NDIA Electronics Division Meeting with Serge Leef

Serge Leef, DARPA/MTO

Virtual Event

July 20, 2021

Mr. Serge Leef

Microsystems Technology Office (MTO)

Program Manager

Mr. Serge Leef joined DARPA in August 2018 as a program manager in the Microsystems Technology Office (MTO). His research interests include computer architecture, chip design tools, simulation, synthesis, semiconductor intellectual property (IP), cyber-physical modeling, distributed systems, secure design flows, and supply chain management. He is also interested in the facilitation of startup ecosystems and business aspects of technology.

Serge Leef Introduction

source: asu.edu

BS & **MS** in *Electrical Engineering & Computer Science*

- Focus: Simulation, Synthesis, Test Generation

source: intel.com

4 years @ **Intel**

- Software Engineer ⇨ Senior CAD Development Engineer
  - RTL Simulation, Behavioral Synthesis, Auto Place & Route, Silicon Compilation, Layout Synthesis

source: microchip.com

2 years @ **Microchip**

- CAD/CAE Group Manager
  - Auto-routing, RTL Simulation, Floorplanning, Place & Route

source: Wikipedia.org

3 years @ **Silicon Graphics**

- Corporate Manager, Design Automation
  - Cycle Simulation, Logic Synthesis, Co-verification

source: Wikipedia.org

28 years @ **Mentor Graphics** (now part of Siemens AG)

- Engineering Manager ⇨ Engineering Director ⇨ GM ⇨ VP
  - Simulation, System Level Engineering, New Ventures

3 years… @ **DARPA** / Microsystems Technology Office

- Program Manager: Design Automation, Security, Trusted Microelectronics, Structured Innovation

DASH (Design Automation & Secure Hardware) Portfolio

Security Portfolio

- **SHIELD** – chiplet containing hardware RoT, key store, RF, anti-tamper
- **AISS** – synthesis of secure SoCs (architecture to RTL)
  - SHINE – security-aware high-level synthesis (algorithm to architecture)
  - ARCHS – security design rule checking at RTL, gate and layout
  - SCATE – side channel and fault injection attack simulator
  - DEMI* – Trojan detection using formal, static and dynamic analysis methods
- **SPIRAL*** – Mitigation of vulnerabilities during physical design (RTL to GDSII)
  - ADVAL* – Detection of vulnerabilities to optical probing & fault injection

EDA Portfolio

- **IDEA** – open-source EDA (analog synthesis, APR, IP integration)
  - COPILOT* – cloud scaling of P&R to performance gains
- **POSH** – open-source IP including CPUs, protocol engines, and other peripherals
- **RTML** – ML models to hardware accelerator synthesis
  - RTML-Plus* – system ASIC flow for task-parallel workloads
- **CHEETA*** – next-gen cloud-scaled digital verification environment
  - Ditto - AI based generation of Reduced Order Models (digital, analog, system)
  - ESANA* – simulation acceleration HPC based on speculative parallelism
  - LILACS* – raising abstraction of all SoC components to improve simulation

Domestic Microelectronics

- **HI³** – domestic foundry for the US
- **SAHARA** – FPGA to Secure Structured ASIC transformation

Miscellaneous

- Toolbox – favorable EDA & IP access agreements for DARPA performers
- O-DRIVE* – distributed system simulation and optimization
- STRIVE* – dual-use chip design innovation accelerators / ventures

*Concept under development

Trends

Market-Driven Eras of Electronic Innovation

Packaged IC

Military & Aerospace

Source: Military.com

Microprocessor

Desktop Computer

Source: computinghistory.org.uk

Digital Signal Processing

Communications

Source: wired.com

Multiprocessing Gigabit, Non-volatile Memory

Consumer

Source: cnet.com

Systems

Defense
Medical
Industrial
Automotive

Sources: Mentor, WD, Canon

Computational Continuum

Cloud
5G
Edge

Vertical Applications

1970 **1980** **1990** **2000** **2010** **2020** **2030**

Chips are Evolving into Heterogeneous SoC Platforms

EDA has been riding on the coat-tails of Moore's Law

- Productivity improvements have largely been the result improved CPU and memory performance
- Innovations in design automation have been largely linear and based on functional aggregation
- Relative simulation speeds are largely unchanged
- Logic synthesis was arguably the most transformative innovation (1988)

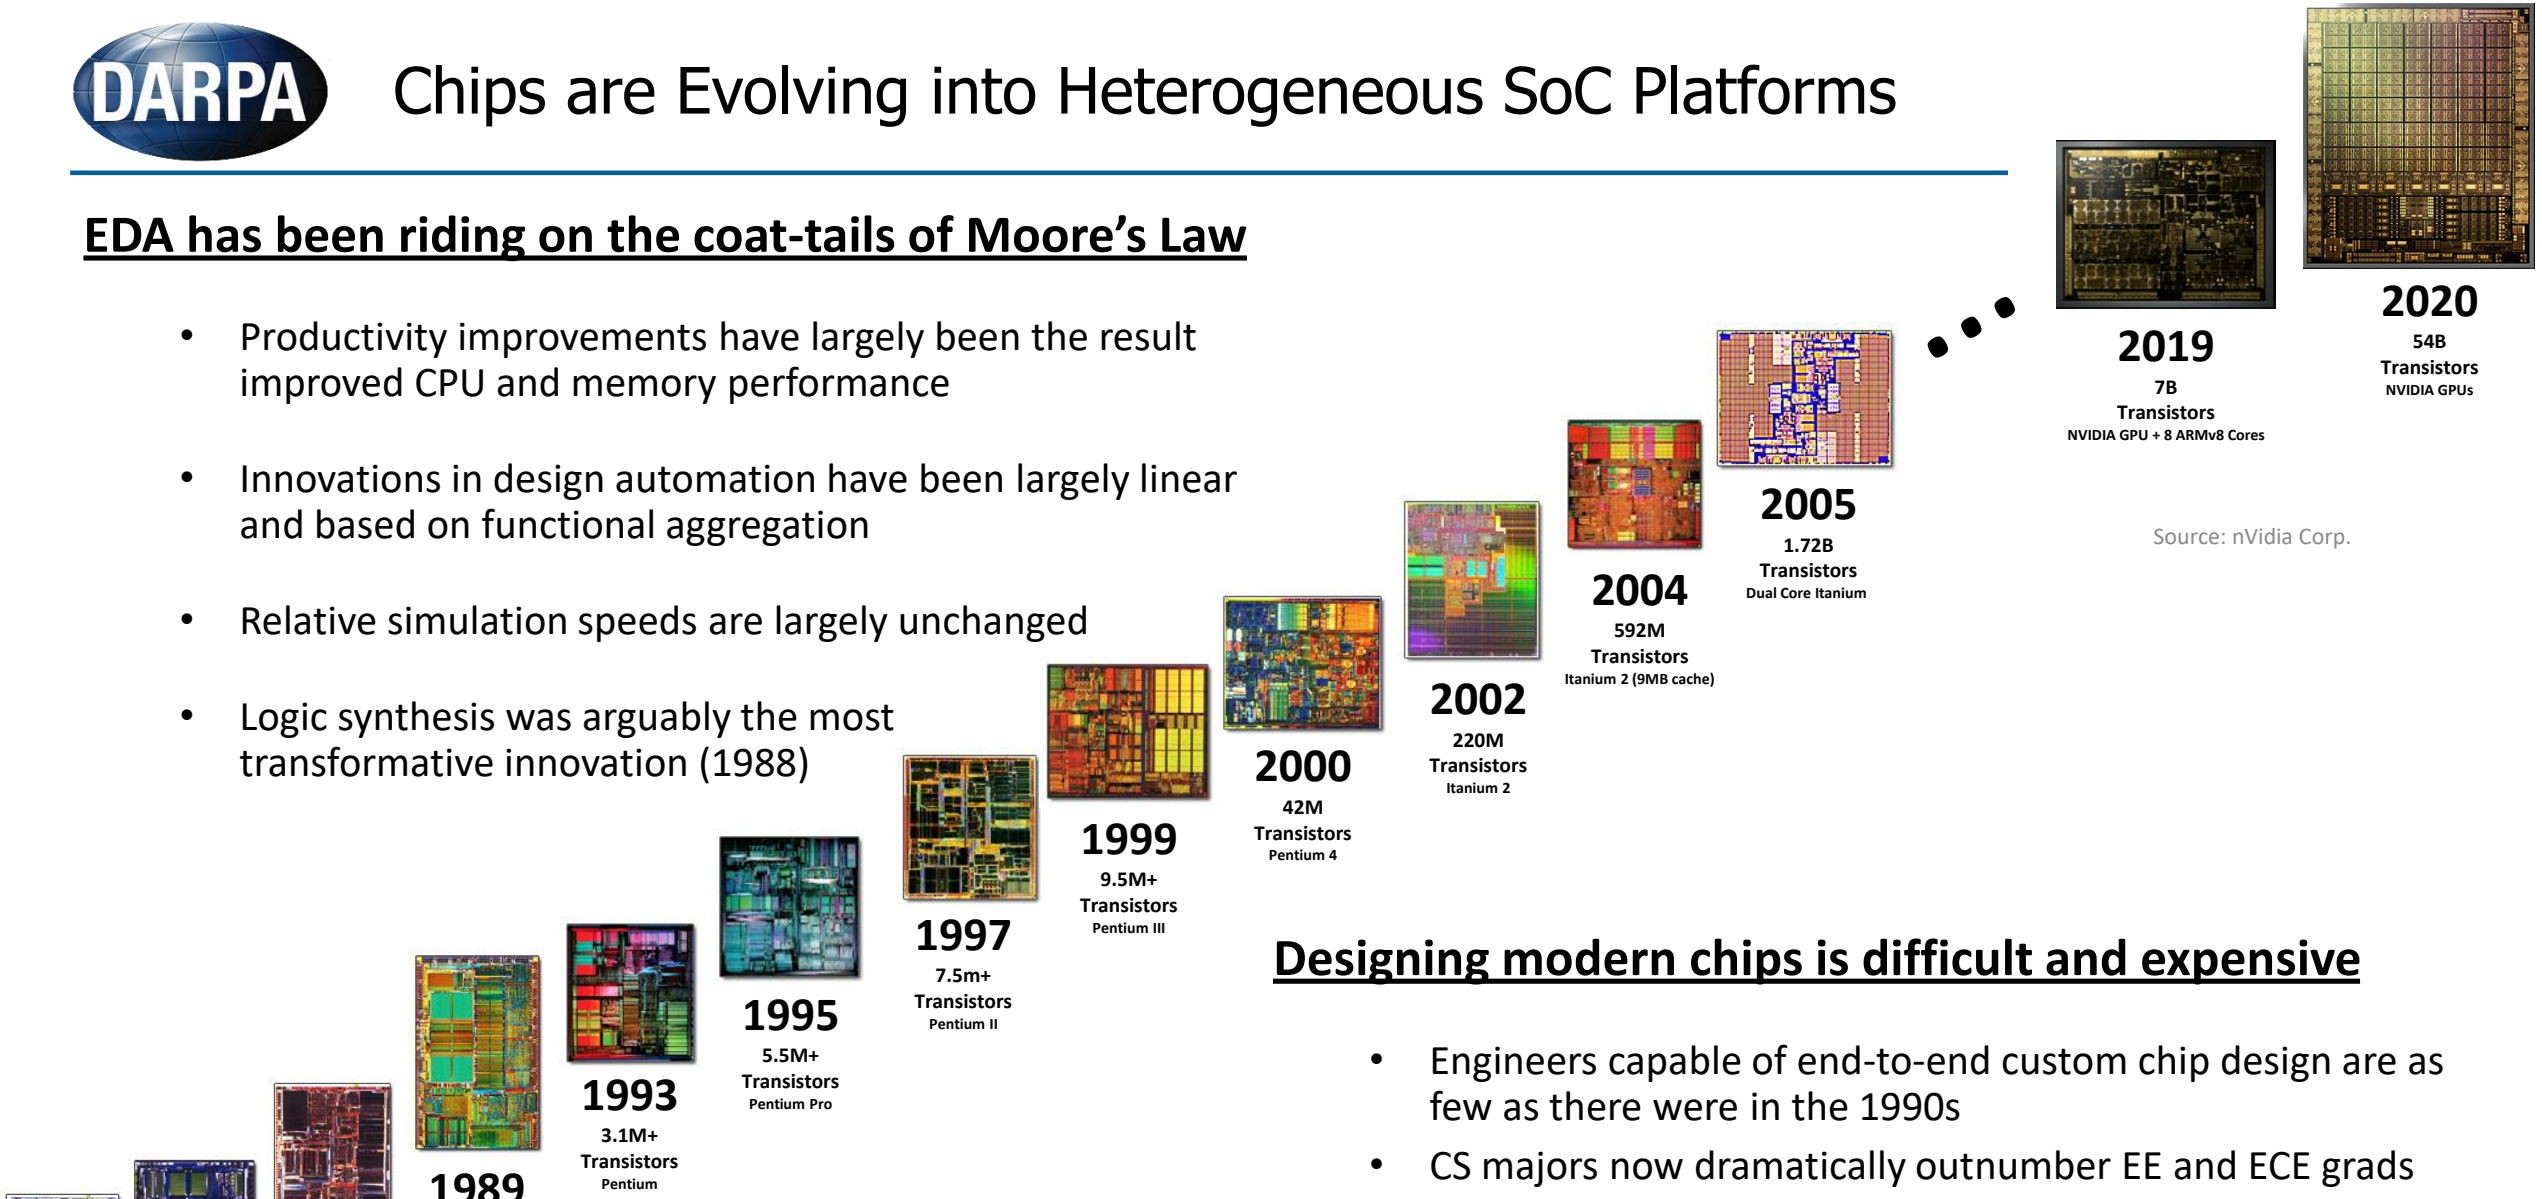

Source: Intel Corp.

Source: nVidia Corp.

Designing modern chips is difficult and expensive

- Engineers capable of end-to-end custom chip design are as few as there were in the 1990s
- CS majors now dramatically outnumber EE and ECE grads
- Few places in the world where SOTA chips can be built
- @ 5nm design cost is projected to be $500M+

Current Methodologies are Severely Strained

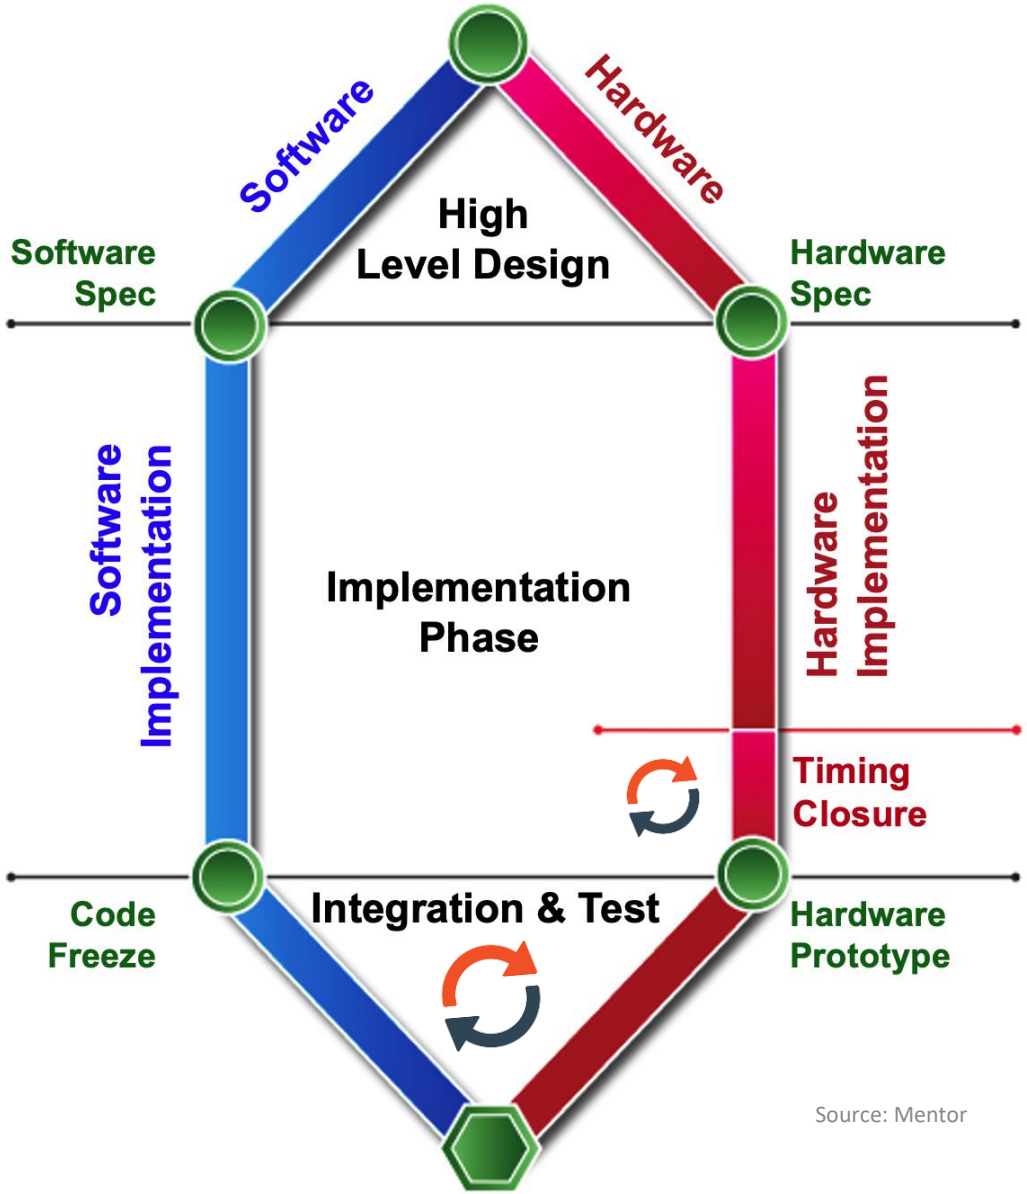

Source: Mentor

- Unpredictable, iterative loop during system integration/test phase
- Poor partitioning decisions at the front end of the process are impossible to overcome during the design
- Functional verification is strained by even today's designs; how to verify multi-discipline systems with billions of gates?
- Even though software is a key, growing system component it is only partially included in hardware verification phases due to insufficient simulation speed
- Interaction with outside world through sensors and actuators is rarely an integral part of the flow

Chip Designs Are Expensive and Getting More Expensive

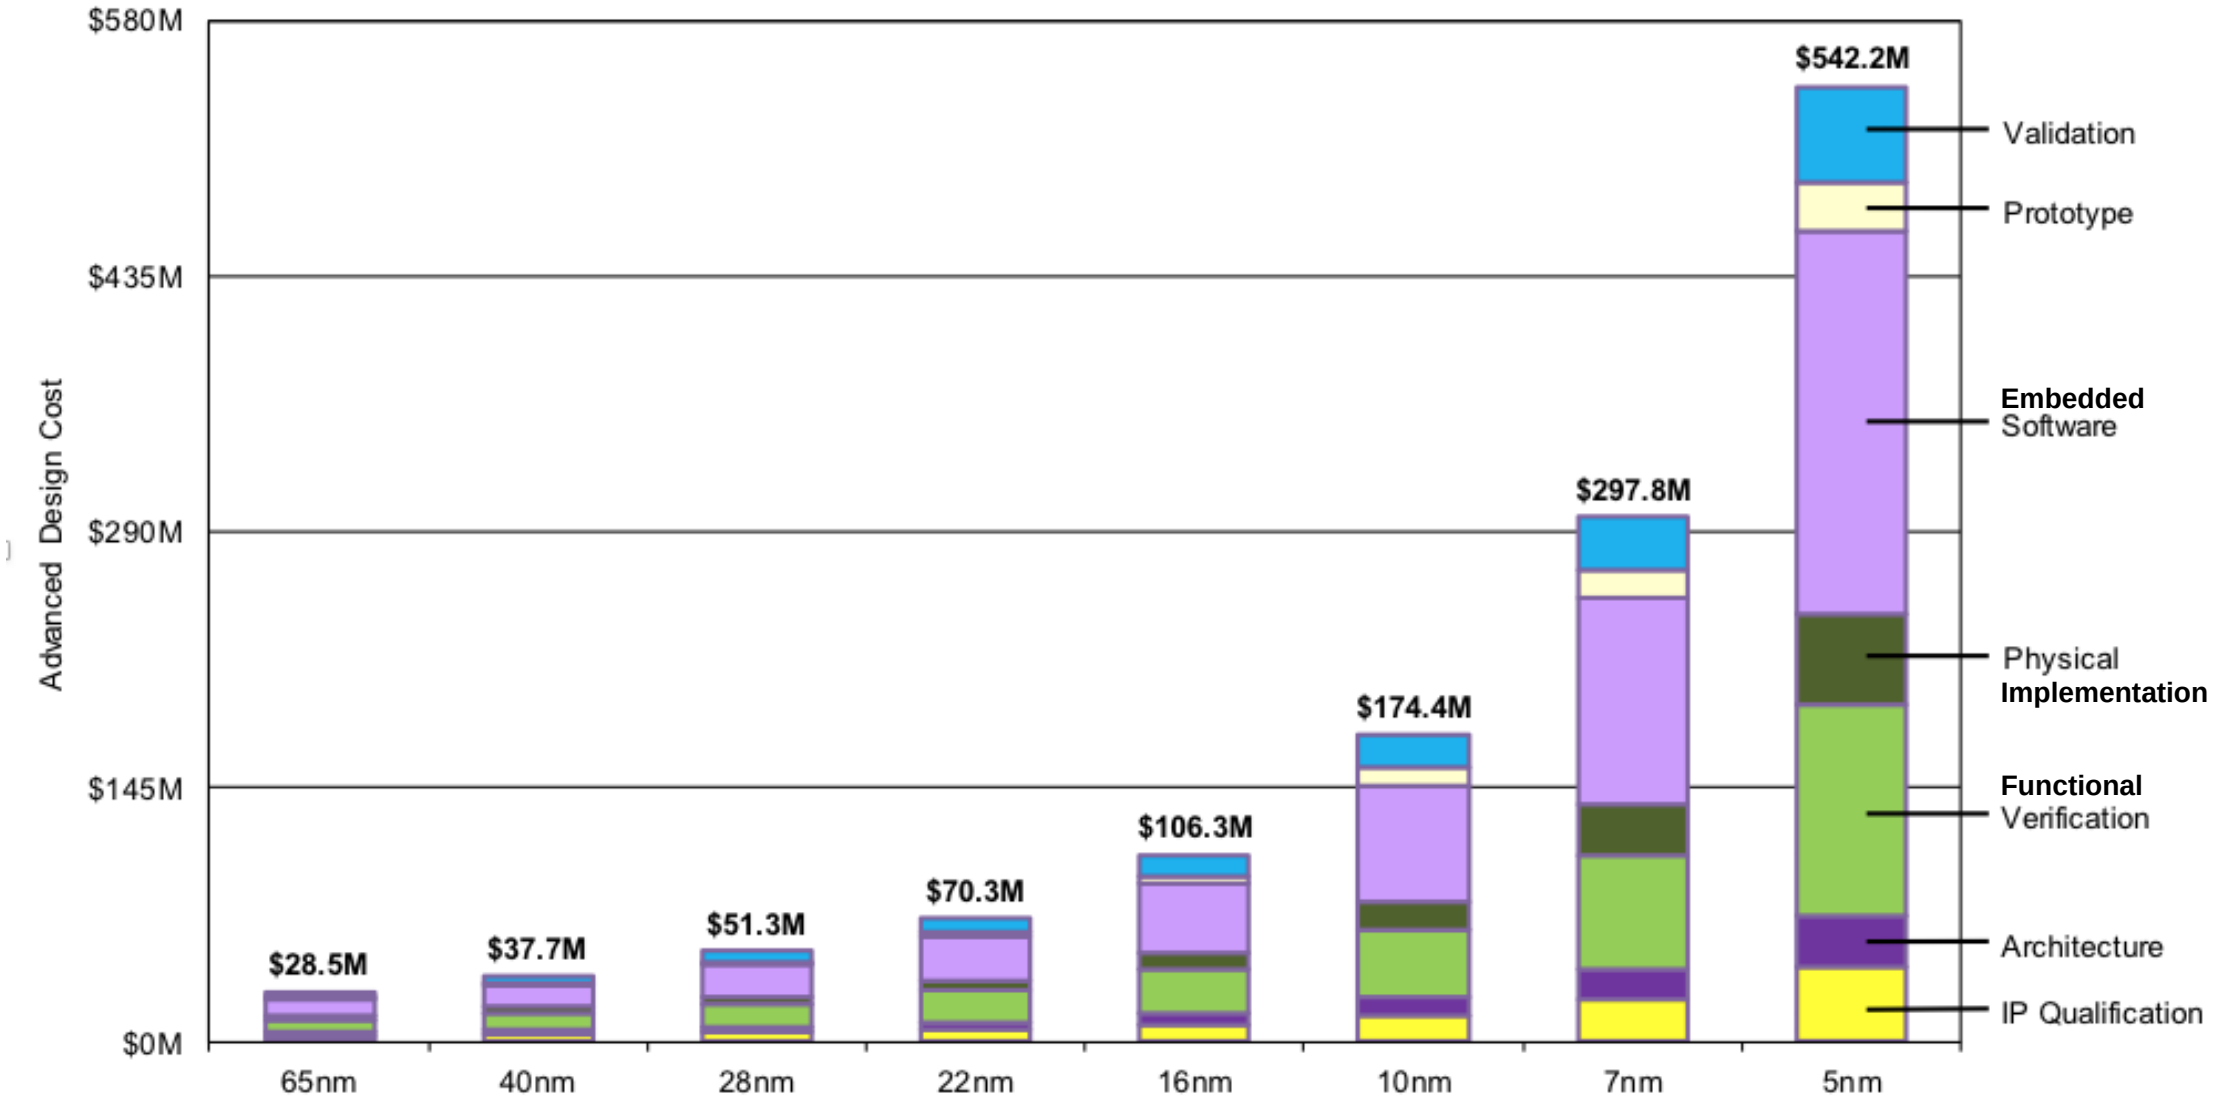

Source: International Business Strategies (IBS); https://www.extremetech.com/computing/272096-3nm-process-node

High Design Costs Make Chip Startups are Driving VCs Away

Conventional Model

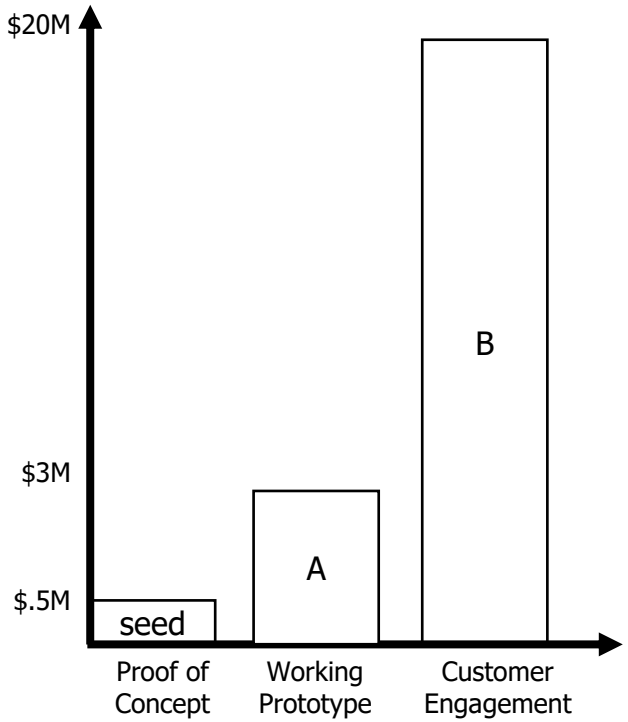

Seed – "A" Activities

- Development of working prototype
- Business model development
- Pricing, packaging, sales strategies
- Market definition & business development

Chip Design Startup

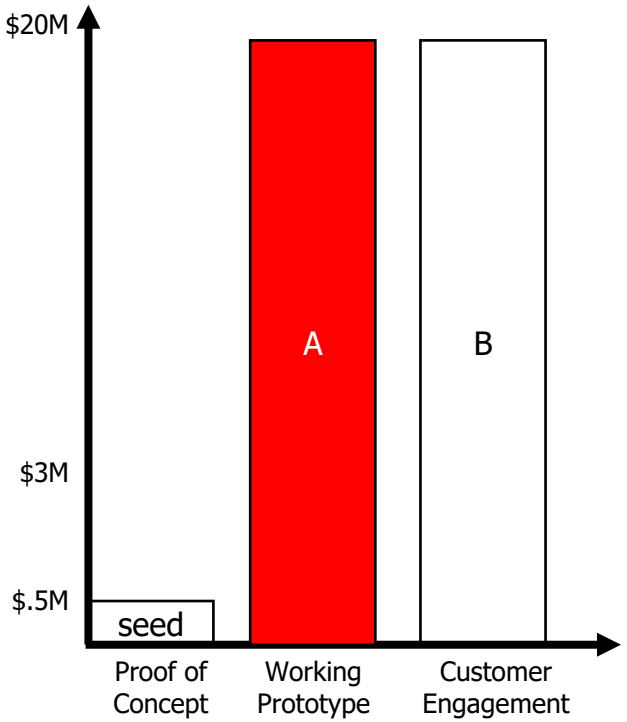

- Proof of concept requires working silicon
- $20M - $30M is too much to risk at early stage
- Rules out friends-and-family bootstrap
- If investment is found, it's highly dilutive
- Personal risk for the entrepreneurs is high

So What Happens Now?
Emergence of SoC Platforms and System-level Differentiation

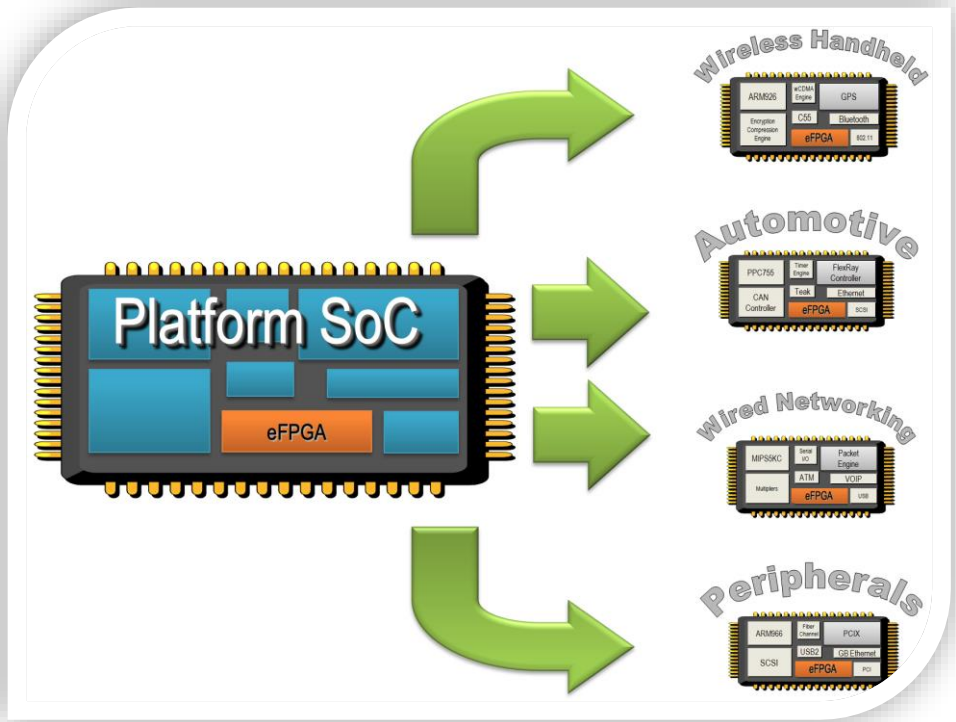

Source: Mentor

Aggregation of IP to Address Verticals

- Tuning of hardware and software interactions and task partitioning
- IP and software integration take on the paramount importance
- Provisioning of IP may differ for sub-segments

Design Challenges Shift Outside of Chip Design

- Ever-smaller, elite engineers will continue advancements in performance, power, integration challenges
- As the number of design starts shrinks there will be few, highly capable platforms around which value can be added
- These system-level capabilities will attract non-chip designers

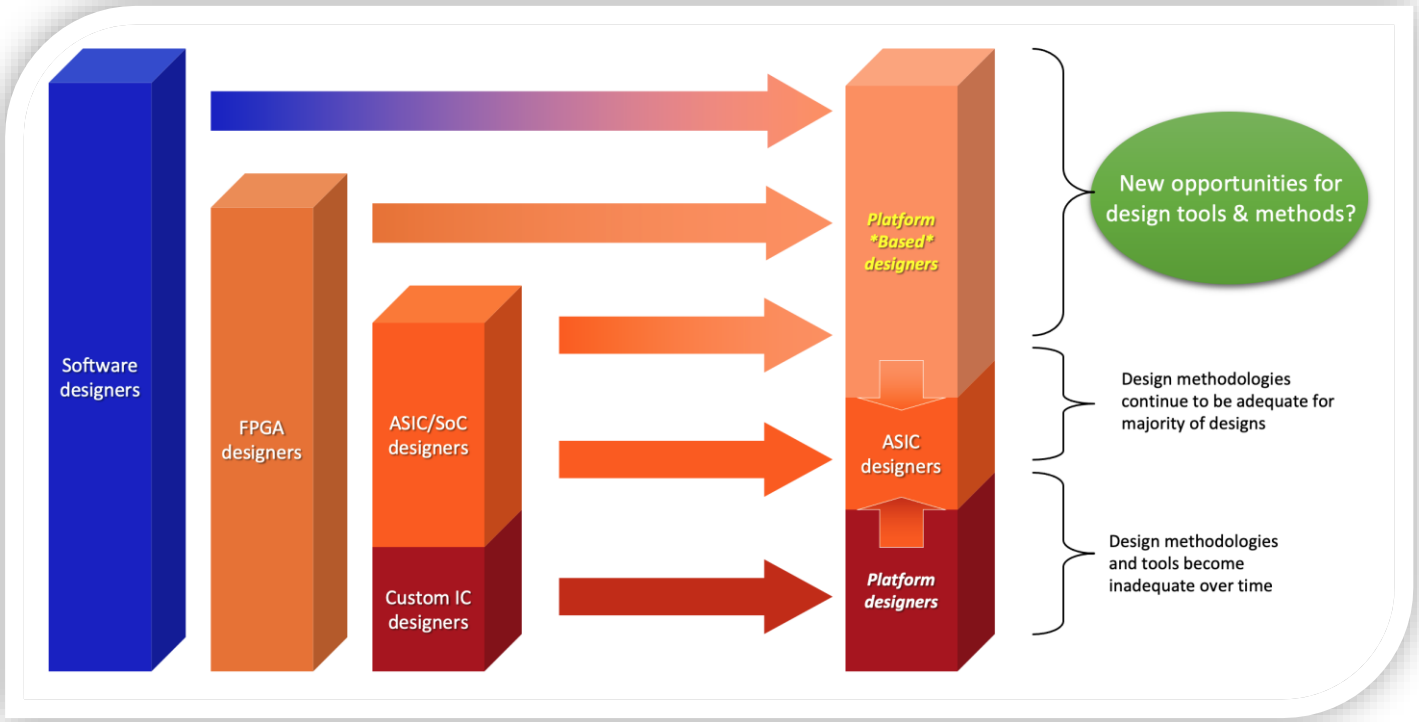

Enablement of Next Generation (System) Designs

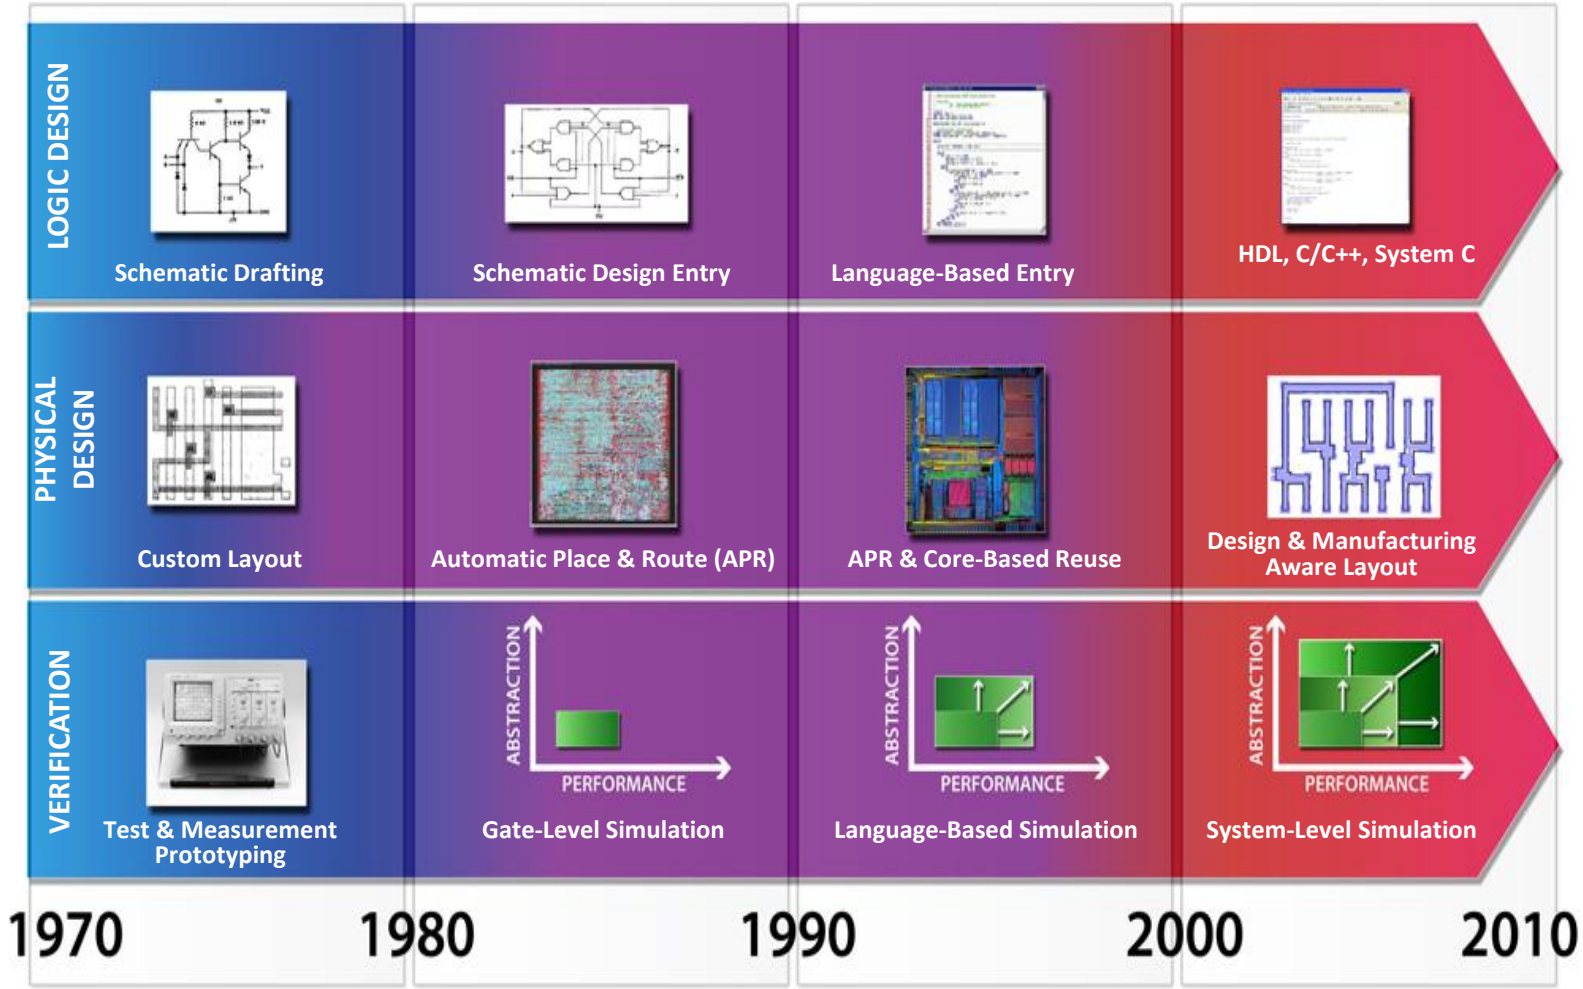

Source: Mentor

Energizing EDA Innovation

- **Platform-based Design**
  - Automated Creation of Accelerators
  - Run-time H/S Reconfigurability
- **System Synthesis**
  - Auto Integration of IP into Platforms
  - System Optimization based on power, area, speed and security
- **Advanced Simulation**
  - Linear or Hyper-linear Cloud Scaling
  - ML Trained Surrogate Models
  - Speculative Parallelism based HPC
- **Secure Silicon**
  - On-chip Security Engines
  - Multi Layer Security-in-depth

2025-2030

Platform-based Design

Platform-based Design: Example Flow

Platform Assembly

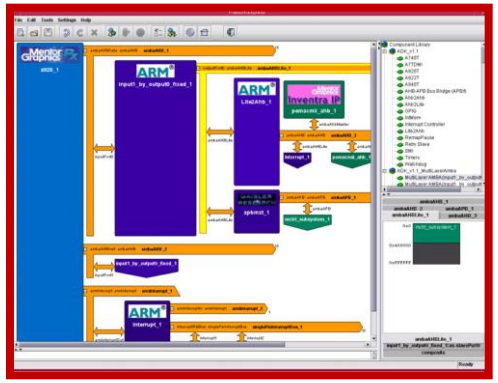

SoC System Simulation

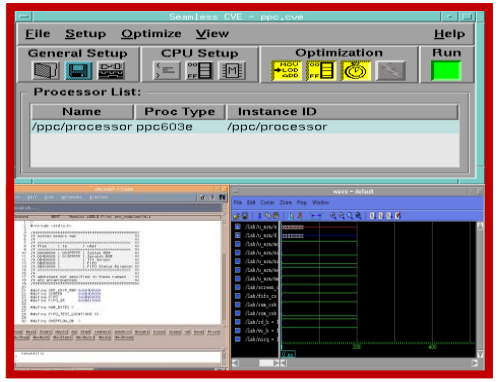

Accelerator Synthesis & Integration

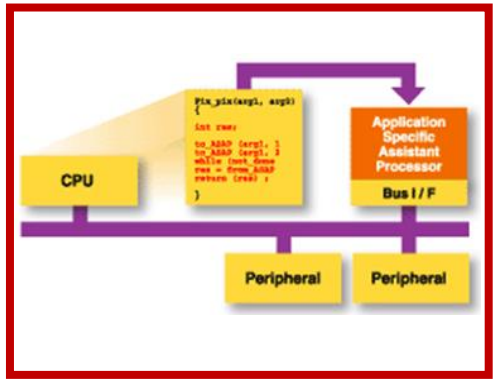

Profiling (Performance, Power…)

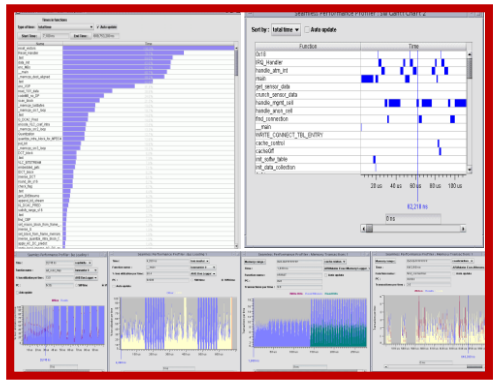

Source: Mentor Graphics, 2003

Platform-based Design: Automated Creation of Accelerators

- Move Software Functions to Hardware
  - Estimate System Level Design Impact
  - Rapidly assess the effect of proposed SW/HW changes
  - Generate fast C model & synthesizable RTL
- From a C function that is part of the software
  - Generate HW/SW interface
    - HW: Bus Interface and interface registers
    - SW: Software Driver to interface new hardware
- Verify the Generated RTL
  - With stimulus and expected responses
  - Automated path from software to hardware
  - implementation and verification
- Opportunity for run-time reconfigurability
  - eFPGA blocks can act as containers ready to house accelerators
  - Potential to use ML to predict ideal hardware/software partitioning DURING the execution in response to real time data and operational scenarios

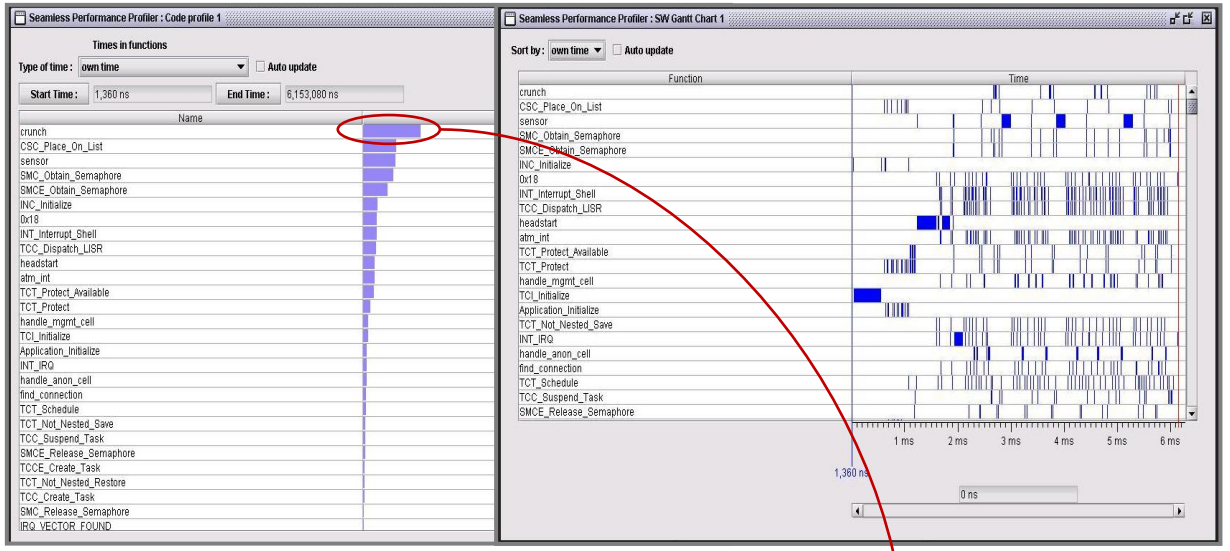

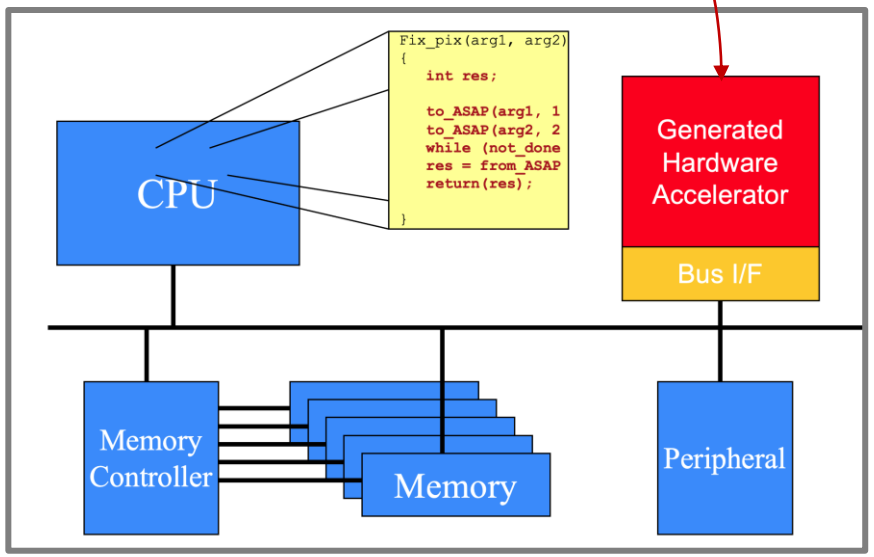

Source: Mentor Graphics, 2003

System Synthesis

System Synthesis: Multi-block IP Assembly, Parametrization, Integration

System synthesis & optimization

1. Σ(a***Performance**, b***Size**)
2. Σ(a***Performance**, b***Size**, c***Power**)
3. Σ(a***Performance**, b***Size**, c***Power**, d***Security**)

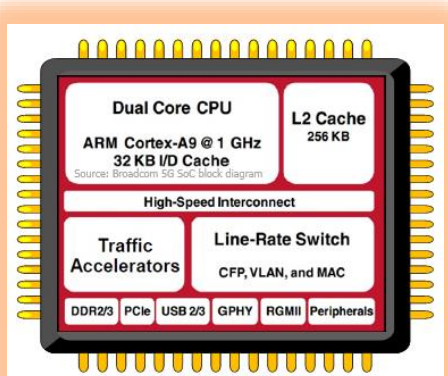

Source: Broadcom

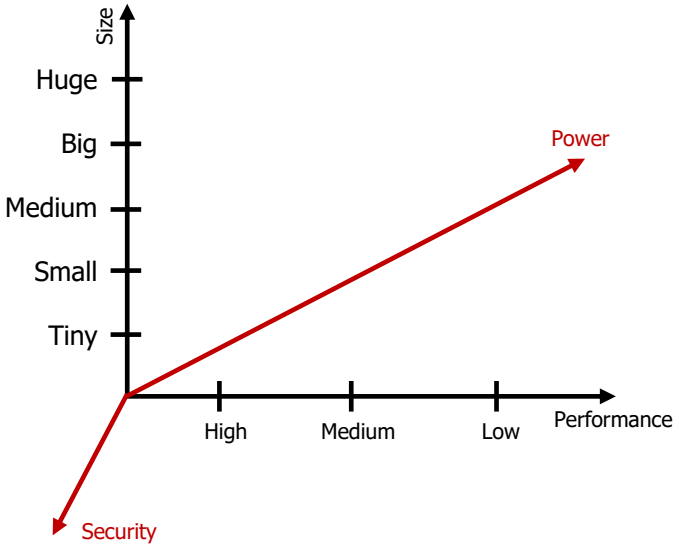

Some challenges:

- *Quantification of security & power*
- *Rapid assessment of candidate architectures*
- *Multi-dimensional optimization*

System Synthesis: Optimized Composition of Platform-based SoCs

- **Architectural Exploration to RTL Synthesis**
  - User selects a platform and supplies a cost function with size, performance, power and security goals to guide combinatorial optimization to find **best architectures** which are presented to the user for assessment and selection

Design: "Power Doors/Windows ECU"

Platform (Automotive Control)

- Performance = 2
- Size = 9
- Power = 3
- Security = 3
  - Supply Chain = 7
  - Side Channel = 2
  - Reverse Engineering = 5
  - Malicious Hardware = 1

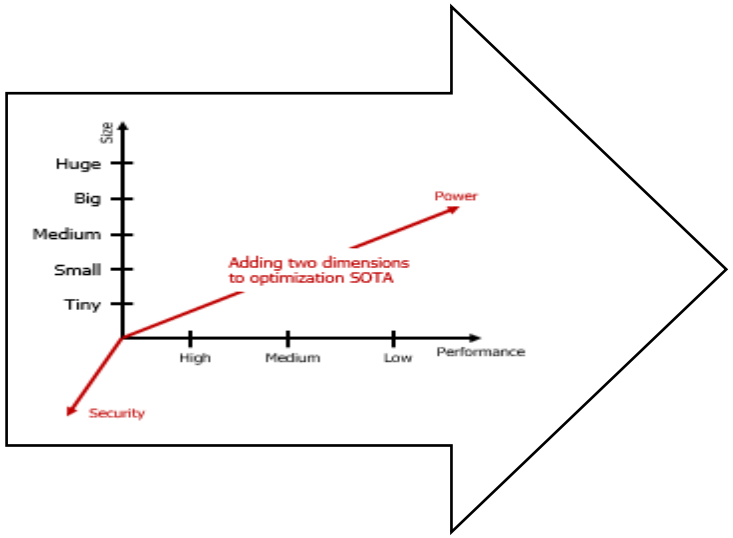

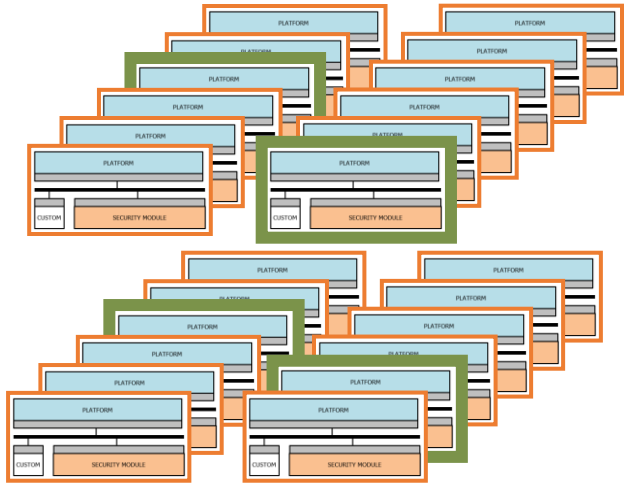

Combinatorial Optimization explores HUGE solution spaces (billions), but requires rapid estimation of "goodness"

Performance and *Size* estimators are well understood and incorporated in modern tools

Need to develop rapid estimation of **power** and **security**

$$f(a,b,c,d) = \sum(a*\textbf{Performance},\ b*\textbf{Size},\ \underbrace{c*\text{Power}}_{\text{estimate}},\ \underbrace{d*\text{Security}}_{\text{estimate}})$$

Advanced Simulation

Advanced Simulation: Reduced Order Models + HPC + Cloud Scaling

- Simulation physics limits reached decades ago and rely on Moore's law for speedups
- Parallel simulation has not been re-visited since emergence of the cloud computing
- Emergence of Machine Learning has not been employed to drive creation of faster models
- Need to re-think simulation in light of advances in HPCs, Cloud, and Machine Learning

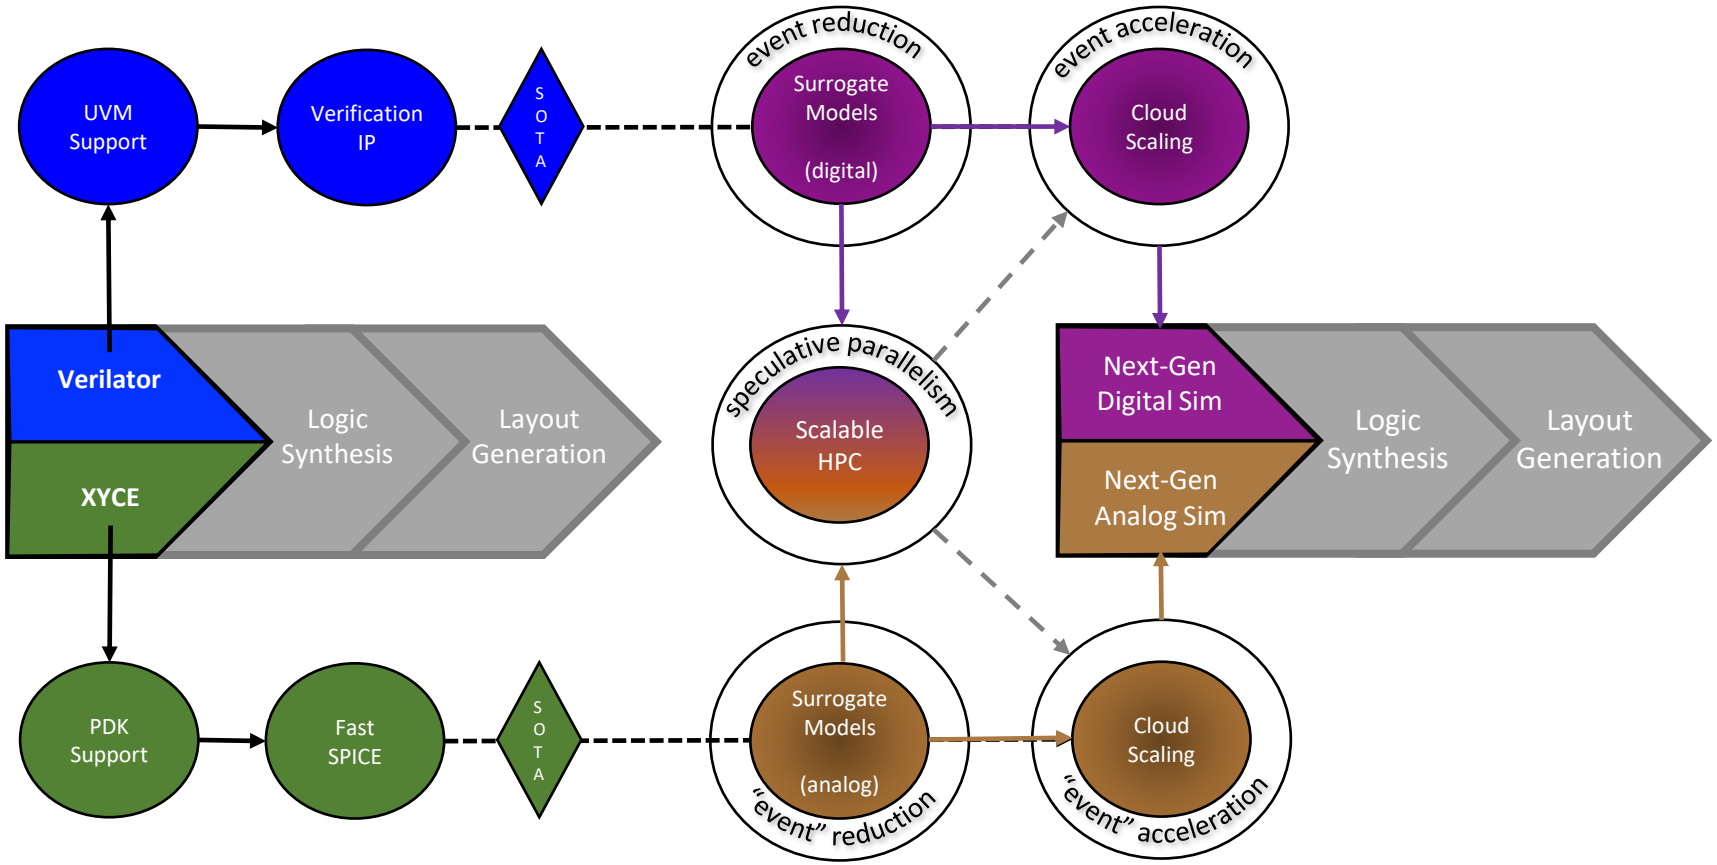

Advanced Simulation: Auto-generation of Reduced Order Models

- Drastically improve simulation performance in digital and mixed-signal SoCs by selectively substituting complex, high-fidelity circuit models with auto-generated, simplified, approximate surrogates

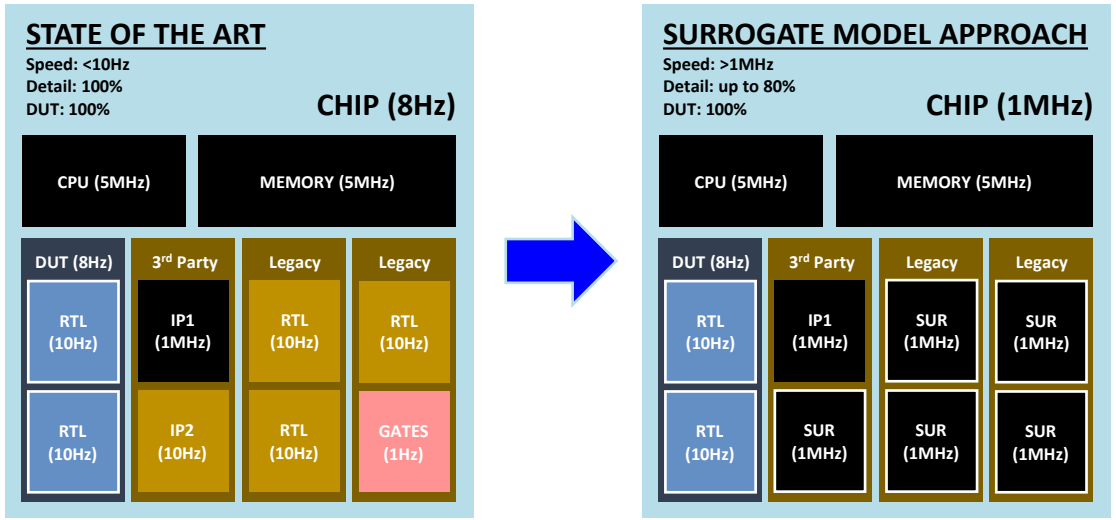

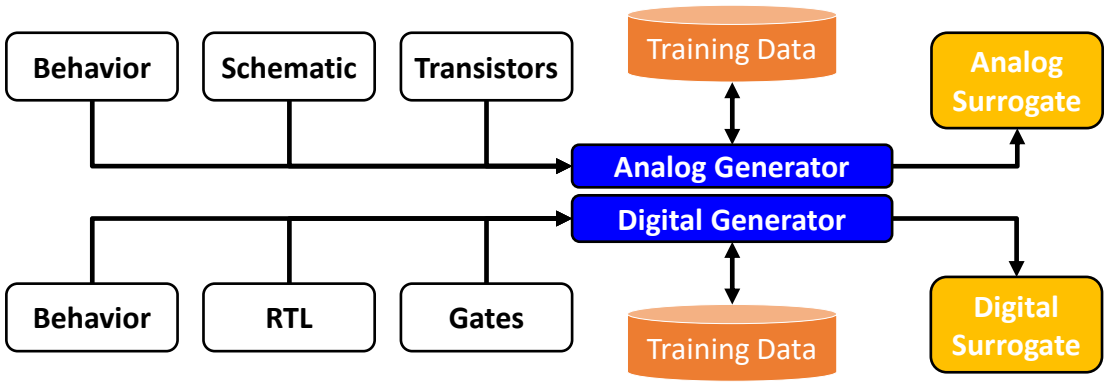

- Deploying state-of-the-art artificial intelligence techniques as the fundamental basis for generating simplified SoC models
- Develop machine learning models that can collapse/expand into different levels of hierarchy & have an awareness of their role within a larger system
- Make intelligent trade-offs between model accuracy and speed to achieve meaningful simulation

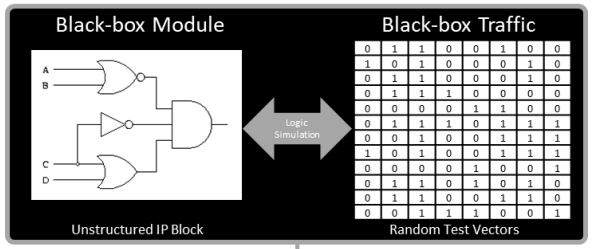

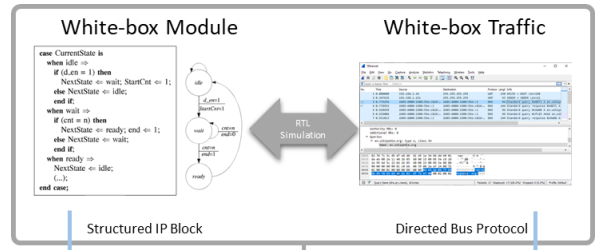

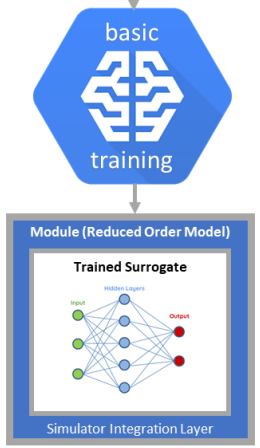

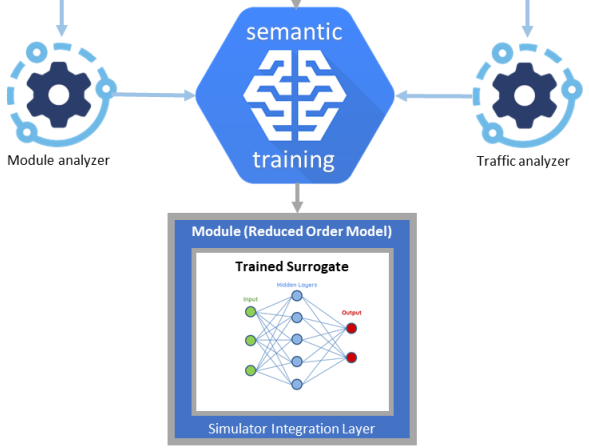

Sources:
https://ai.googleblog.com/2016/03/machine-learning-in-cloud-with.html
https://cloud.ibm.com/docs/AnalyticsEngine?topic=AnalyticsEngine-best-practices
https://i.stack.imgur.com/MG3ps.gif
https://ai.plainenglish.io/neural-networks-simplified-df7407201e88
https://en.wikipedia.org/wiki/File:Wireshark_Example_Decode.png

Nguyen, Duc-Minh & Thalmaier, Max & Wedler, Markus & Bormann, Jörg & Stoffel, Dominik & Kunz, Wolfgang. (2008). Unbounded Protocol Compliance Verification Using Interval Property Checking With Invariants. Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on. 27. 2068 - 2082. 10.1109/TCAD.2008.2006092.

Advanced Simulation: Enable Linear Scaling / Hyperscaling

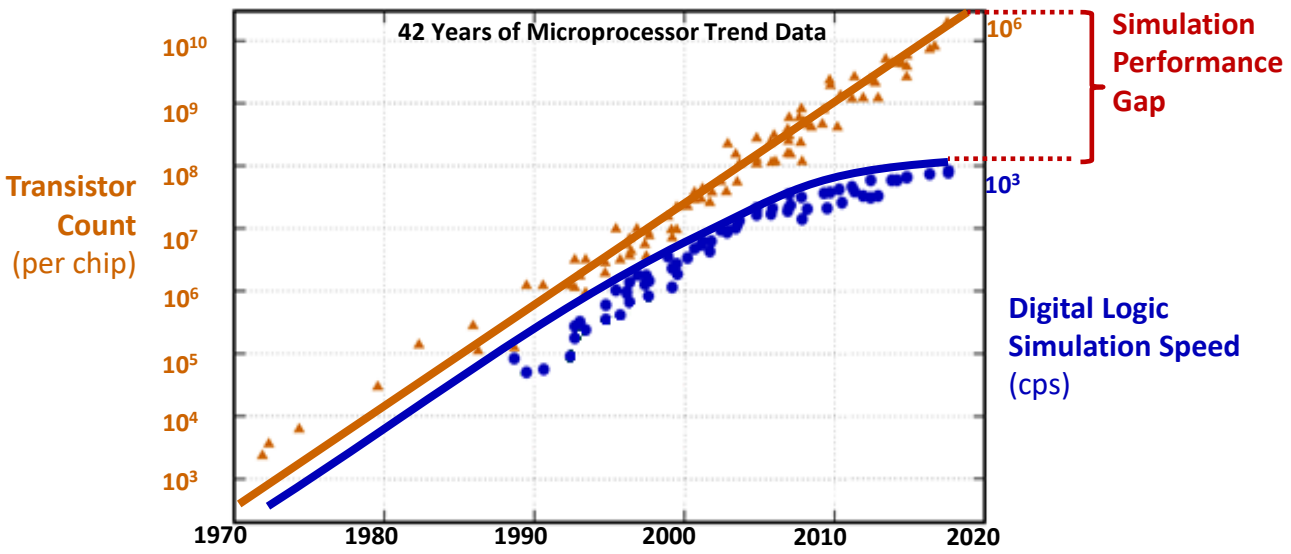

Source: https://github.com/karlrupp/microprocessor-trend-data

- Chip development timelines are bottlenecked by functional verification, where simulation speed is center-stage
- Simulation speed grew with Moore driven platform performance to ~ 1K cps
- Post-Moore, the architectural direction is multicore/cloud, but modern simulation algorithms are not designed take advantage of distributed computational fabric

- Combine limitless and elastic compute and storage available on the cloud with one or more of the following innovations:
  - Simulation engine novel algorithms
  - Parallel partitioning schemes
  - High-performance-compute (HPC) architectures and programmable cloud based FPGA fabric
  - ML-driven simulation partitioning

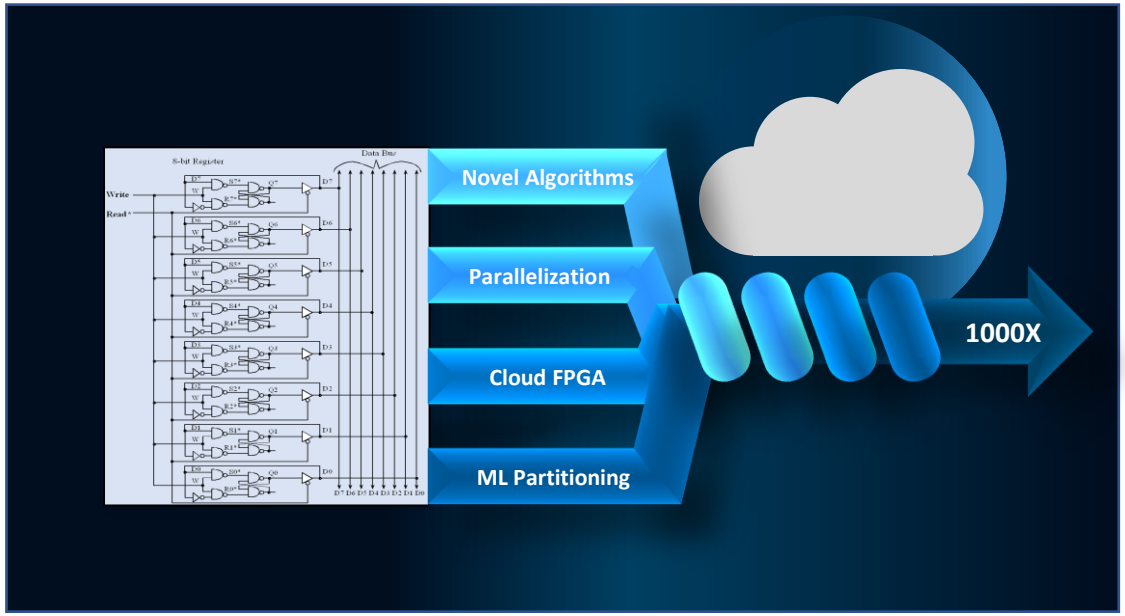

Source: http://users.ece.utexas.edu/~valvano/Volume1/

Advanced Simulation: HPC to Exploit Speculative Parallelism

- ASIC verification relies on single-threaded simulation algorithms (100 cps) and/or high-cost, inflexible, dedicated hardware emulators (1.5M cps)

- Concept of speculative parallelism dates back to 1987, but since then the computational fabric and opportunities have substantially evolved (Cloud, eFPGA, GPUs, etc.)

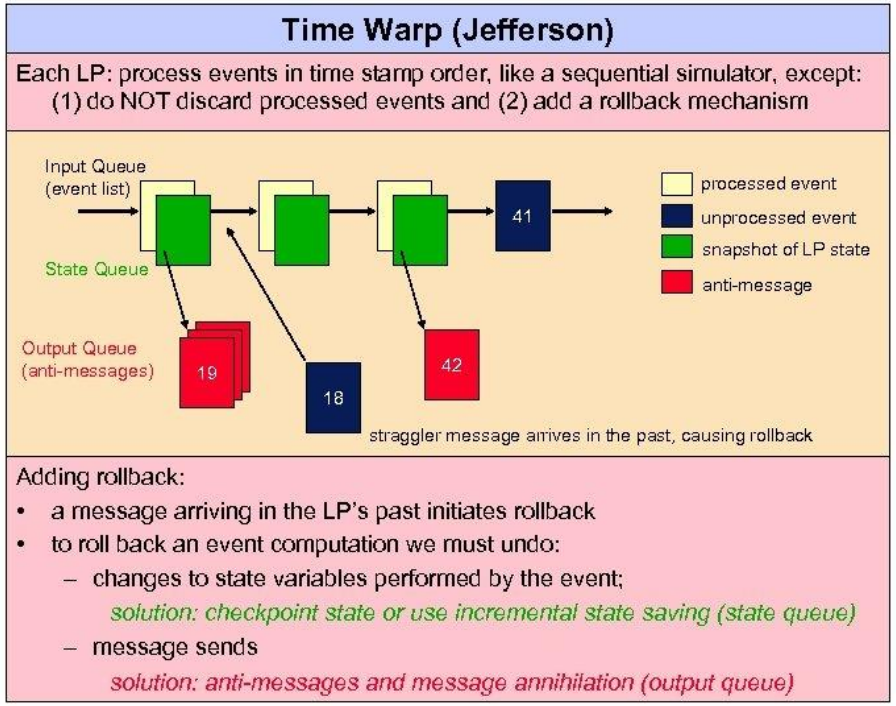

Source: R. Fujimoto, Georgia Institute of Technology

- Revisiting this strategy via specialized hardware → MIT Chronos

- Single-FPGA-based digital simulation accelerators have been prototyped, but have not demonstrated scaling off-chip or ASIC implementation

- Can we do better and achieve smooth system-level scaling?

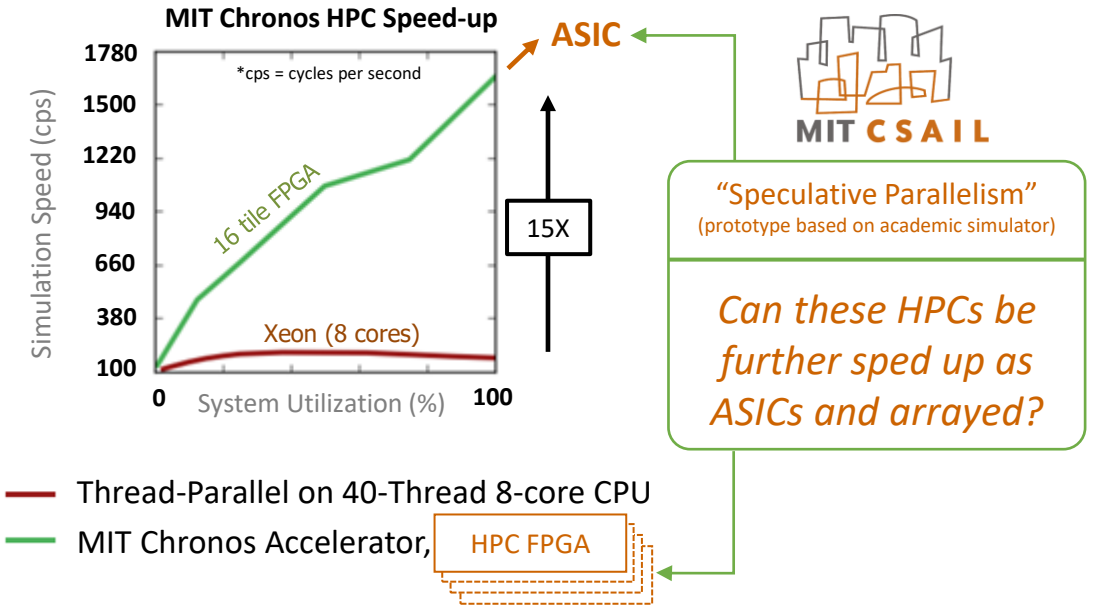

Source: MIT CSAIL

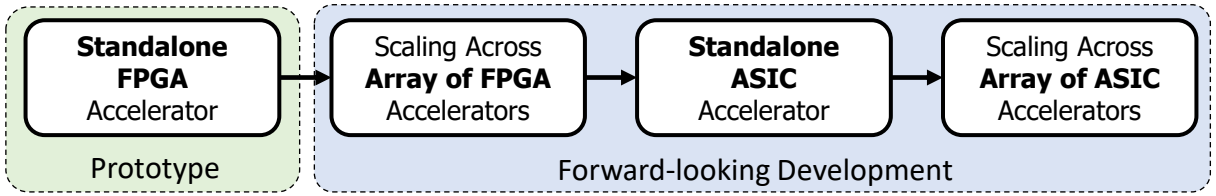

Secure Silicon

Secure Silicon: On-Chip Security Engine

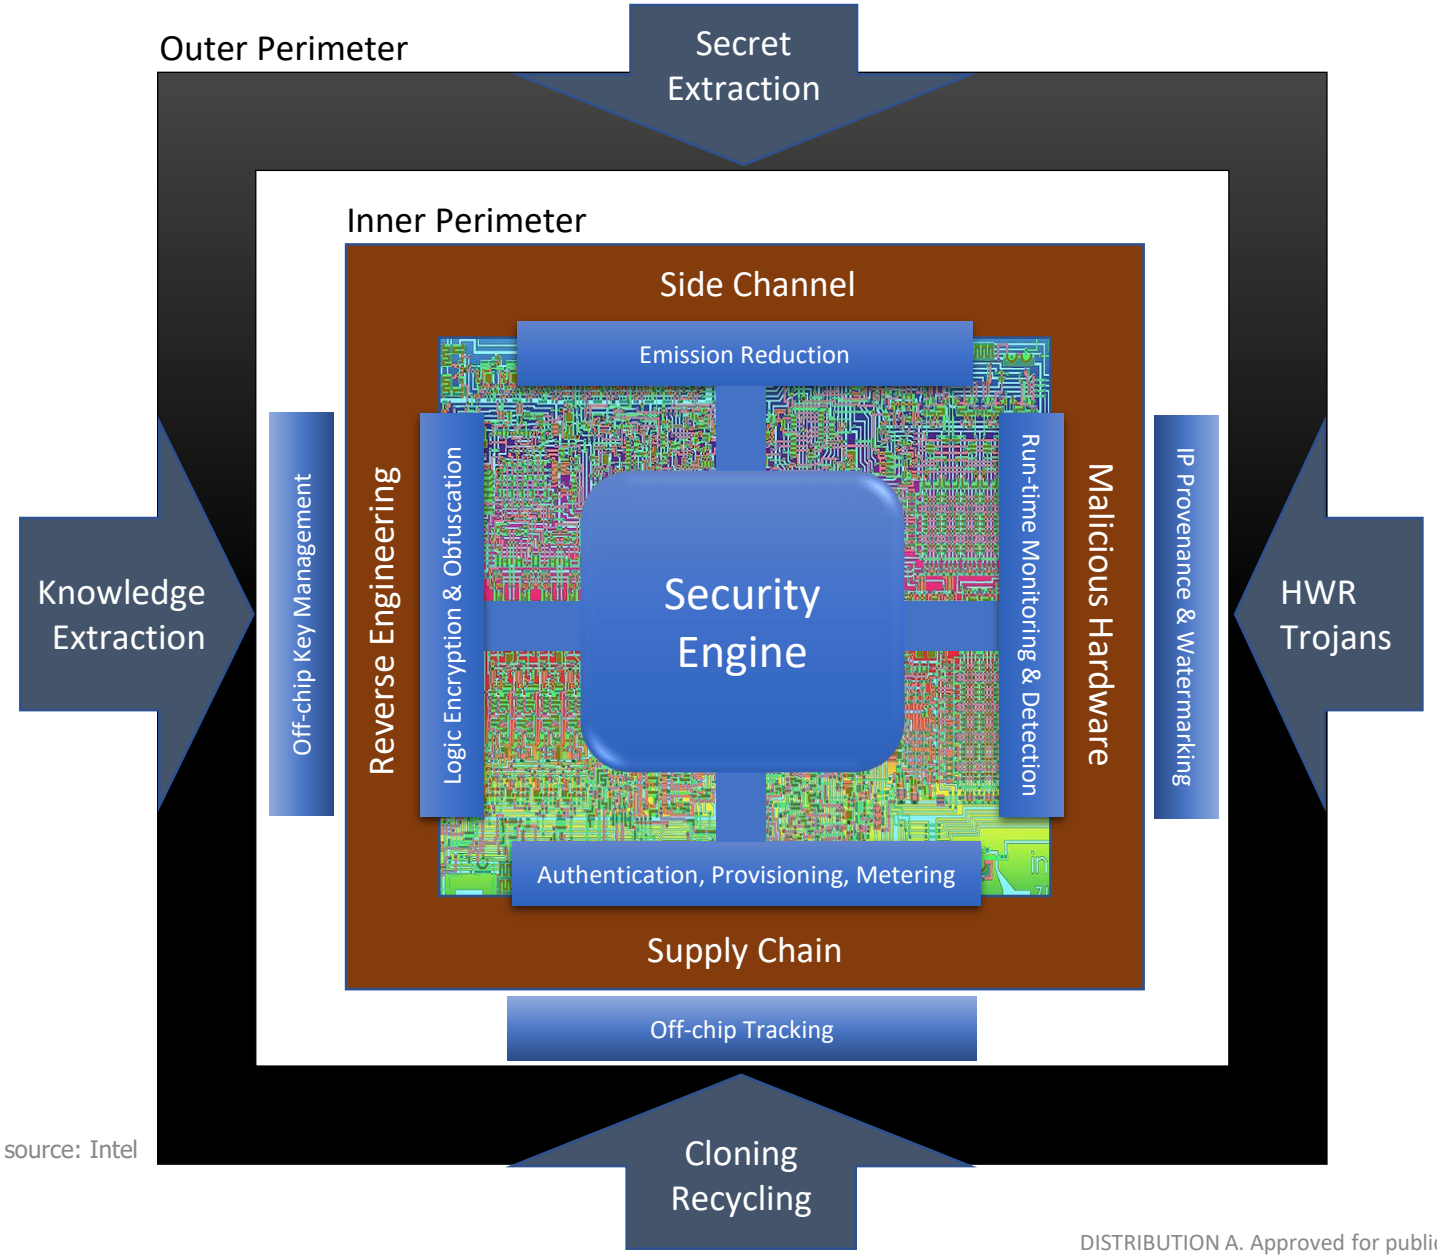

Image source: Intel

- There are many effective outer perimeter attack strategies
- We are assuming outer perimeter is penetrated or compromised
- AISS focus is only on securing inner perimeter with on-chip structures
- Some level of off-chip support is also needed

Secure Silicon: Three Security Engine (SE) Examples

- SE's serve as a Root-of-Trust
- Modular on hardware and software level
- Three ISA's: Arm, Synopsys ARC, RISC-V
- Each SE has its own unique architecture/features
- Each SE configurable based on PASS constraints
- Field upgradable security policies

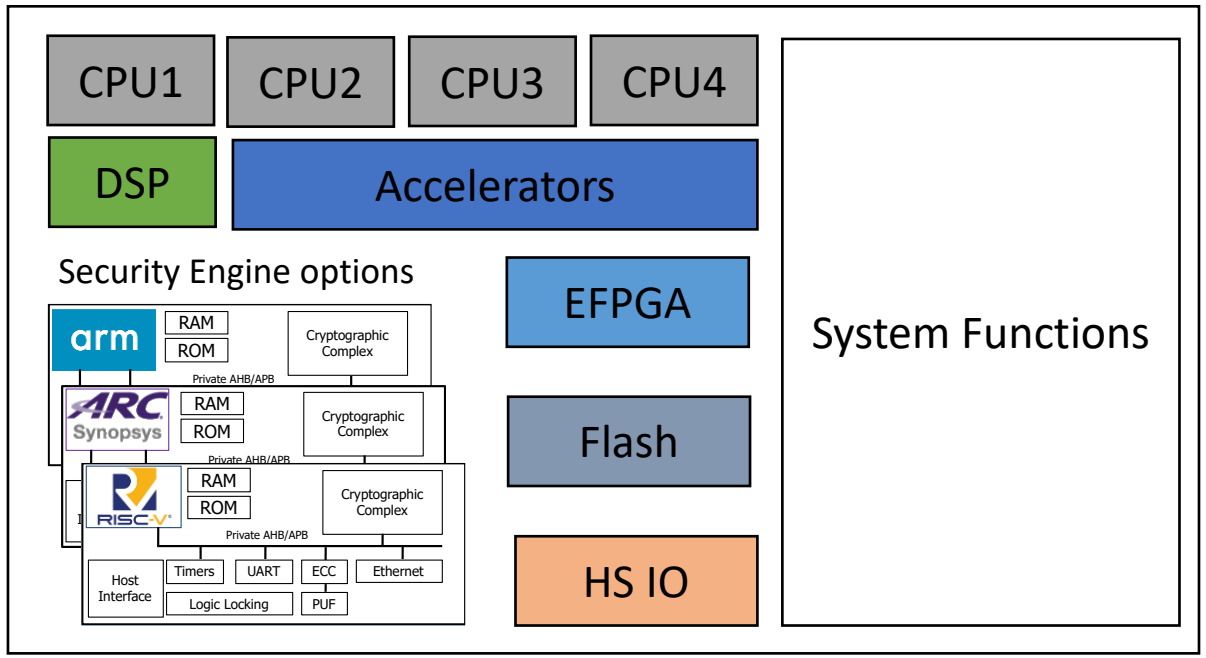

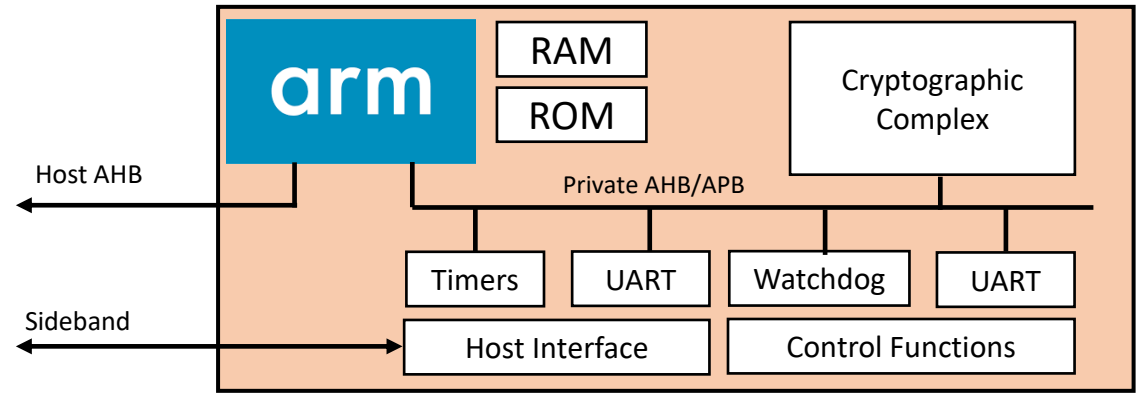

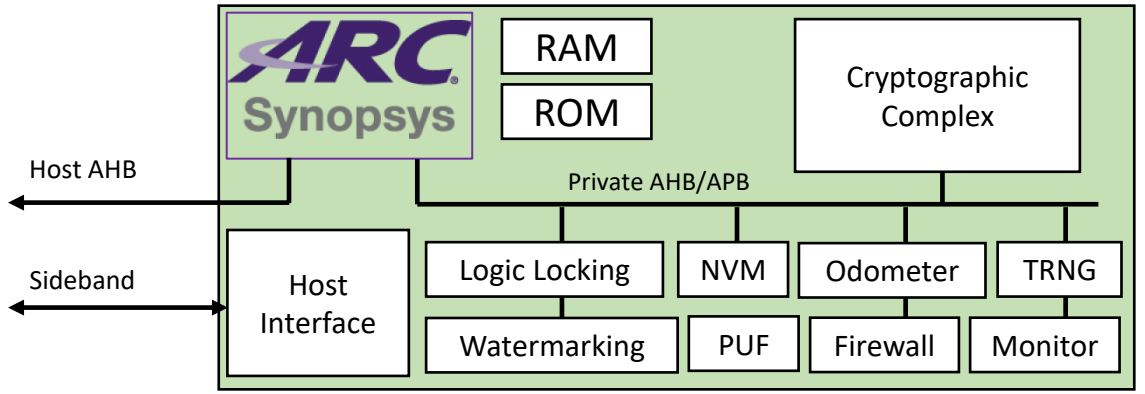

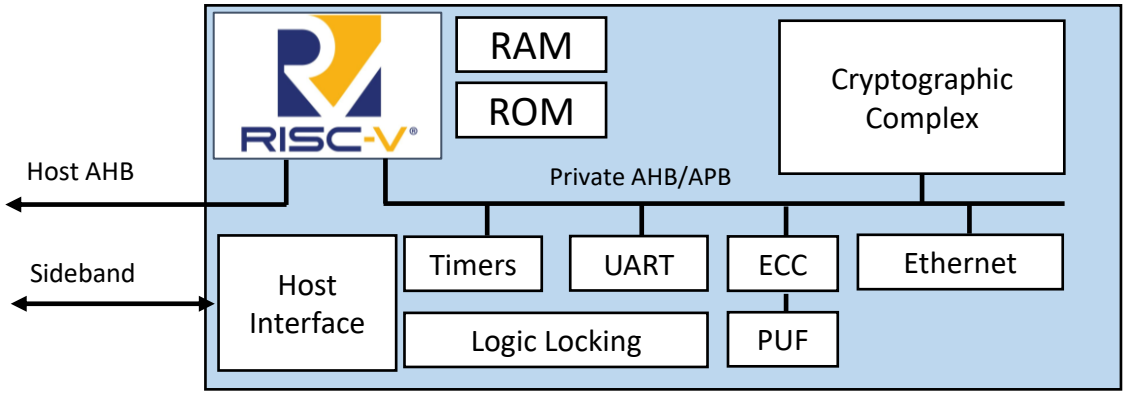

Sources: Synopsys, Arm, Northrop Grumman, RISC-V Foundation

Enabling Supply Chain Security

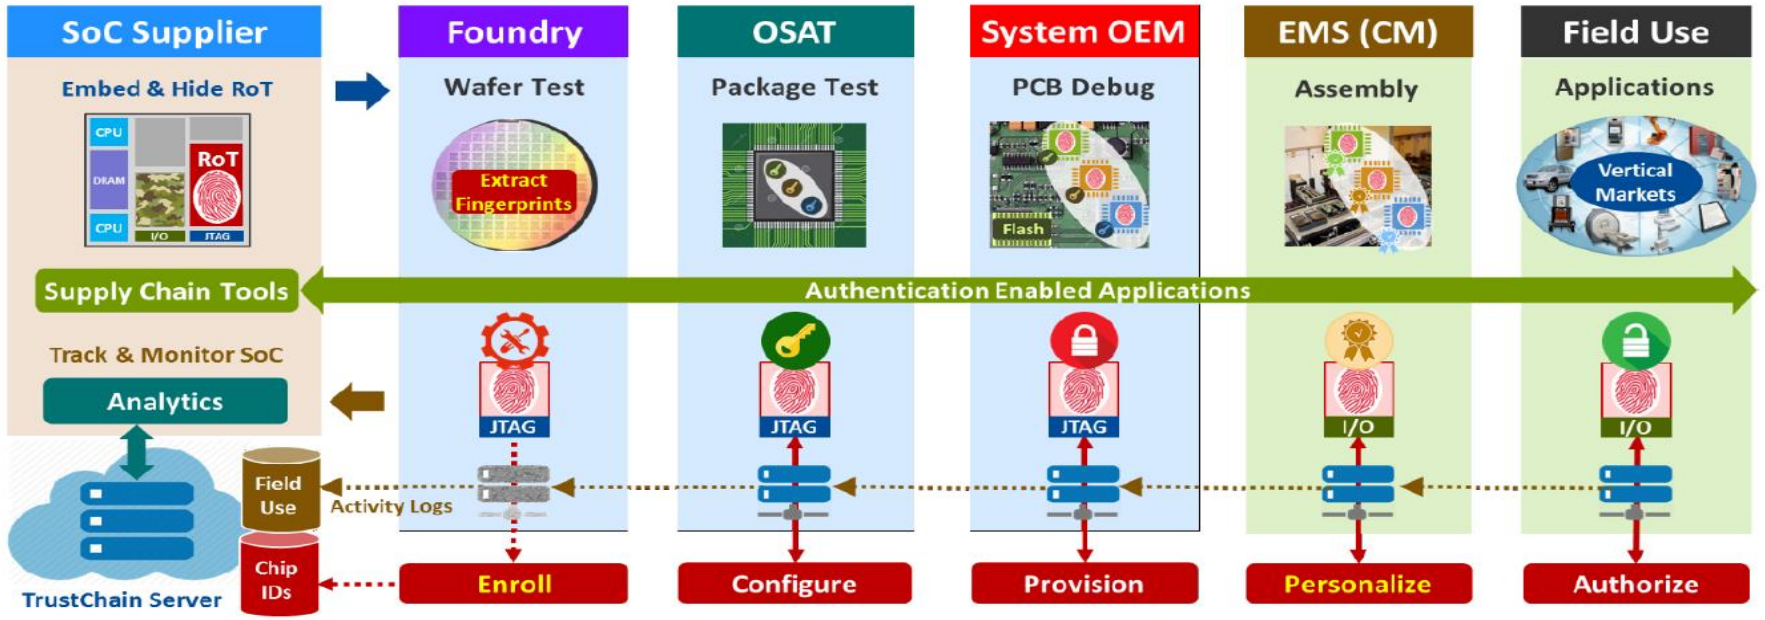

- **Design**: Create secure-reconfigurable SoCs with a unique ID based on an inborn Root of Trust
- **Enroll**: Extract chips unique ID into a secure server during first power up at wafer test
- **Configure**: Inject keys to encrypt, sign, or decrypt content for devices or end-applications
- **Provision**: Program SKUs downstream to reduce inventory risk and exploit volume ramp
- **Personalize**: Enables secure device identity during PCB assembly based on the chip's Root of Trust
- **Authorize**: Allow authorized parties to securely sign devices based on the SoC Root of Trust
- **Update**: Securely update firmware and provision SOC hardware features in the field
- **Monitor**: Track field use and evolve Big Data analytics on field failures, intrusions, counterfeits

Source: Mentor Graphics, 2017

Can DARPA Energize or Facilitate EDA Innovation?

Summary: Energizing EDA Innovation from DARPA

- **Platform-based Design** – AISS, IDEA, POSH, RTML
  - Automated Creation of Accelerators
  - Run-time H/S Reconfigurability
- **System Synthesis** – AISS, IDEA
  - Auto Integration of IP into Platforms
  - System Optimization based on power, area, speed and security
- **Advanced Simulation** – Ditto, POSH, <more to come>
  - Linear or Hyper-linear Cloud Scaling
  - ML Trained Surrogate Models
  - Speculative Parallelism based HPC
- **Secure Silicon** – SHIELD, AISS, ARCHS, SCATE, <more to come>
  - On-chip Security Engines
  - Multi Layer Security-in-depth

Call for Action

Advance EDA to help jump-start chip design companies and…

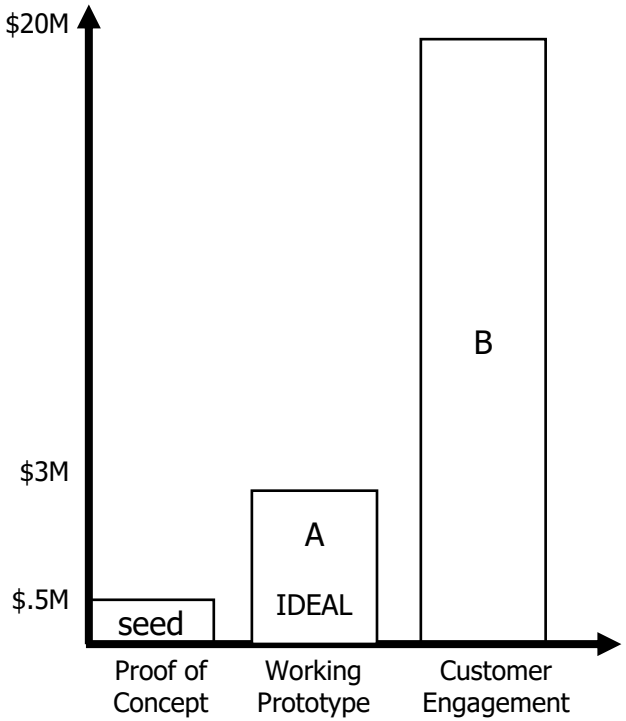

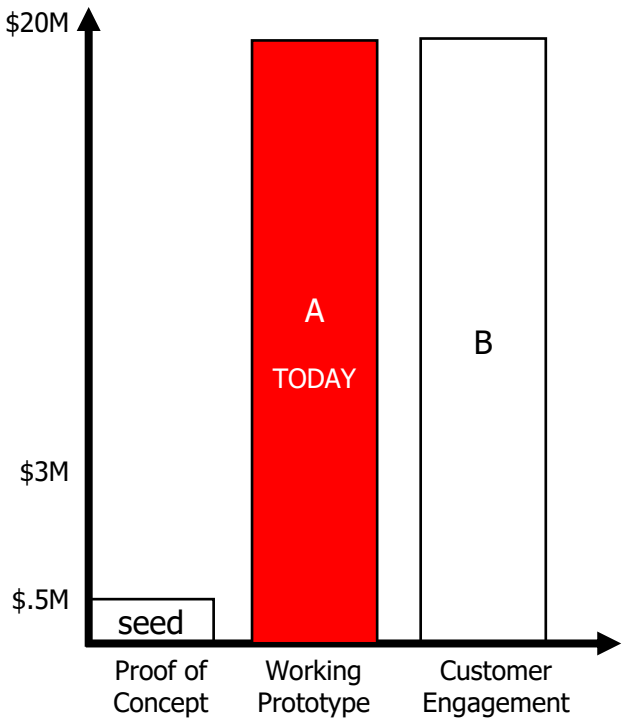

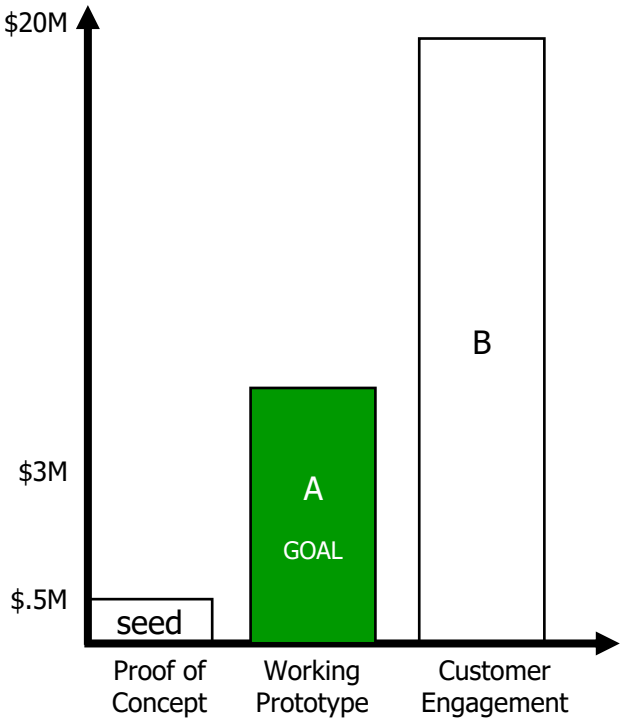

…substantially improve the economics for future US chip startups

DARPA Toolbox

Toolbox: DARPA Toolbox Initiative

What Reduce performers reliance on low-quality, low-cost tools & IP that increase execution risks & complicate post-DARPA transitions

Provide easy, low-cost, scalable access to SOTA tools & IP under predictable legal terms and streamlined acquisition procedures

Why Performers have hard time getting access to commercial grade design tools and IP

- Business model mismatch – pricing levels are based on expectations of customers' downstream revenue
- No volume purchasing – since the Government entities are not coordinated, they get no consistent pricing and all purchases are subject to protracted, one-of negotiations for each project
- Complex legal structures – all members of the performer teams are expected to execute agreements with ominous terms requiring lengthy multi-level reviews and approvals.
- Government contracting – is not set up for purchasing, owning and managing tool and IP licenses which further complicates access

How Framework that gives DARPA performers simplified, non-production, heavily discounted access to SOTA tools & IP

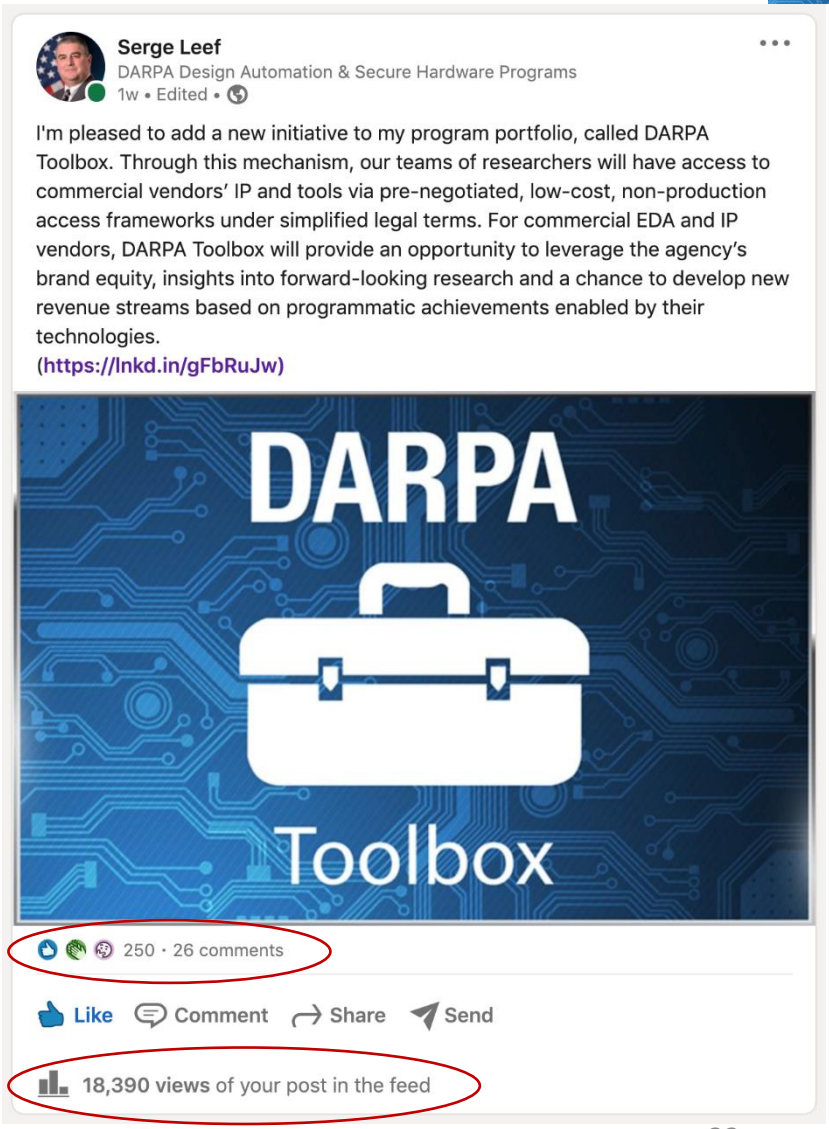

Toolbox: Framework

Value Proposition to Partners

- Strategic marketing to gain visibility and PR
- Brand association with the premier research funding agency
- Early visibility into leading edge research to gain insights into next generation requirements
- Establishment of new, tool/IP enabled capabilities in the DoD performer community
- Presence in the program makes partner a natural supplier in the transition phase where opportunities for production license sales exist

- **Approach**: XXXX creates a special non-production license for DARPA-funded program performers
- **Licensee**: Prime contractor who leads a team of DARPA performers under a direct contract with the Government
- **Duration**: Length of the program or 4 years whichever is shorter
- **Support**: One-time fee of XXXX paid by the prime and covering itself and/or all subcontractors for 1 year of support & maintenance, subsequent support years can be purchased as an option by the prime at the same or lower cost
- **Consulting**: XXXX will offer a unit of consulting that costs XXXX and represents XXXX hours – this will be optionally available for purchase by the prime for itself and/or all subcontractors at a cost representing a substantial discount relative to list
- **Purpose**: Demonstration of capability / proof-of-concept as required by the DARPA programs
- **Restrictions**: Derivative products may not be distributed under commercial terms or for use in production
- **Contracting**: DARPA & XXXX would execute a 3-year, no-cost agreement that capture these terms
- **Promotion**: Access to XXXX technology will be communicated to the proposers using Industry Day or FAQ vehicles
- **Acquisition**: Prime proposers contact XXXX directly for quotes and licenses and roll the costs into proposals

Toolbox: Methodology

- **External facing web pages**
  - Explain the initiative to the PMs, Proposers and Suppliers
  - Provide summary of all vendors and products available to the DARPA community
- **At Proposer/Industry Day**
  - If PM deems Toolbox elements to be relevant, he/she can
    - Discuss them in their presentation or
    - Give stage time to a Supplier's representative
    - Explain general structure and terms of the business model
    - (Re-iterates lack of endorsement or vendor preference)
  - CMO presentation will include a page summarizing all suppliers and assets available in the Toolbox
- **During proposal development**
  - A prime proposer executes an NDA with the relevant Toolbox supplier for technical planning and pricing discussion
  - A supplier issues a price quote to the the prime to ultimately be included in the proposal document set
- **During program execution**
  - If selected, the prime and the Supplier execute a licensing agreement
  - The Supplier delivers products, licenses, support and maintenance to the prime & its team members

Toolbox: Current Partners

DARPA Toolbox Initiative – Suppliers

Participation in DARPA Toolbox provides commercial vendors – or Toolbox Suppliers – with early insight into DARPA innovations and research findings, which can help inform next-generation requirements and potential product advancements. Through this effort, Suppliers also have the opportunity to access the Agency's ecosystem of innovators, researchers, and thought leaders. Further, your presence on DARPA programs makes you a natural supplier in the transition phase of the Agency's programs where opportunities for production license sales exist.

Instructions for New Suppliers

- Email Toolbox@DARPA.mil to
  - Describe the proposed products, representative customers, and use cases
  - Explain why you believe these are of interest to the DARPA community
  - Provide a brief summary supporting that these are **proven** technologies **commonly used** in state-of-the-art commercial microelectronics or system design methodologies
- You will receive an email with disposition of your submission. If approved, a call will be set up to discuss business model alternatives as they apply to licensing, support, and maintenance and consulting services.
- Once an agreement of business terms is reached you will
  - Download and fill out the sample agreement
  - In a separate file, itemize each product that is proposed for the DARPA Toolbox Initiative
  - Provide a "pricing comparison example" that clearly demonstrates that DARPA performers are getting substantially more favorable pricing than commercial customers
  - Submit these items to Toolbox@DARPA.mil. You will be contacted by the Contract Management Office (CMO) to finalize and execute the agreement. The expected timeline for this interaction should be 1-2 weeks.

RESOURCES

- Overview
- Suppliers
  - Andes
  - Arm
  - CEVA
  - eMemory
  - Flex Logix
  - Lattice Semiconductor
  - QuickLogic
  - Rambus
  - Riscure
  - Secure-IC
  - SiFive
  - Tortuga Logic
  - Verific Design Automation
- Proposers
- Programs Using Toolbox

- If you wish to publicly announce your involvement in DARPA Toolbox, the announcement content must first be reviewed and approved by DARPA
- Contributed products will be added to the web pages and to the presentation that the Contract Management Office makes at ALL future DARPA Industry Days
- **Promotion to the PMs is planned for late summer**
- **Promotion to the community is planned starting with the in-person Industry Days to be hosted at DARPA**

Q&A

www.darpa.mil