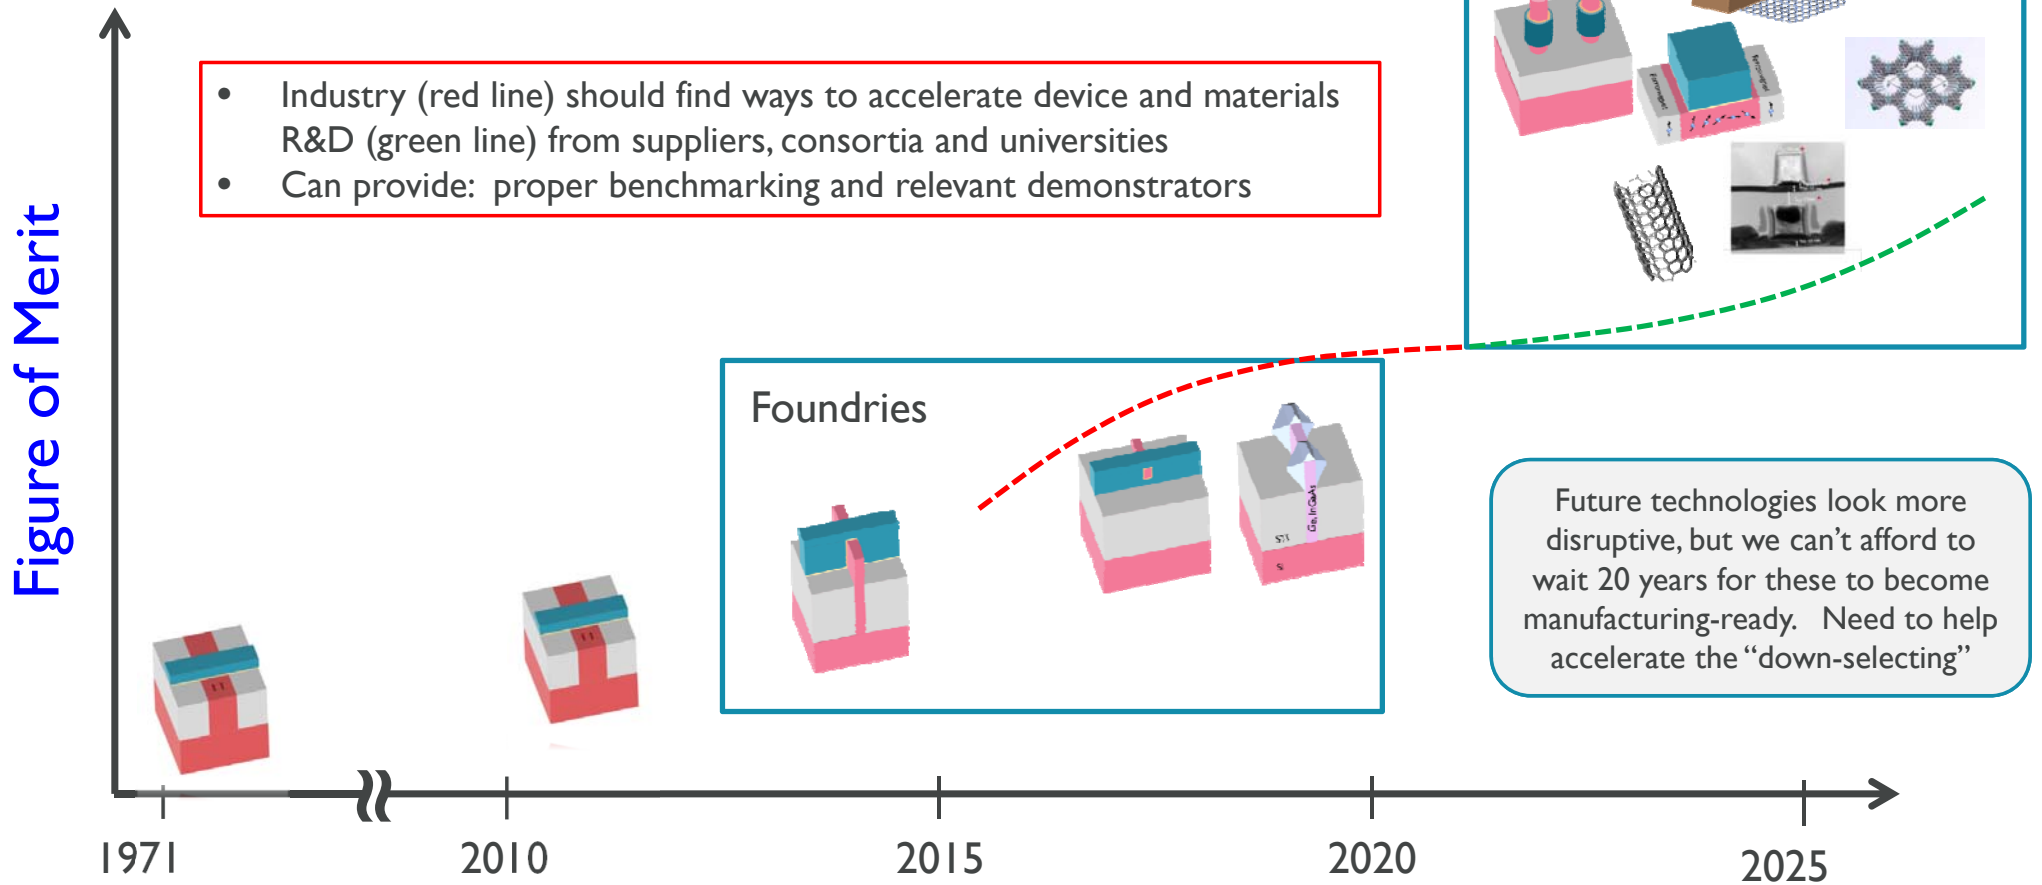
http://theconfab.com/wp-content/uploads/2013/07/cf13_lode_lauwers.pdf

source: IHS iSupply
Sematech Report
IBS 2011

Industry has its hands full with 5nm

The needs and complexities of solutions for 3/2 nm solutions will be unprecedented

Future technology pipeline



Examples of ARM activities along this theme:

ARM Cortex™-M0 DesignStart

> 150 university downloads!



IP and Development Boards

ARM has announced the availability of the ARM® Cortex™-M0 DesignStart Processor via the [ARM DesignStart™ online IP portal](#) as a download after a simple click-through EULA. The processor is configured as a synthesizable, obfuscated verilog netlist designed for academics, start-ups and ad-hoc technology teams looking to teach or prototype with a real ARM processor. Implementation is possible in almost any FPGA to enable SoC design courses and projects around a real, instruction set compatible ARM core. It can also be manufactured according to the EULA. From there, anything from embedded systems courses or microprocessor applications can be taught using the same platform. Through this online access model, ARM will accelerate the proliferation of ARM technology in university curriculums and research projects.

The Official ARM Cortex™-M0 DesignStart Example Design Kit (EDK) is Now Freely Downloadable!

The [ARM Cortex-M0 DesignStart Example Design Kit \(EDK\)](#) helps to start designing complete systems (hardware + software) with the Cortex-M0 DesignStart Processor on an FPGA board. It includes a selection of simple AHB-Lite peripherals, AHB-Lite bus infrastructure components, example systems usable as design templates, software examples based on the popular Keil MDK-ARM software development tools, and associated quickstart guides and documentation. It's everything you need to get an example system based on the M0_DS running out-of-box in a [Xilinx-based Digilent Nexsys3 FPGA board!](#)

<https://www.arm.com/support/university/researchers/soc-fpga/index.php>

2015: ARM joins Semiconductor Research Corporation

The screenshot shows the SRC website's member roster page. The breadcrumb trail is "Home » SRC » Member » Roster". Under "SRC Members", there are tabs for "GRC", "STARnet", and "NRI". The "GRC" tab is selected, displaying a table of members.

Members	
Advanced Micro Devices, Inc.	Analog Devices, Inc.
ARM, Incorporated	Freescale Semiconductor, Inc.
GLOBALFOUNDRIES	IBM Corporation
Intel Corporation	Mentor Graphics Corporation
Microsoft	Mubadala Technology
Qualcomm	RTI International
Texas Instruments Incorporated	Tokyo Electron Limited

The screenshot shows the "GRC THRUSTS" section of the website, listing various research areas:

- AMS - CSD**
Analog/Mixed-Signal Circuits, Systems, and Devices
- CADT**
Computer-Aided Design and Test
- EP3C**
Efficiency and Performance for Connectivity Constrained Computing
- ESH**
Environment, Safety, and Health
- I3T**
Innovative and Intelligent Internet of Things
- LMD**
Logic and Memory Devices
- NMP**
Nanomanufacturing Materials and Processes
- PKG**
Packaging
- SemiSynBio**
Semiconductor Synthetic Biology
- SLD**
System Level Design
- T3S**
Trustworthy and Secure Semiconductors and Systems

The screenshot shows the "GRC CENTERS" section of the website, listing various research centers:

- ACE4S**
ATIC-SRC Center of Excellence for Energy Efficient Electronics Systems
- CAIST**
The New York Center for Advanced Interconnect Science and Technology
- EBSM**
SRC Engineering Research Center for Environmentally Benign Semiconductor Manufacturing
- TxAACE**
Texas Analog Center of Excellence

2nd Annual ARM Research Summit: Sept 11-13, 2017

ARM Research Summit Computing for the next decade


15-16 September 2016 | Churchill College, Cambridge, UK

Over 70 talks

Example “Moore’s Law” area papers:

Scaling and Variability of Multi-Gate Non-Planar III-V	Karol Kalna	Swansea University
Optical Computing: A technology Whose Time Has Come?	C David Wright	University of Exeter
Design Technology Co-Optimisation in Advance Technology Nodes	Asen Asenov	University of Glasgow
3D IC Revolution: Where Are We Now, and What is Next?	Sung-Kyu Lim	Georgia Tech

Broad ecosystem DTCO: ARM / ASML collaboration



14nm node ARM M1 clip without OPC, 46nm minimum pitch, exposed on an NXE:3300B with conventional illumination

van den Brink, ASML, ISSCC 2013 Keynote



2016: ARM joins IMEC Logic INSITE program



ARM is contributing by helping IMEC improve benchmarking of “scaling boosters”

News & Analysis
IMEC, ARM Collaborate on 7nm Design

Peter Clarke
7/11/2016 05:09 PM EDT
Post a comment

NO RATINGS
LOGIN TO RATE

Like 16 Tweet in Share 8 G+1 1

LONDON--Processor intellectual property licensor ARM Holdings plc (Cambridge, England) has signed on to the INSITE collaborative research program on design at nanoelectroics research institute IMEC (Leuven, Belgium).

The INSITE program was started in 2009 and now has more than 10 participants with a focus on enabling design for chip process nodes at 7nm and beyond.

INSITE makes use of the knowledge gained in IMEC's lithography and logic device scaling programs to help companies anticipate design parameters and options for next-generation systems and applications. Developers are now faced by a large number of potential design choices at 7nm including: the required number of lithography exposures, device architecture such as FinFETs or lateral nanowires, the local interconnect scheme, cell architecture and the metallization scheme.

2016: ARM joins Stanford SystemX Alliance

Stanford | ENGINEERING
SystemX Alliance

- **Bio Interfaces:** Transformative technologies for healthcare | **LEAD:** Prof. Shan Wang and Prof. Tom Soh | **DETAILS**
- **Design Productivity:** Improving critical area productivity | **LEAD:** Prof. Mark Horowitz | **DETAILS**
- **Energy/Power Management Systems:** Efficient energy power supply | **LEAD:** Prof. Juan Rivas-Davila | **DETAILS**
- **Heterogeneous Integration:** Bridging everything into anything | **LEADS:** Prof. Eric Pop and Prof. H.S. Philip Wong | **DETAILS**
- **Internet of Everything:** Universal connectivity of devices and systems | **LEAD:** Prof. Amin Arbabian and Prof. Thomas Lee | **DETAILS**
- **Photonic and Quantum Technologies:** More efficient devices and systems; optics | **LEAD:** Prof. Jelena Vuckovic | **DETAILS**
- **Computation for Data Analytics:** Optimizing application, architecture & technology | **LEAD:** Prof. Chris Ré and Prof. Subhasish Mitra | **DETAILS**

To be presented at IEDM (December)

32-bit Processor Core at 5-nm Technology:
Analysis of Transistor and Interconnect Impact on VLSI System Performance
 Chi-Shuen Lee¹, Brian Cline², Saurabh Sinha², Greg Yeric², H.-S. Philip Wong¹
¹Deptment of Electrical Engineering, Stanford University, Stanford, CA 94305, USA. Email: chishuen@stanford.edu
²ARM Inc., Austin, TX 78735, USA.

This paper to be presented at IEDM is a good example of vetting future device candidates in a full core design context, to help get more accurate projections of value

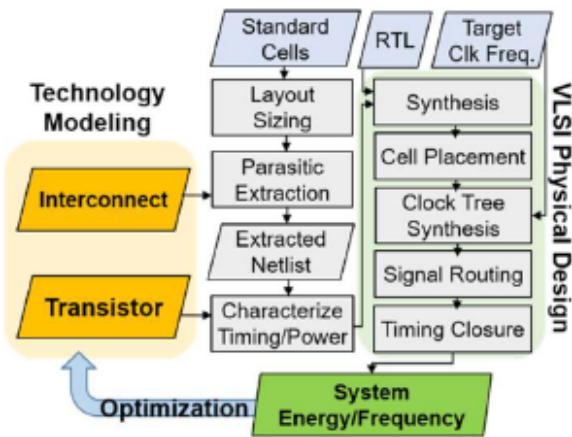


Fig. 1. Integrated design flow for technology modeling and VLSI system implementation.

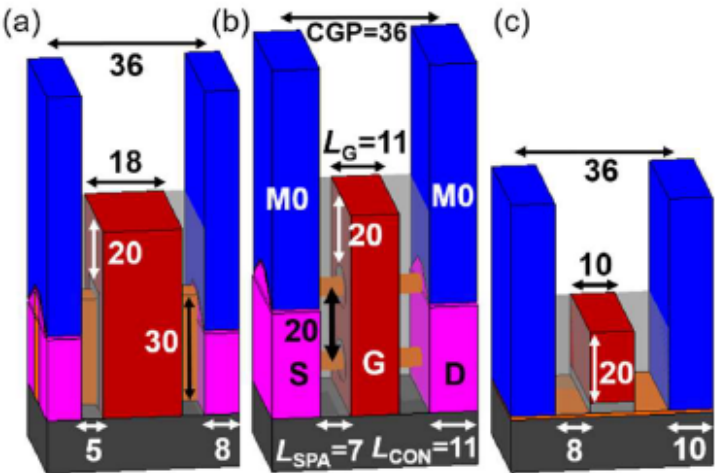



Fig. 2. FET device structures of (a) FinFET, (b) NWFET, and (c) 2D-FETs. All labeled numbers are in the unit of nm. The Fin and NW widths are both 5 nm.

ASU ASAP 7nm predictive PDK

<http://asap.asu.edu/asap/>

ASAP/ASU 7nmPDK
In partnership with ARM



ASAP7: Full 7nm 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



Contents lists available at [ScienceDirect](#)

Microelectronics Journal

journal homepage: www.elsevier.com/locate/mejo



Microelectronics Journal 53 (2016) 105–115

ASAP7: A 7-nm finFET predictive process design kit



Lawrence T. Clark^{a,*}, Vinay Vashishtha^a, Lucian Shifren^b, Aditya Gujja^a, Saurabh Sinha^c,
Brian Cline^c, Chandarasekaran Ramamurthy^a, Greg Yeric^c

^a Arizona State University, Tempe, AZ 85287, USA

^b ARM Inc., San Jose, CA 95134, USA

^c ARM Inc., Austin, TX 78735, USA

ARTICLE INFO

Article history:

Received 22 December 2015

Received in revised form

30 March 2016

Accepted 10 April 2016

Keywords:

Predictive process design kit

7-nm technology

Process scaling

Extreme ultraviolet lithography

ABSTRACT

We describe a 7-nm predictive process design kit (PDK) called the ASAP7 PDK, developed in collaboration with ARM Ltd. for academic use. The PDK is realistic, based on current assumptions for the 7-nm technology node, but is not tied to any specific foundry. The initial version assumes EUV lithography for key layers.

(SAQ) ...
imm ...
line ...
MOL ...
layer ...
81138

Free download for university researchers.
Allows novel technology benchmarking in an
accurate circuit context that is more indicative
of a real 7nm/5nm design.

University eNVM incubation project

Too early to understand if this particular material can be successfully integrated (and the road is littered with well-intentioned NVM development projects), but this is a good example of the materials-to-design expediting theme. Perhaps tune in to TechCon 2017 to find out.

- Find a better memory
 - PCM, RRAM, MRAM: Nothing yet is a “super” memory. (for on-die cache. SCRM is great, but another topic)
 - Endurance will be a key system factor (i.e., can we work with a memory less robust than SRAM)
 - Non-volatility is a big wild card, both in traditional cache hierarchy and power states
- Utilize 3DIC
 - Look at additional partitioning options
 - There is a wide variety of present and future technologies need design to

Lab-to-fab proof of concept material to design: what a company like ARM can provide:

- Design domain knowledge
- Test chip design / test / analysis

Summary

- **Traditional Moore's Law cost reduction is slowing down**
 - Broad ecosystem progress plus scaling boosters can achieve 5nm (3nm?) with manageable disruption
 - Slowing of traditional scaling means more radical scaling options will have opportunity
- **Power/performance scaling may be more difficult than cost (transistors + wires)**
 - Scaling obstacles will be increasingly require novel integration, materials, devices, and physics
- **Heterogeneity: challenges and opportunities**
 - Different materials for NMOS / PMOS, LVT / HVT, local wires vs. global wires. Varying pitches.
 - 3D-SiP → 3D-SOC → 3DIC: Huge design challenges. DTCO → STCO
 - Novel memories, photonics, (3D-SOC, moving to 3D-IC), other wild cards
- **Continued progress will require managing the rate and level of change**
 - Increasing “top to bottom” DTCO including semiconductor tool manufactures, EDA, etc.
 - Accelerate novel technology from universities and research consortia: “lab to fab”
 - Benchmarking is a key topic. DTCO → STCO

Among the topics I did not talk about:

- The IoT:
 - Reducing cost/function (more Moore) vs. enabling new functions (more than Moore)
 - Moore's Law progress enabled the Internet of Things
 - Vanishing cost of compute, plus reasonable cost of big compute (data analytics)
 - Ubiquitous internet, and vanishing cost to connect
 - **Children of Moore's law:** Energy harvesting, sensors, printed electronics
 - These new capabilities are heavily leveraging the progress of the semiconductor industry
 - 10um to nm progress in 50 years: sensors can span from cell to virus to protein / molecule
→ Materials “functionalization”, add smell and taste, beyond human capabilities. (medical)
- Wacky things that could re-write the scaling math
 - Bending physics with nm-scale capabilities: Plasmonics, energy filtering, metamaterials, ...
 - Spintronics, non-volatile transistors, BEOL eNVM as FPGA, neuromorphic devices, ...



Thank you