

Heterogeneous Integration at DARPA: Pathfinding and Progress in Assembly Approaches

Daniel S. Green

U.S. Defense Advanced Research Projects Agency (DARPA)

Arlington, VA

(now with Office of Naval Research, Arlington, VA)


Presented by Jeffrey C. Demmin, Booz Allen Hamilton
68th IEEE Electronic Components and Technology Conference

29 May 2018



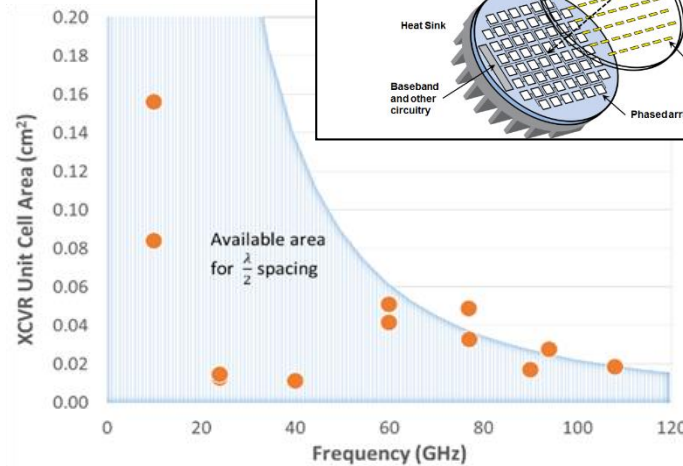



Why Heterogeneous Integration?




- Si CMOS/SiGe BiCMOS
- InP HBTs/HEMTs
- GaN HEMTs
- RF MEMS/High-Q passives

Best of breed technology



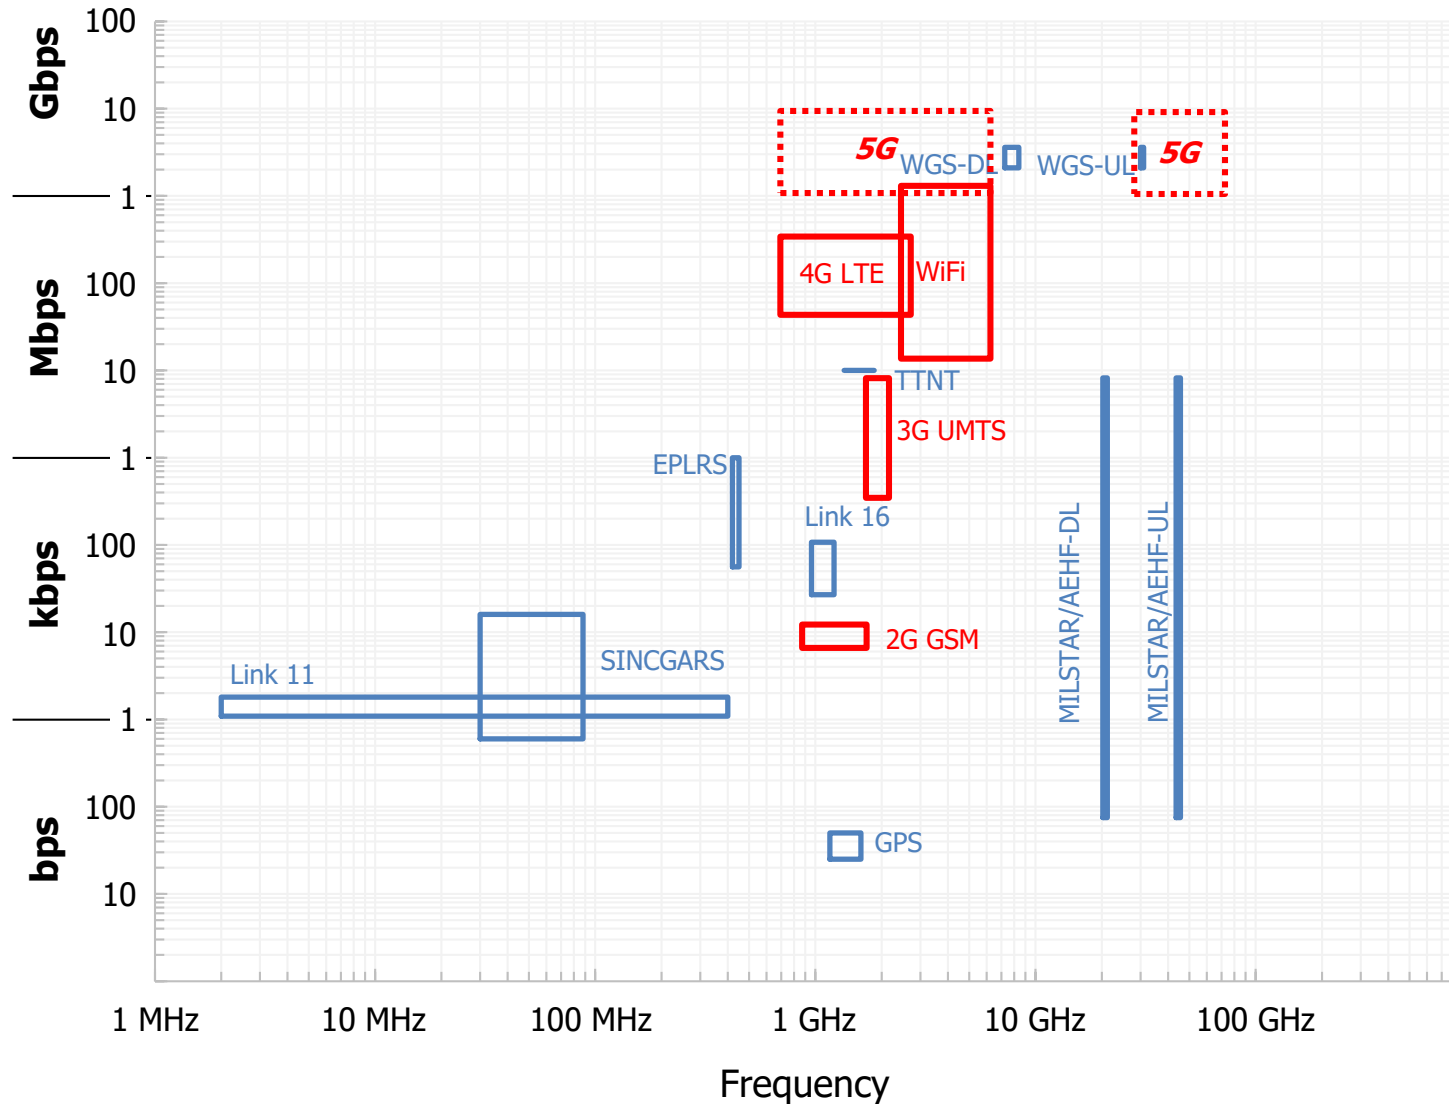
Density of integration



Modular design




Civilian and defense needs overlap at the leading edge






DARPA's long history of innovation in integration


ASEM




COSMOS




CHIPS




VISA




3D-IC




DAHI



E-PHI



MOABB



- ASEM: Application Specific Electronic Modules
- E-PHI: Electronic-Photonic Heterogeneous Integration
- VISA: Vertically Integrated Sensor Arrays
- COSMOS: Compound Semiconductor Materials on Silicon
- DAHI: Diverse Accessible Heterogeneous Integration
- MOABB: Modular Optical Aperture Building Blocks
- CHIPS: Common Heterogeneous Integration and IP Reuse Strategies

1990s


2000s

2010s


2020s



What is heterogeneous integration?



CHIPS: modular design

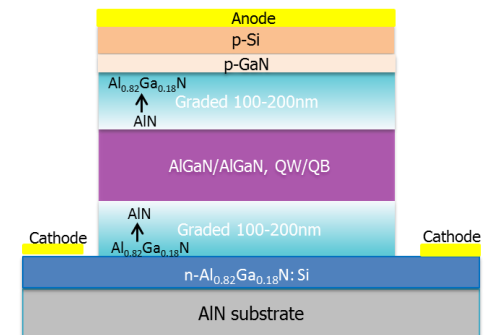
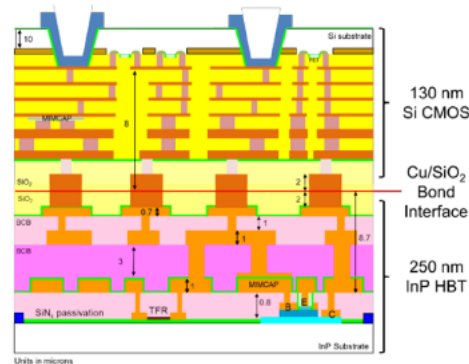


Heterogeneous devices

DAHI: RF power

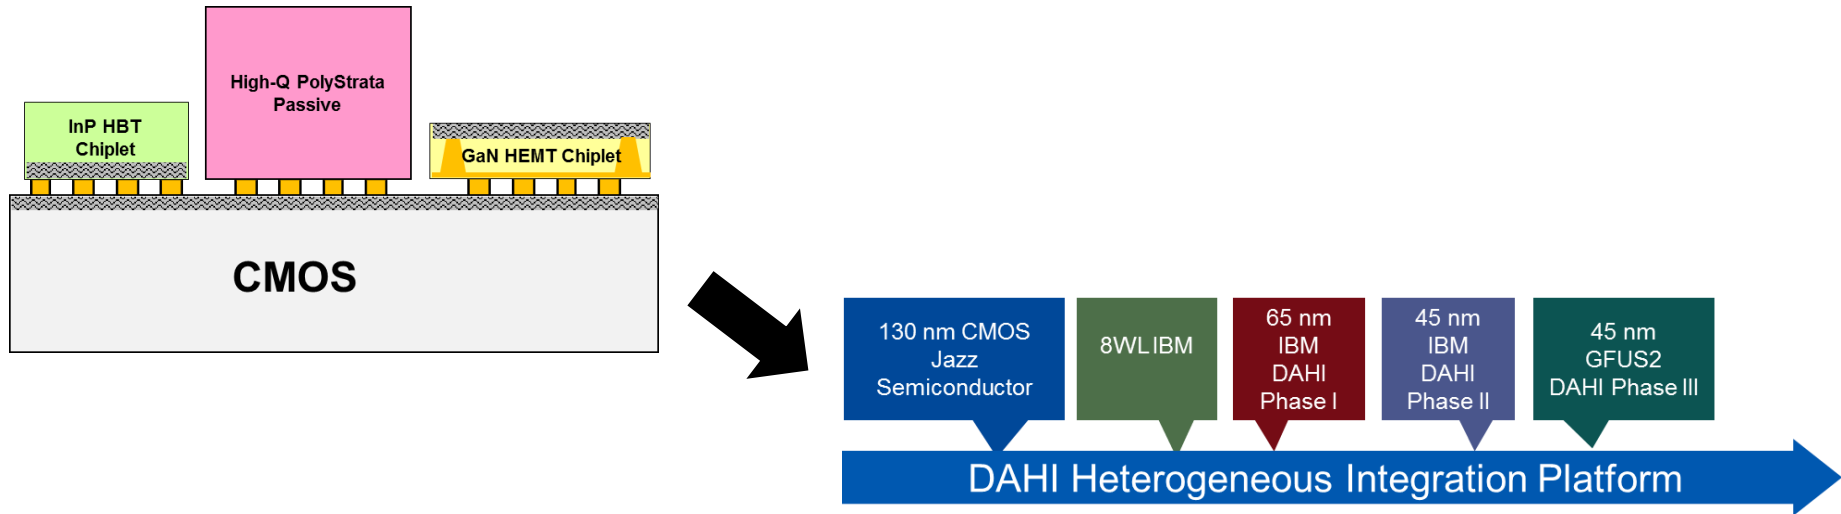


DAHI: wafer scale







DARPA's DAHI program: Diverse Accessible Heterogeneous Integration




Picture of MPW0 Reticle



Picture of MPW1 Reticle




MPW2




MPW3 in Fab



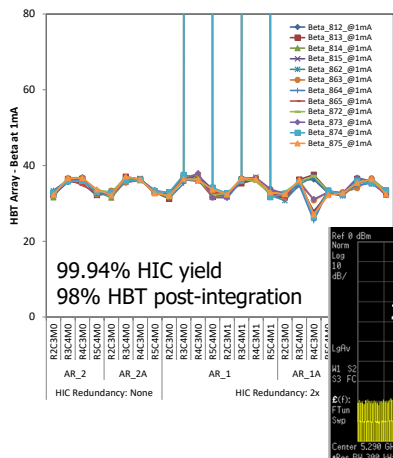
DAHI snapshot: Excellent yield, demonstrated RF performance





300mm diameter Si CMOS wafer (45nm node)



DAHI integration: Si (45nm), InP (TF5 HBT), GaN (GaN20 HEMT)



DAC with very low digital noise (-70dBc)



Successful testing identified optimal S/H circuit for ADC (>65dB SFDR @ 2GHz)

Sources: DARPA, Northrop Grumman



DAHI simplicity enables rapid evolution

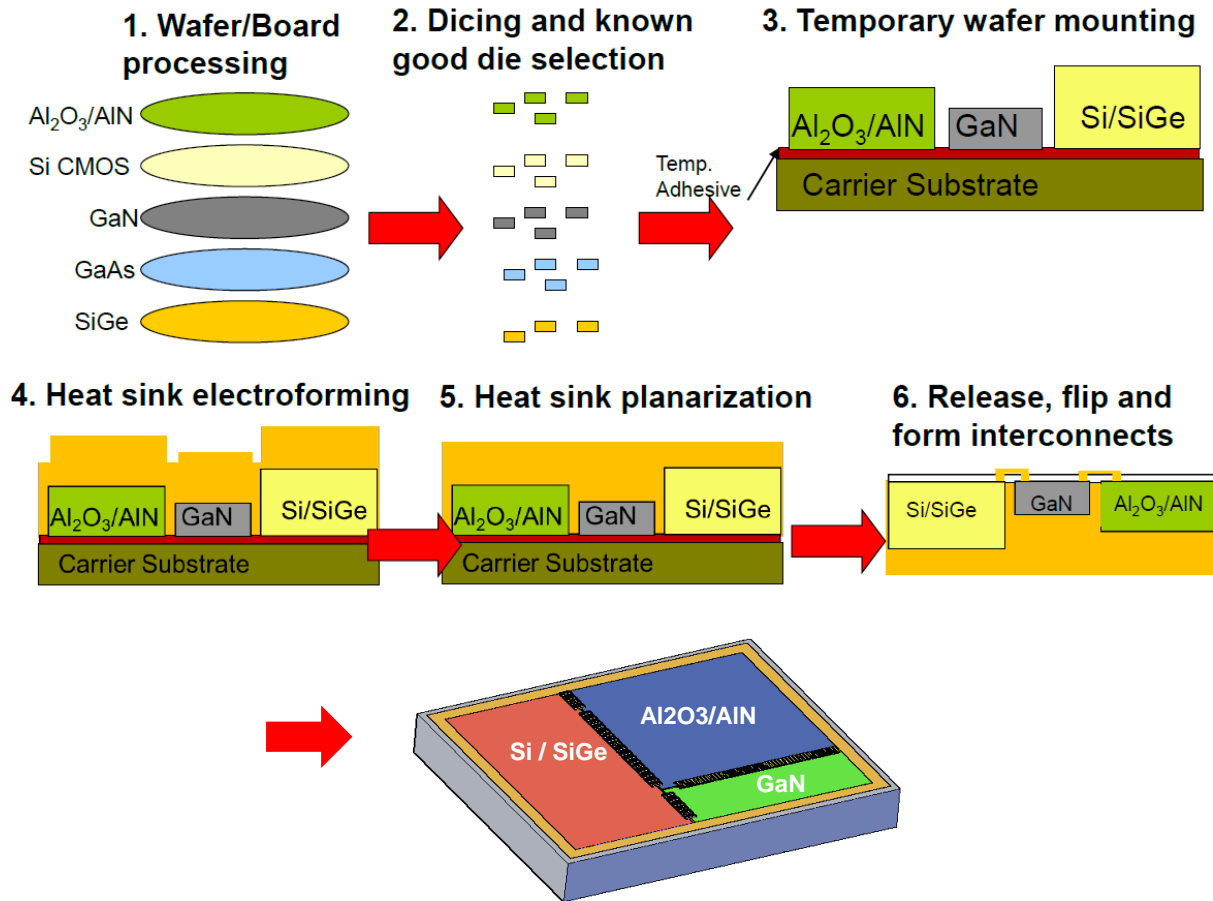
Technology	MPW0	MPW1	MPW2	MPW3	Future MPWs
CMOS	IBM 65nm	GF 45 nm	GF 45 nm	GF 45 nm	GF 45 nm
InP HBT	TF4 (2 metals)	TF4 (3 metals)	TF4 (4 metals)	TF4 (4 metals)	TF4 (4 metals)
		TF5 (3 metals)	TF5 (4 metals)	TF5 (4 metals)	TF5 (4 metals)
InP Varactor Diode					AD1
GaN HEMT	GaN20	GaN20	GaN20	GaN20	GaN20
	T3 (HRL)	T3 (HRL)	T3 (HRL)	T3 (HRL)	T3 (HRL)
GaAs HEMT				P3K6	P3K6
Passive Components		PolyStrata (Nuvotronics)	PolyStrata (Nuvotronics)	PolyStrata (Nuvotronics)	PolyStrata (Nuvotronics)
Base Substrate	CMOS	CMOS	CMOS	CMOS	CMOS
				SiC Interposer (IWP5)	SiC Interposer (IWP5)
				<p>In fab</p>	

Sources: DARPA, Northrop Grumman




Metal Embedded Chip Assembly (MECA)

MECA enables heterogeneous integration with a metal interconnect platform for high-power requirements.



MECA-integrated heterogeneous module



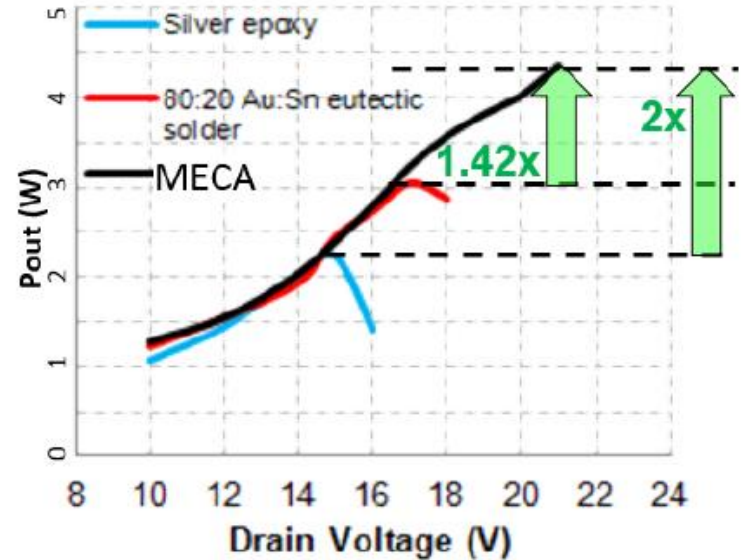
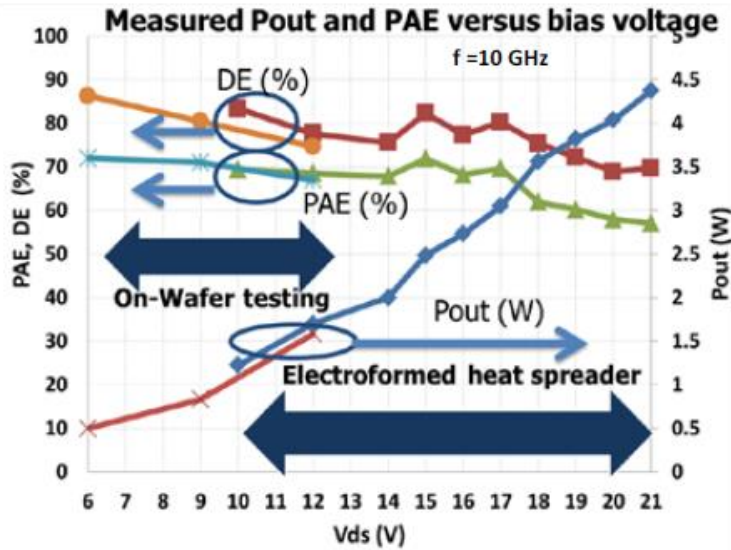



Sources: HRL, Solid State Technology



MECA-enabled performance upgrade

Integration in electroformed heat spreader: 1.4-2x improvement in PA performance





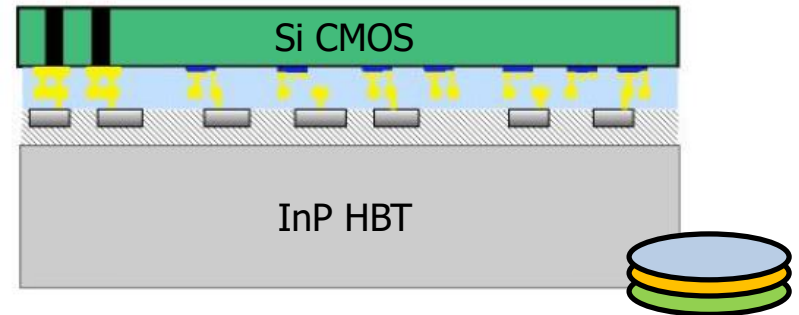
DAHI chip-scale phased arrays

Heterogeneous integration for mm-wave:

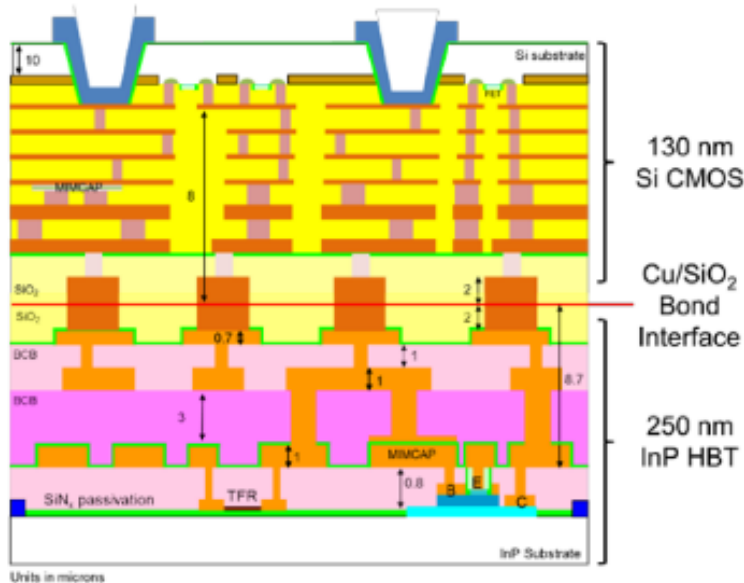
Phased array beamformers

- Can maintain $\lambda/2$ channel spacing as frequencies increase
- CMOS control circuitry closely integrated with RF chain
- Improved channel performance and efficiency with addition of III-V devices
- Fully integrated beamformer channels demonstrated with integrated InP devices and Si control electronics
- >100mW Pout Tx channel, 4.5 dB NF Rx


Wafer-level heterogeneous integration



Integration schematic



InP/CMOS with DBI Process




DBI = Direct Bond Interconnect



DAHI InP/CMOS beamformer performance

Q-band InP/CMOS Tx Channel Layout





Channel dimensions: 3.0x0.6mm²


Q-band Tx Channel

- 57 HBTs, 1704 CMOS gates, 76 HICs
- P_{diss} ~ 1W
- 28-35 dB gain variation
- 5° RMS phase error at 40 GHz
- 20.3 dBm P_{sat}, 18.1 dBm P_{1dB}

Measured Gain and Phase Variation





Measured Tx Channel Power



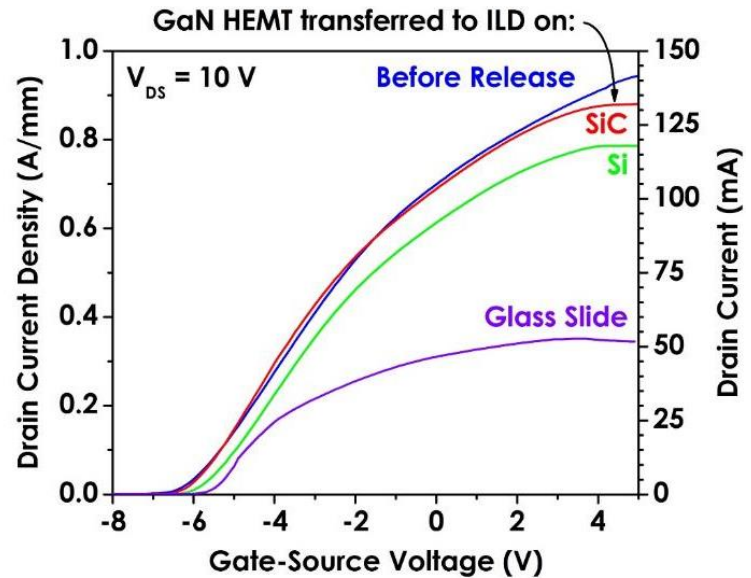


Advancing integration to the device level




Microtransfer Printing



Microtransfer Printing: GaN HEMT

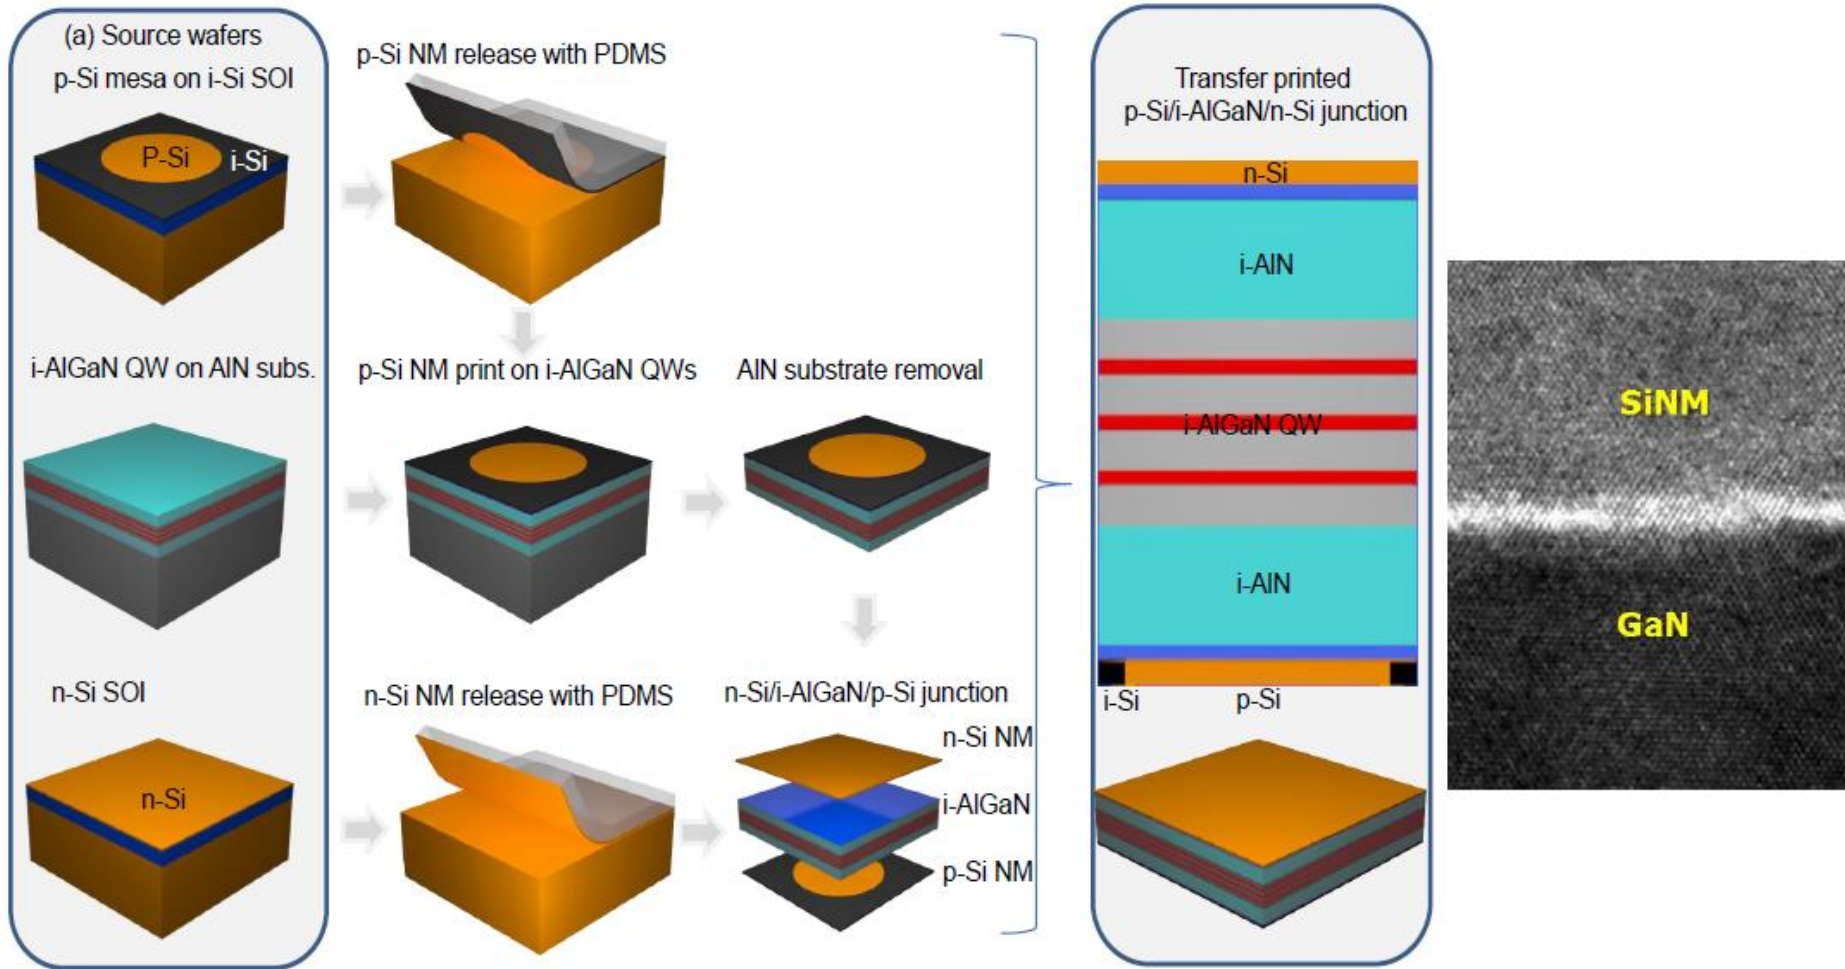


Challenges to managing thin-film strain



Source: Naval Research Laboratory


Transfer printing can stack nano-membrane layers to create junctions and devices




Sources: U. of Wisconsin, Michigan State U.



Common Heterogeneous Integration and IP Reuse Strategies: The next step in heterogeneous integration



CHIPS: modular design



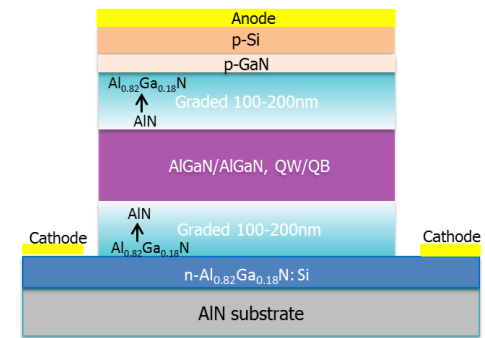

Integration technologies

Heterogeneous devices

DAHI: RF power






DAHI: wafer scale





Integration impact driven by **Standards** and **Modularity**



"Printed circuit board" invented by Paul Eisler.

Early PCB demo in a radio.

First HVM PCBs enable proximity fuze during WWII.

Patent to US Army for PCB assembly.

IPC (Institute for Printed Circuits) founded; standards follow.

Image: Intel Multi-layer PCB invented.

Surface Mount Technology on PCBs revolutionizes manufacturing.

HDI / Microvia technology enables further integration.


First package-on-package standard from JEDEC

DoD jump-start



1936 1941 1943 1956 1957 1960 1980s 1995 2006

PCB industry sees steady expansion with DoD origins, standardization, and technology development.



Global PCB Revenue (\$M)



Computers, smartphones, and tablet sales: 1975-2011



Moore's Law




Faith no Moore

Selected predictions for the end of Moore's Law

Cited reason:

■ Economic limits ■ Technical limits



Sources: Press reports; *The Economist*

Economist.com

Changes in silicon industry will be felt by compound semiconductors



Moore's Law INCLUDES Heterogeneous Integration

It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected. The availability of large functions, combined with functional design and construction, should allow the manufacturer of large systems to design and construct a considerable variety of equipment both rapidly and economically.

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 logic of integrated electronics is the logic 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 laser computers—each at least ten times smaller than a conventional computer—automatic controls for automobiles, and personal portable communications equipment. The electronic revolution needs only a decade to be feasible today.

One of the biggest potential lies in the production of large systems. In telephone communications, integrated circuits in digital filters will separate channels in complex equipment. Integrated circuits will also replace telephone circuits and produce data processing.

Computers will become powerful, and will be organized in completely different ways. For example, numerous kinds of integrated electronics may be distributed throughout the

area instead of being concentrated in a unit within. The improved reliability made possible by miniaturization will allow the construction of larger systems. Machines designed for this use in the future will have more and will have more speed.

Research and development

In the past, many of the integrated circuits that are referred to as microcircuits were in fact discrete components that result in a circuit when applied to the user as individual units. In fact, the first integrated circuit was developed in the late 1950s just as a miniature electronic component to carry out certain logic functions in a test instrument. Several approaches evolved, each using a different technique for individual component integration and semiconductor technology.

Each approach evolved rapidly and cost-effectively, but each had its own limitations. Many believe the way of the future is to be a combination of these approaches.

The advantages of semiconductor technology are being realized by applying each of these approaches to a variety of substrates. These advantages are being realized by applying each of these approaches to a variety of substrates. These advantages are being realized by applying each of these approaches to a variety of substrates.

Such approaches have varied and will continue to vary in the future.

Electronics, Volume 36, Number 8, April 19, 1965

The establishment

Integrated electronics is an established science. Its development is almost mandatory for the military system, since the reliability, size and weight required by some of these systems are not achievable only with integration. Study programs at Apple, for example, have demonstrated the reliability of integrated electronics by showing that complete circuit functions are as fast as those in the best individual transistors.

Microcomputers 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. Parameters of various sorts, especially the speed of processing, are being improved by digital techniques, an indication of integration because it cuts costs of both manufacturing and design.

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

Reliability concerns

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 often holds the same cost, and in many systems improved performance has been realized.

Integrated electronics will make electronic techniques more generally available throughout all of society, performing more functions that practice is not yet equipped by other techniques or not done at all. The principal advantages will be lower costs and greatly diminished size—results from a ready supply of low-cost functional packages.

For most applications, semiconductor integrated circuits will produce more. Semiconductor devices are the only transmittable solid state primary in an electronic 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 practice 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. The silicon will predominate at lower frequencies because of the technology which has already evolved around it and its wide application in 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 silicon chip. For example, a single, three-foot per centimeter is usually necessary, proportional to the number of components, the cost of the


chip and the price of semiconductor in the equivalent package containing more components. The as components are added, the cost per component falls more than proportionally for the complexity, tending to make the cost per component fall more than proportionally for the complexity.

There is a minimum cost per component for the cost of the technology. At present, it is reached when 30,000 units are used per circuit. But the minimum is rising, while the cost per component falls, over a period of about five years, a plot of cost suggests that it may reach one cent per component for a cost of about 1,000 components per circuit (including lead and functions can be produced in moderate quantities, if the manufacturing cost per component can be kept only a half of the present cost.

The complexity for minimum component cost, around a size of roughly a factor of two per a silicon chip. Certainly over the short term can be expected to continue. If not to increase, it may even decrease. The rate of increase is a bit more rapid than that in the past 10 years. The trend by 1975, 100,000 components per integrated circuit for minimum cost is not far off.

Two- and three-dimensional

With the dimensional reduction already being realized in integrated circuits, reduced high performance packages will be built in custom to thousands of an inch square.



Electronics, Volume 36, Number 8, April 19, 1965




A two-and-a-half square inch area contains several thousands of resistors 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-tenth a square inch.

On the silicon wafer currently used, used by a wafer or more in diameter, there is ample room for such a structure if the components can be closely packed with no space wasted for interconnective purposes. This is possible, since efforts to achieve a level of complexity above the present available techniques are already underway using reduced dimensions and patterns imposed by automatic means. Such a density of components can be achieved by present optical techniques and does not require the more exact 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 do not exceed the cost of the semiconductor structure itself that there is no incentive to improve yields, but they can be raised to high levels.



Electronics, Volume 36, Number 8, April 19, 1965

It is economically justified. The itemization of equipment, as is not even necessary to do any fundamental research in to replace present practices. Only the engineering effort is needed.

In the early days of such markets, when yields were extremely low, the cost of integrated circuits was high. Today, yields are high, and the cost of integrated circuits is low. The same pattern will make integrated electronics economical, if other conditions make such an area desirable.

Power problems

Power will be a problem in the future. The heat generated by tens of thousands of components in a single item chip? If we could shrink the volume of a standard high-speed digital computer to the size of a single item chip, we would have a great deal of power to dissipate. But it would happen with integrated circuits. Since integrated electronic structures are two-dimensional, they have a surface area for heat conduction close to each corner of heat generation. Heat dissipation is not a problem in integrated structures unless the heat sink is not connected with the system. As long as a function is confined to a small area on a multi-layered structure which may be driven in a distributed fashion, the heat dissipation is not a problem. It is possible to operate the structure at higher speed for the same power per unit area.

Cost of packaging

Clearly, we will be able to build such component-united equipment. Now, we ask under what circumstances we should do it. The cost of making a particular system function must be minimized. To do so, we could simulate the engineering over and over identical items, or a single function for the assembly of large functions or that as appropriate and it may be done by a particular army. Perhaps newly devised design means from procedures could be used to speed up the engineering.

It may prove to be more economical to build large



G. 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.



What is CHIPS?

CHIPS will develop design tools, integration standards, and IP blocks required to demonstrate modular electronic systems that can leverage the best of DoD and commercial designs and technology.

Today – Monolithic






Image: Intel




Tomorrow – Modular





CHIPS enables rapid integration of functional blocks at the chiplet level

Custom chiplets



Commercial chiplets




COMM



RADAR EW




SIGINT





Adaptive filter	SerDes	SerDes
Beam forming	Beam forming	Adaptive filter
QR Decomp.	QR Decomp.	QR Decomp.



CHIPS developing interface standard



CHIPS Target



CHIPS Program Interface Standard Metrics	
Data rate	10 Gpbs
Energy efficiency	< 1 pJ/bit
Latency	< 5 ns
Bandwidth density	> 1000 Gbps/mm


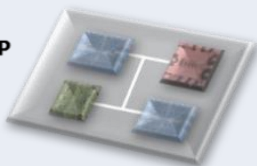
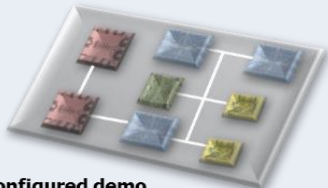
- 28nm SOI, Single-ended, Al on Si
- 28nm, ground-ref., single-ended, organic PCB
- 45nm SOI, differential, Cu on Si
- 32nm, differential, 32AWG cable
- EMIB
- 14nm SERDES, PCB
- 14nm HBM

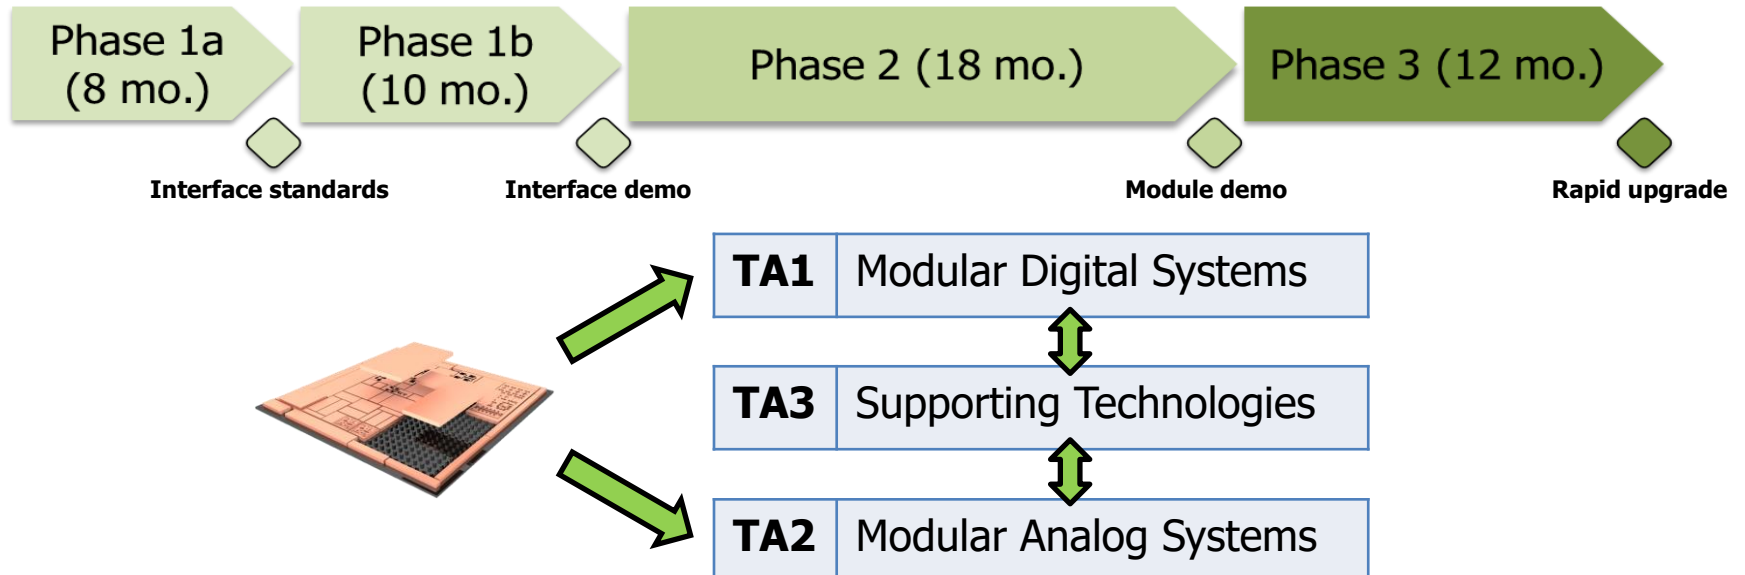
Sources:
 1. 2016 JSSC, Dehlaghi
 2. 2013 JSSC, Poulton
 3. 2012 JSSC, Dickson
 4. 2013 JSSC, Mansuri
 5. 2016 ECTC, Mahajan

CHIPS interface is one of many possible routes for efficient interdie communications



CHIPS program: structure and timing


PHASE 1	PHASE 2	PHASE 3
Interface and IP Block Demo	Module Demo with IP Blocks	Rapid Module Upgrade
 Integration platform Interface demo	 Full system IP reuse demo	 Reconfigured demo



Seeking CHIPS collaboration to help drive a common interface



CHIPS: August 2017 Kickoff



CHIPS Team

Designs

- Boeing
- Intel
- Lockheed Martin
- Northrop Grumman
- Univ. of Michigan







Chiplets

- Intrinsix
- Jariet
- Micron
- North Carolina State
- Synopsys

Tools

- Cadence
- Georgia Tech

CHIPS Approach Modularity Standards



CHIPS Results Fast Cost-Effective Best-in-Class



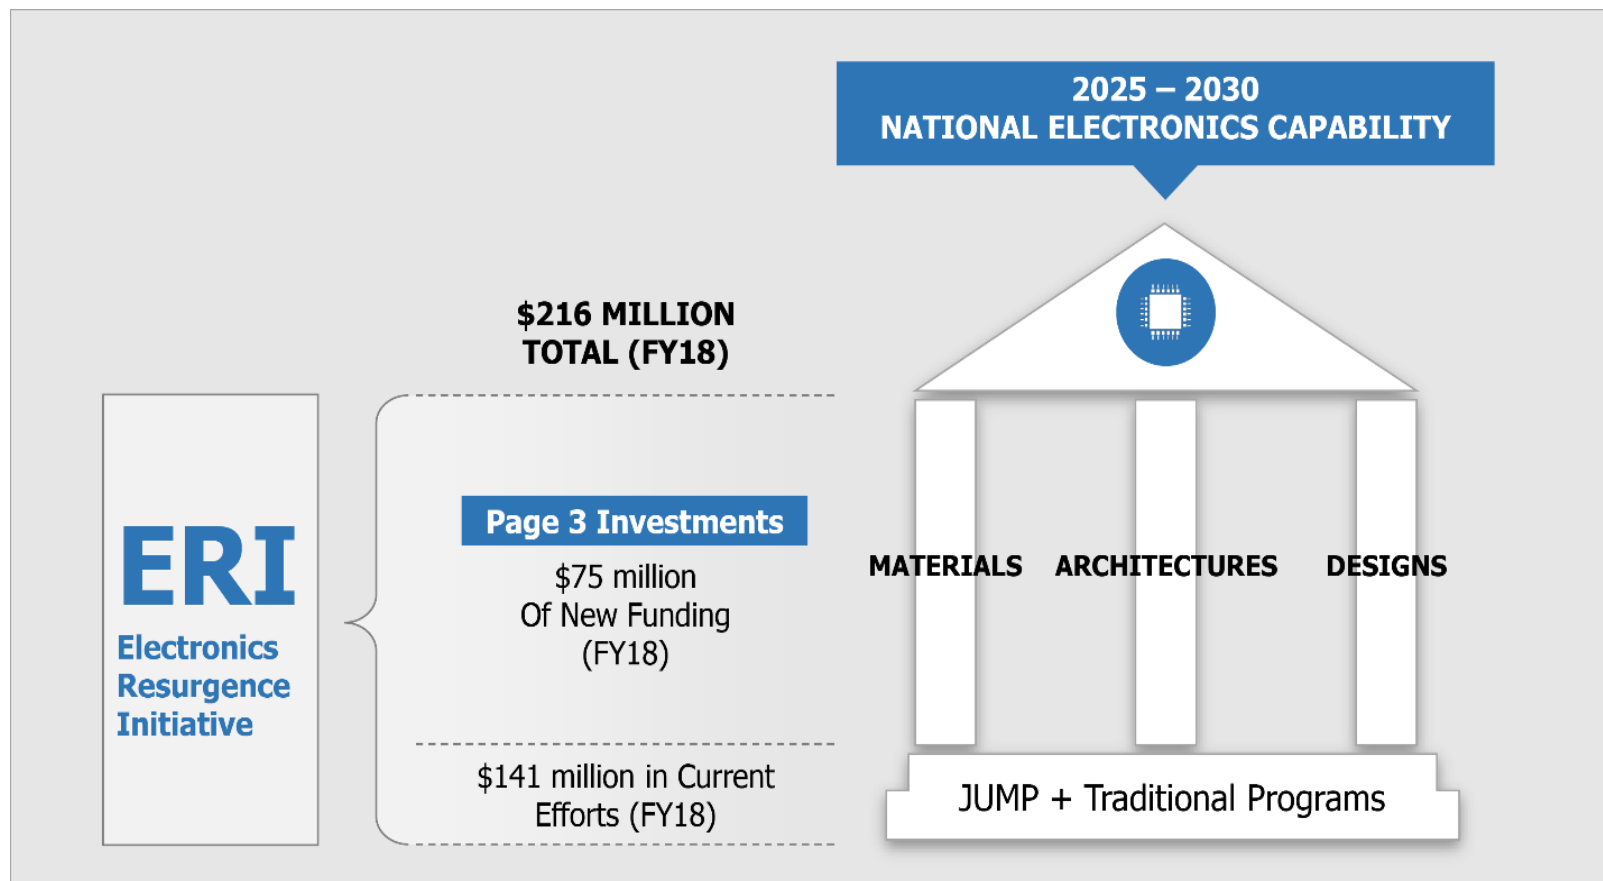


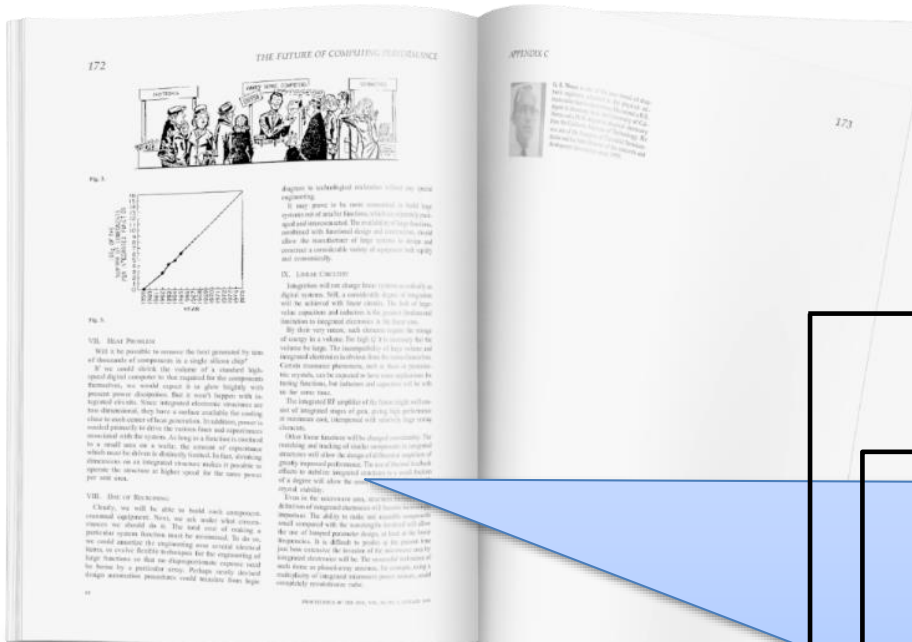
What's next?



Electronics Resurgence Initiative

- ERI—creating an electronics capability that will provide a foundational contribution to national security
- Three thrust areas: Materials and Integration, Architectures, Designs





Electronics, April 19, 1965: Cramming More Components onto Integrated Circuits; Gordon Moore

P.3

VIII. 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 systems out of smaller functions, which are separately packaged** and interconnected. The availability of large functions, combined with functional design and construction, should allow the manufacturer of large systems to design and construct a considerable variety of equipment both rapidly and economically.

Architecture

Maximizing specialized functions

Design

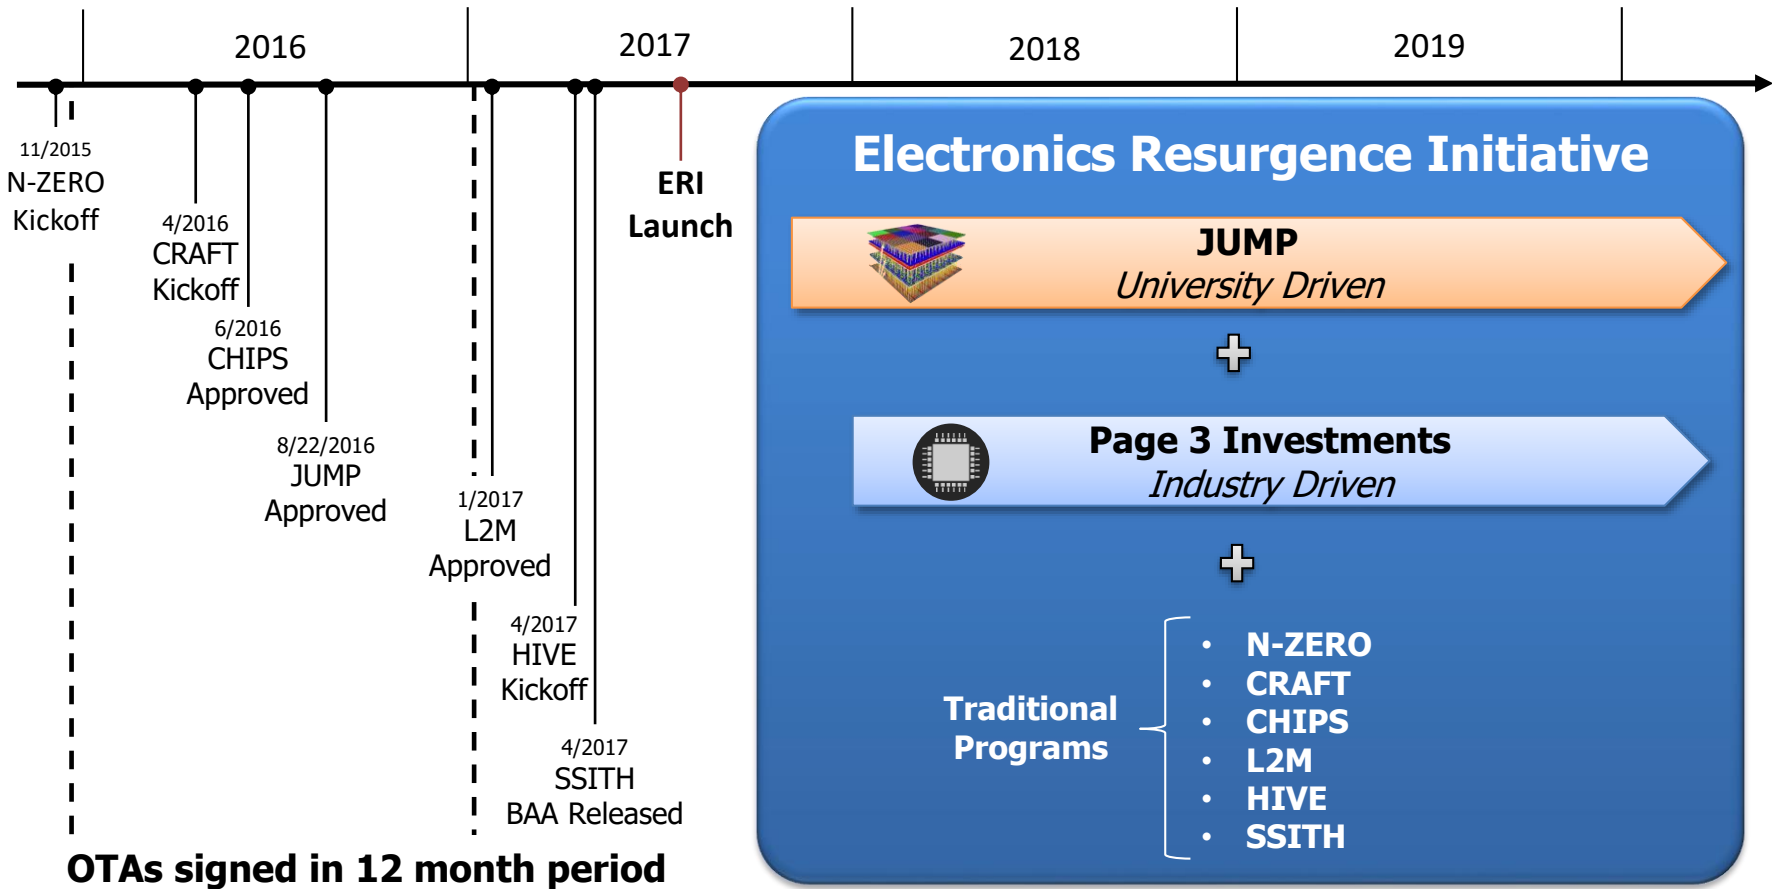
Quickly enabling specialization

Materials & Integration

Adding separately packaged novel materials and using integration to provide specialized computing



Recent DARPA investments and momentum



Electronics Resurgence Initiative

JUMP
University Driven

+

Page 3 Investments
Industry Driven

+

Traditional Programs

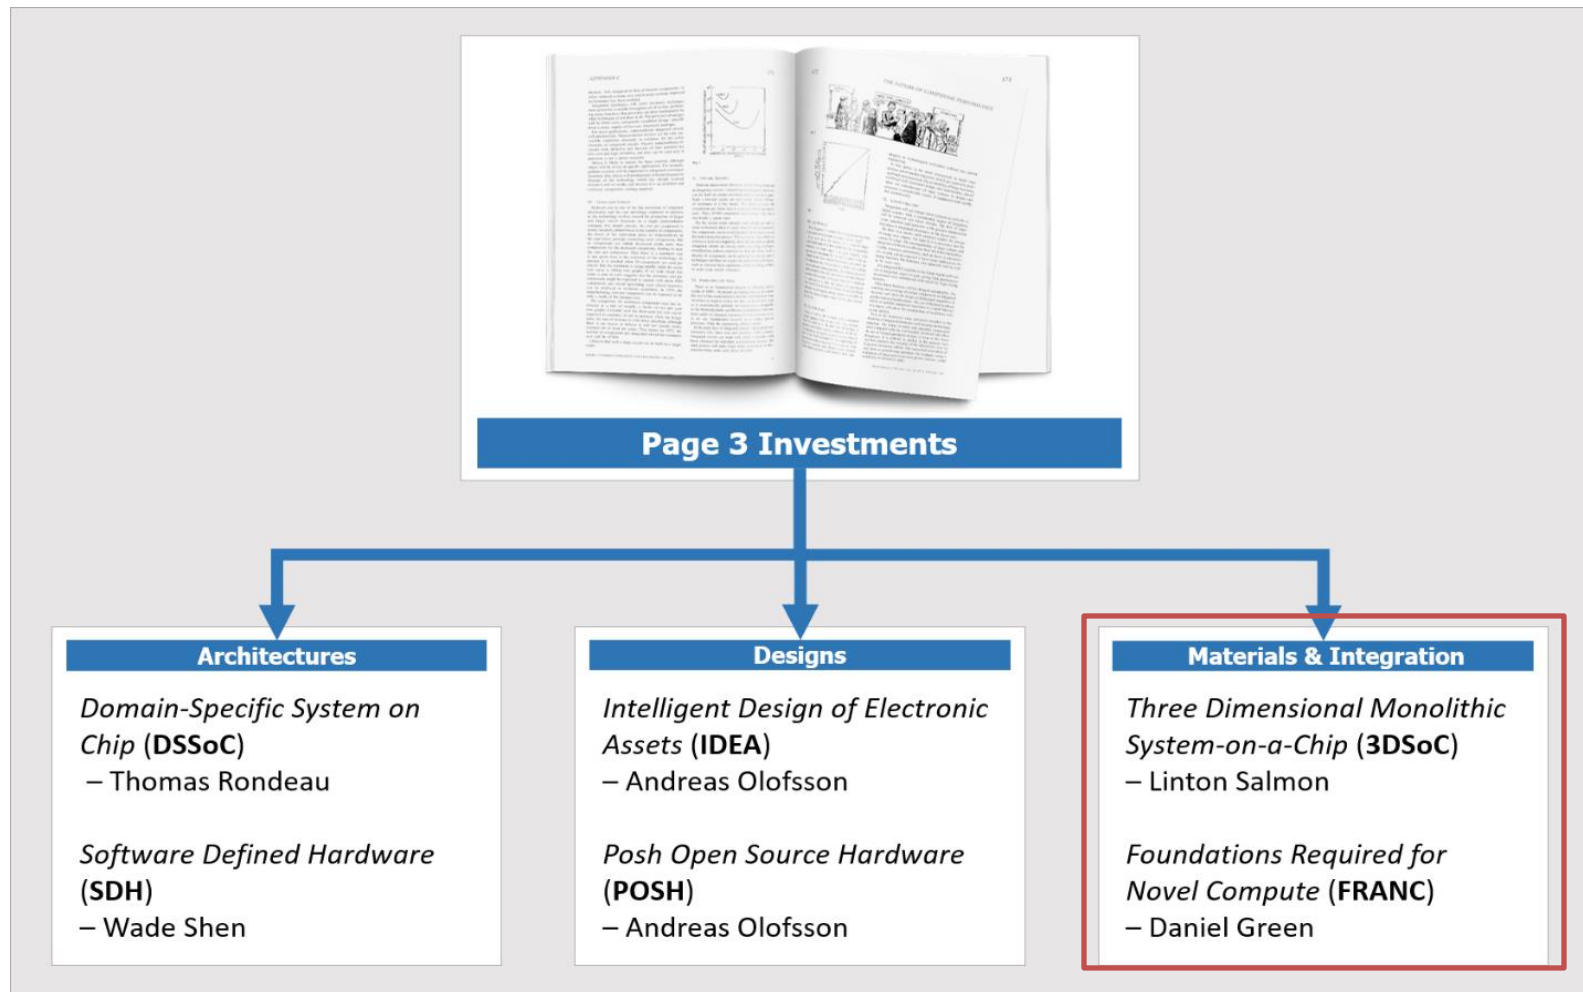
- N-ZERO
- CRAFT
- CHIPS
- L2M
- HIVE
- SSITH

Intel Qualcomm Rambus XILINX Micron

Cadence Keysight Technologies NVIDIA Flexlogix Technologies, Inc



Investments inspired by Moore's "Page 3"




Heterogeneous Integration directly part of Materials and Integration thrust and indirectly part of Architectures and Design thrusts as well




The Challenge: Overcoming Memory Bottleneck

Current von Neumann architecture spends more time moving data than processing it


Neural Programmer (LSTM)



ResNet-152 (CNN)



Alex Net (CNN)



■ Compute ■ Memory

Data from S. Mitra of Stanford

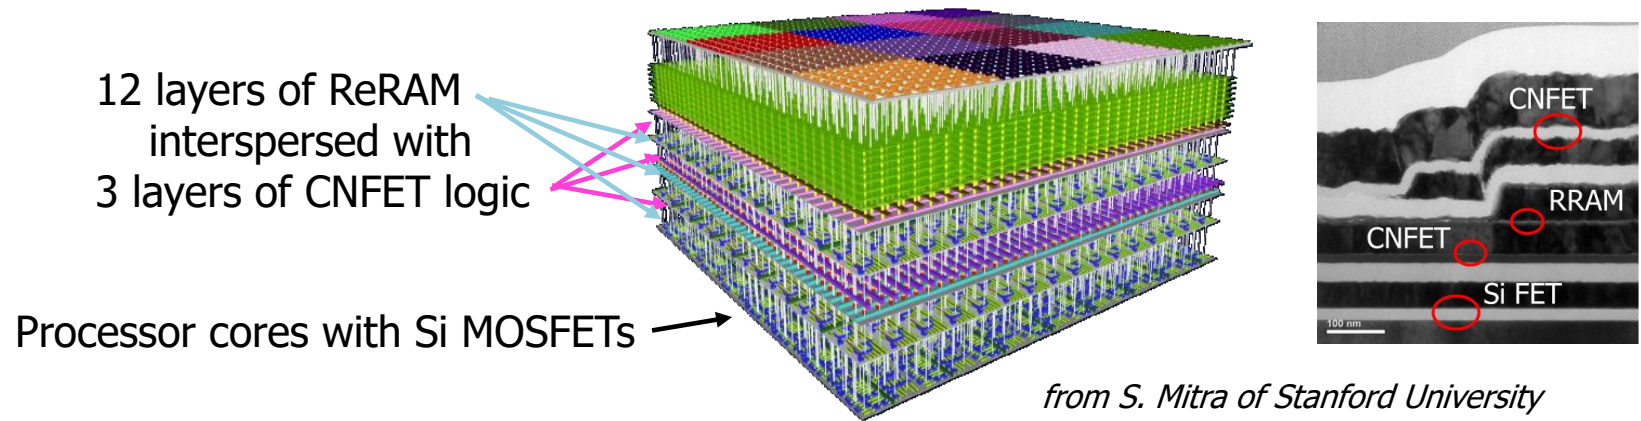
Data for 7nm instantiation of a state-of-the-art Machine Learning accelerator

Accelerators don't help (enough) if using the same architecture



An Integrated, Monolithic SoC (3DSoC) Solution

An example of an integrated flow that fabricates 3D logic and memory on a single die



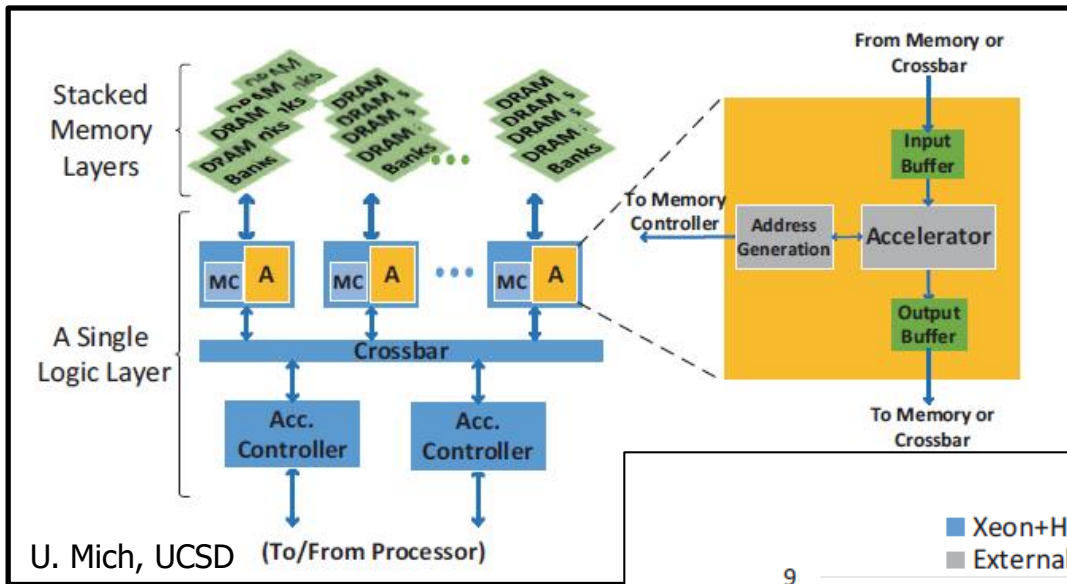
Note: This is an example only. Other technical approaches are expected.

Critical characteristics for a monolithic solution

- Must permit new architectures that leverage fast, configurable access to non-volatile main memory
- Stackable 3D logic and memory functions that allow new architectures
 - Low temperature formation
 - Logic AND memory
 - High density of memory – at least 4GB (Giga-Byte)/die
- Possible to fabricate in existing domestic, commercial, high-yielding infrastructure
 - 90nm on 200mm wafers
 - High yield on large SoCs




FRANC leverages demonstrated benefits of beyond von Neumann topologies



Near-memory processing provides dramatic performance and energy improvements

U. Mich, UCSD (To/From Processor)

- Xeon+HMC—Quad Xeon and Hybrid Memory Cube
- NM Atom—16 Intel Atom Cores
- External Acc—16 external accelerators, SerDES linked
- NM Acc—Near memory accelerator



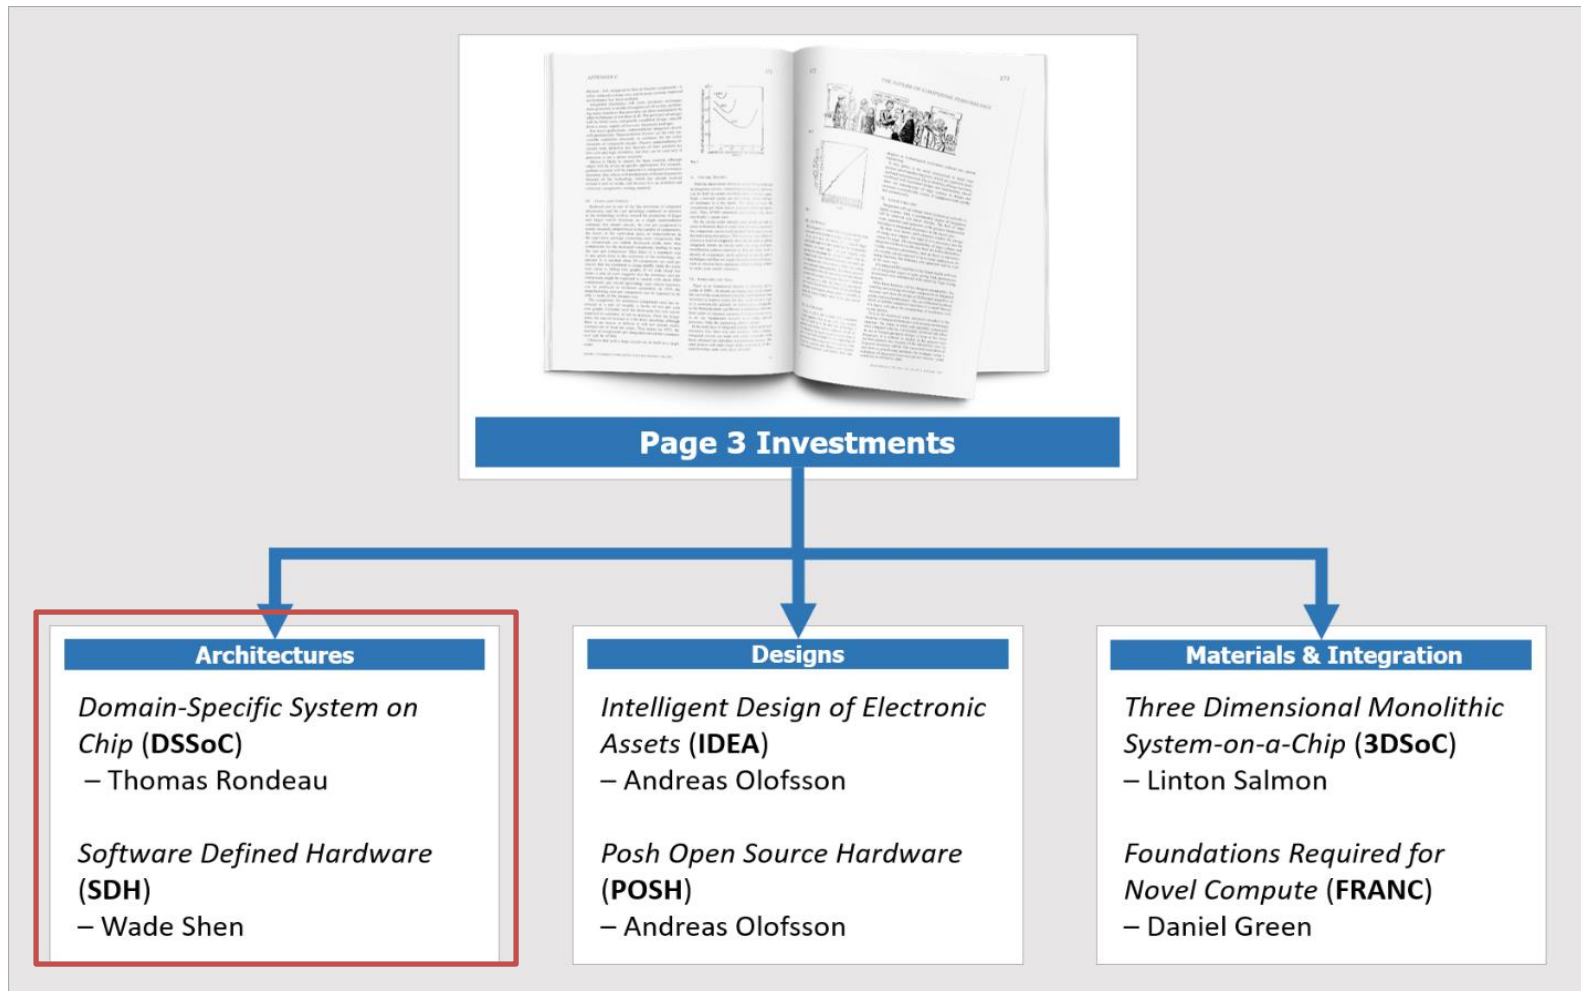
Performance Comparison (higher better)

Source: S.F. Yitbarek, et al., DATE 2016.

Key result: near-memory processing provides dramatic performance improvements



Investments inspired by Moore's "Page 3"




Heterogeneous Integration directly part of Materials and Integration thrust and also indirectly part of Architectures and Design thrusts as well

Build new processors that solve the significant computing needs of today's and tomorrow's applications.

1: Domain Specific System on Chip (DSSoC)



Streaming Data is latency sensitive, small but many work loads



DSSoC

2: Software Defined Hardware (SDH)

Big Data is efficiency sensitive, large and repeatable work loads




SDH



DSSoC program will...

- Create a development ecosystem that takes advantage of the specialized hardware with no added burden to the programmer
- Design an intelligent scheduler for efficient data movement between DSSoC processor elements
- Build a DSSoC of advanced, heterogeneous processors and accelerators for software radio



Examples of Processor Elements (PE)



DSSoC will enable rapid development of multi-application systems through a single programmable device



Software-defined hardware (SDH)

High-level program



Dynamic HW/SW compilers for high-level languages (TA2)




1. Generate optimal configuration based on static analysis code
2. Generates optimal code
3. Re-optimize machine code and processor configuration based on runtime data

Code₁

Code₂

Code₃

Code_N



Time



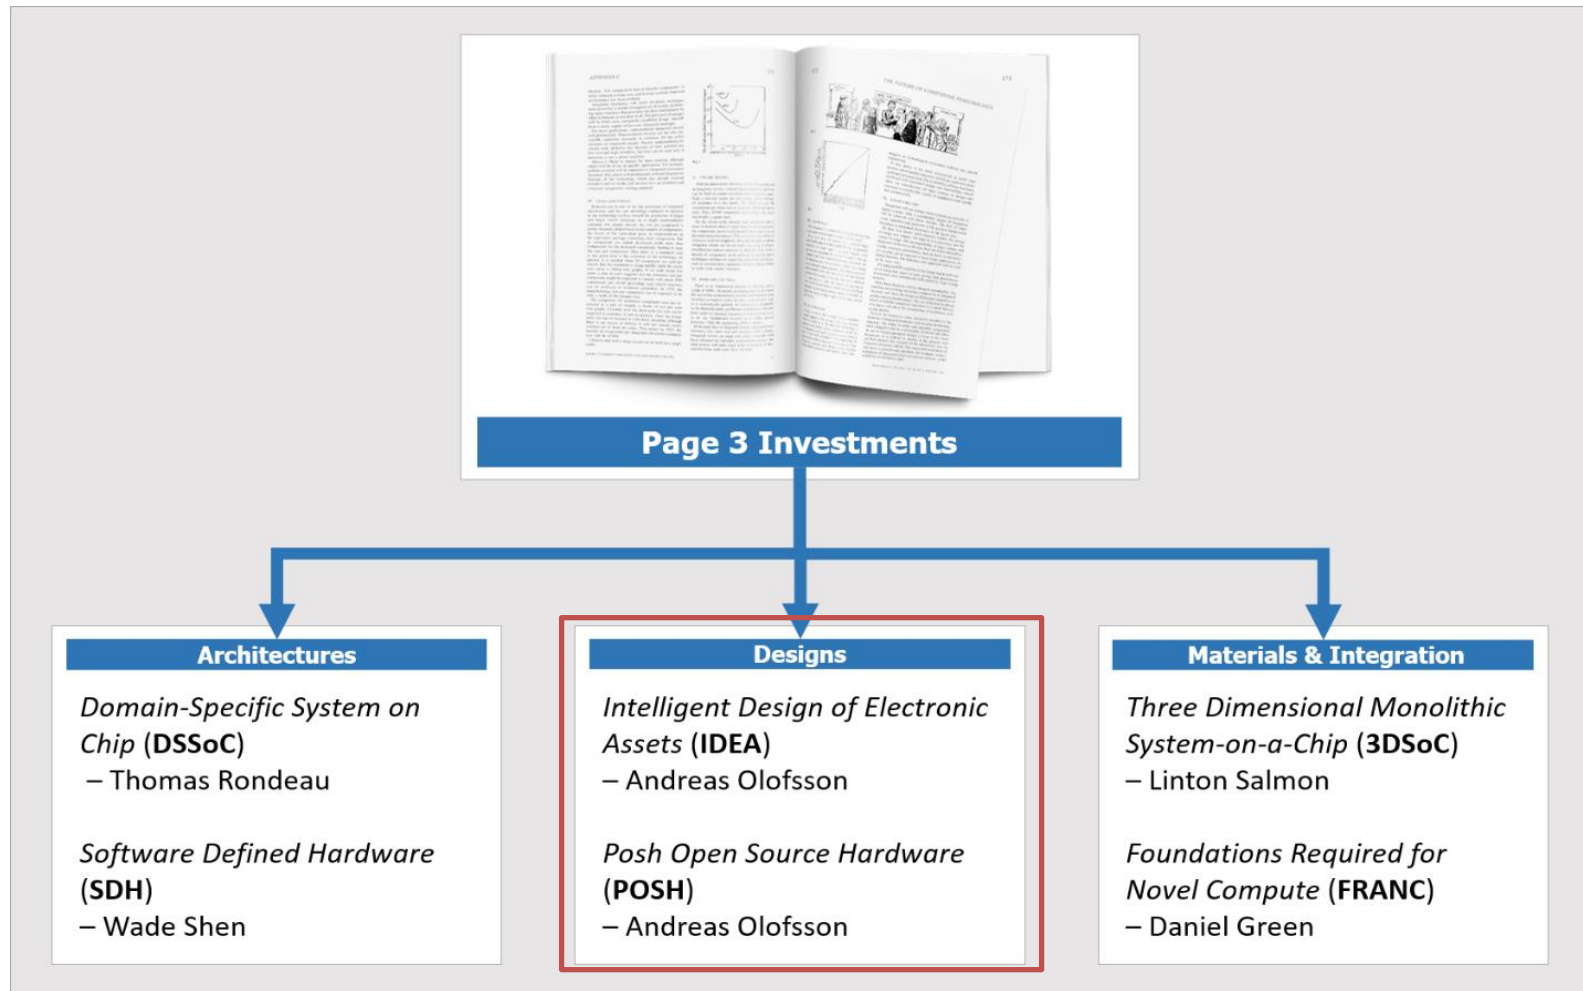
Reconfigurable processors (TA1)

Properties:

1. Reconfiguration times: 300 - 1,000 ns
2. Re-allocatable compute resources – i.e. ALUs for address computation or math
3. Re-allocatable memory resources – i.e. cache/register configuration to match data
4. Malleable external memory access – i.e. reconfigurable memory controller



Investments inspired by Moore's "Page 3"




Heterogeneous Integration directly part of Materials and Integration thrust and also indirectly part of Architectures and Design thrusts as well



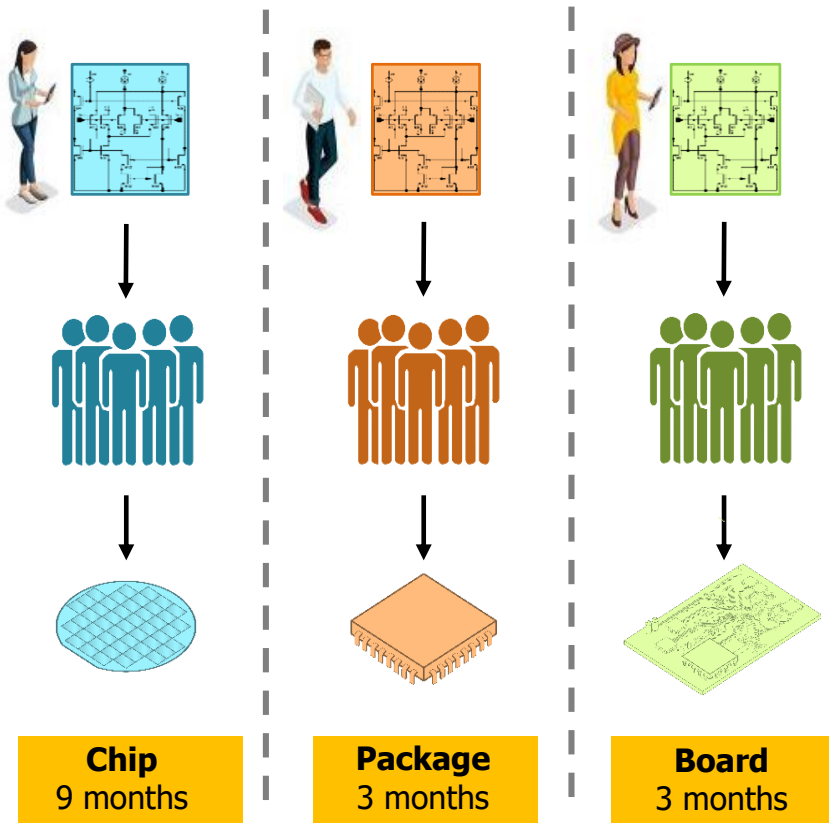
IDEA: High level objectives

IDEA will completely automate the layout of electrical circuits and systems



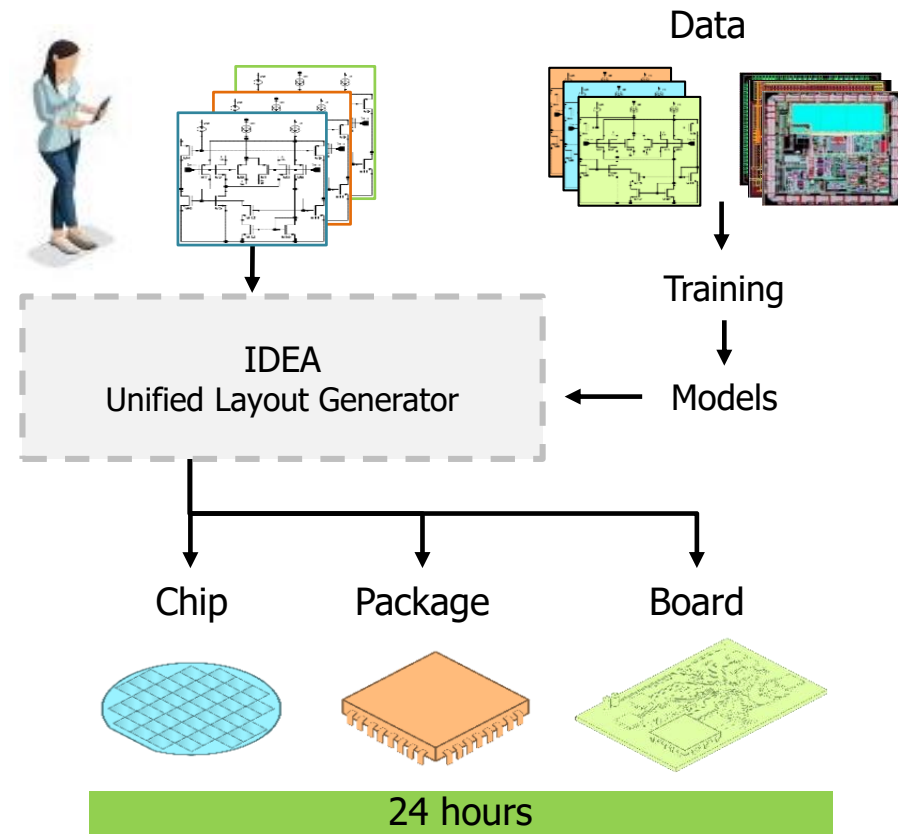
Sources: Raspberry Pi

Today



- Knowledge embedded in humans
- Limited knowledge reuse
- Reliance on scarce resources


IDEA



- Knowledge embedded in software
- 100% automation
- 24 hour turnaround



Future of heterogeneous integration



Requires a lot of pieces coming together!



www.darpa.mil



GIGAOM.com