

# Nearly “Forgotten” Memories about Parallel Performance from 30+ Years Ago

Vladimir Getov

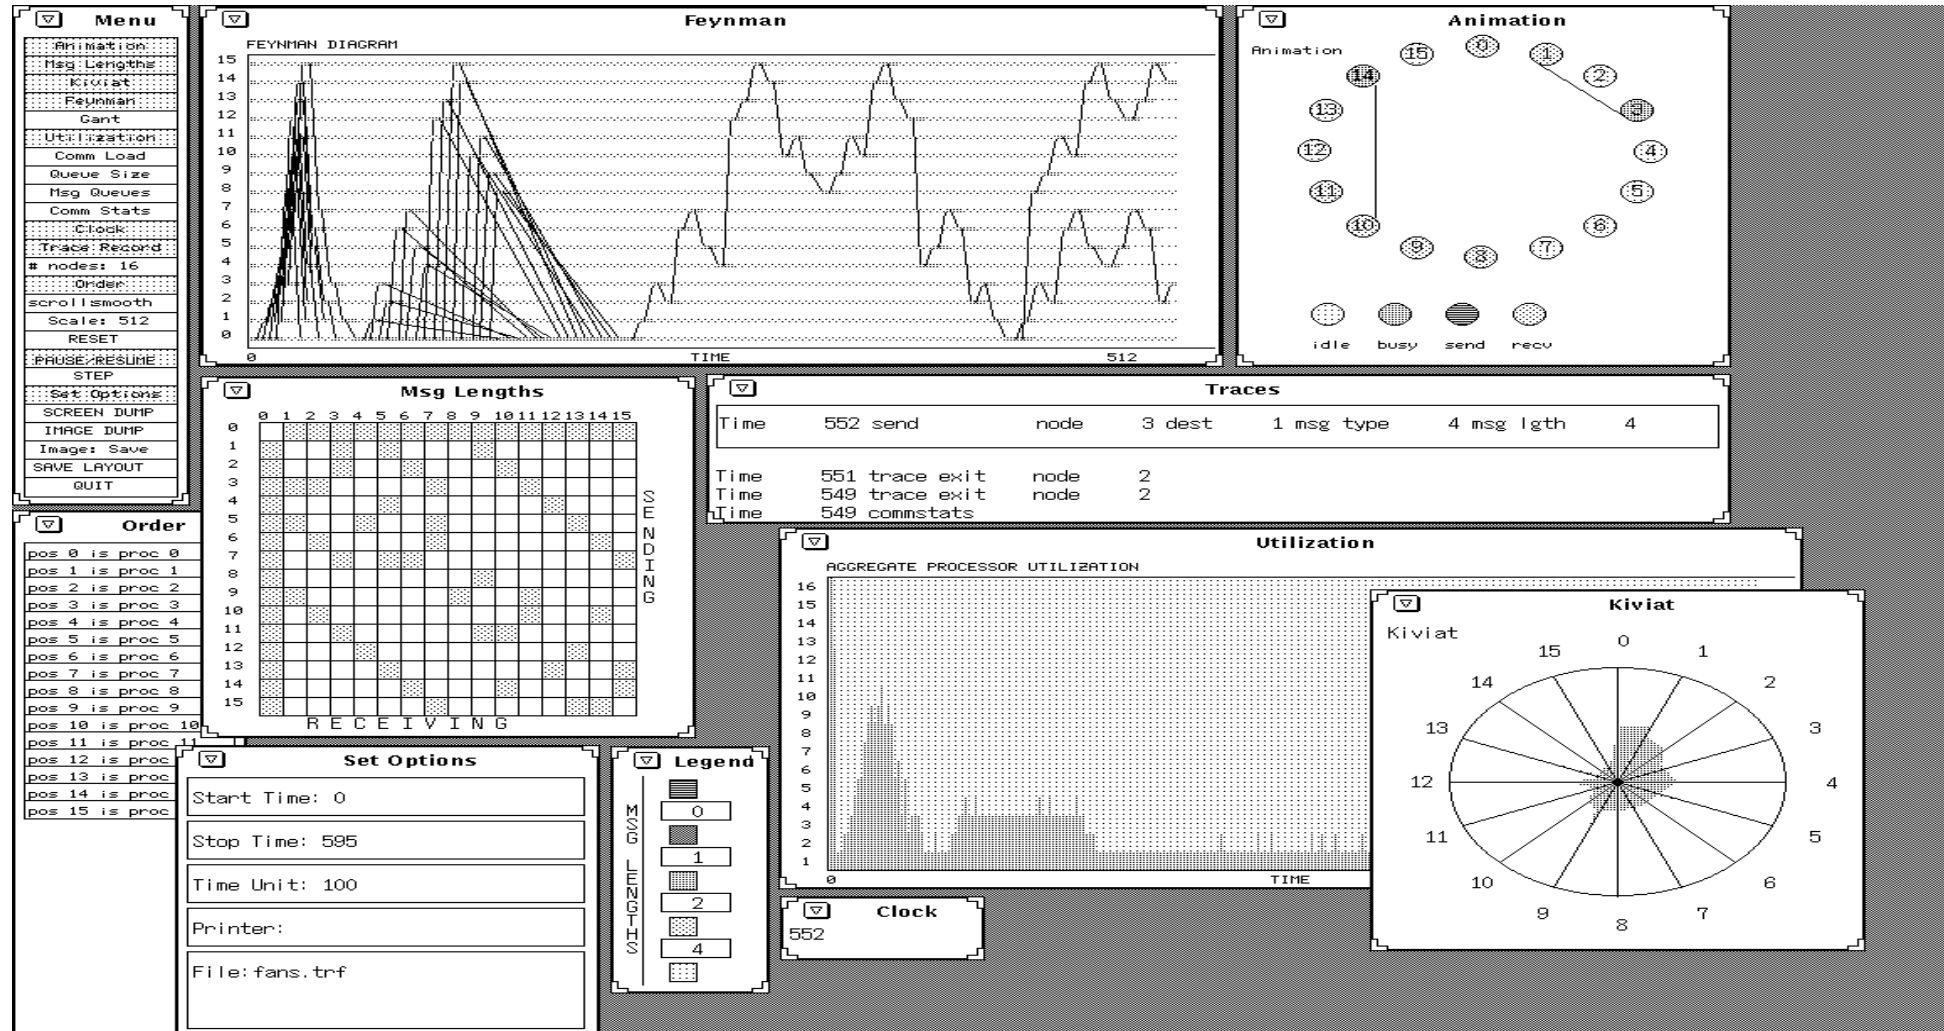
Distributed and Intelligent Systems Research Group  
University of Westminster – London

APART 25th Anniversary Workshop  
Obergurgl, Austria, 12-15 February 2024

# Early Projects, Tools, and Environments

- GENESIS project (1989 – 1992) – SUPRENUM computer, PARMACS message passing, Meiko CS2
- GENESIS parallel benchmarks (1989 – 1994)
- PPPE – Portable Parallel Programming Environments, HPF-like programming environments, ParaGraph monitoring and visualisation (1992 – 1994)
- PVM + PARMACS = MPI-1 (1992)
- PARKBENCH committee and codes (1994 – 1997)
- RAPS benchmarks -> Pallas benchmarks -> Intel benchmarks -> ...

# ParaGraph (1992) based on PICL – M.T. Heath



# Novel Methodology for Application Performance Modelling and Evaluation

Vladimir Getov

Distributed and Intelligent Systems Research Group  
University of Westminster - London

APART 25th Anniversary Workshop

Obergurgl, Austria, 12-15 February 2024



# Outline

- IRDS: Overview and Structure
- International Focus Teams and Key Roadmap Elements
- Interaction Between Applications Benchmarking and Systems and Architecture IFTs
- Selection of Representative Application Domains – Market Drivers
- Critical Issues and Tradoffs – Moore's 'Law'
- Example: Physical System Simulations
- Key Messages
  - New Technology Requirements
  - Breakthroughs in Technology, Research
- Appendix:
  - Team Members
  - Collaborative Alignments



IEEE  
COMPUTER  
SOCIETY



IEEE



IEEE COMPUTER SOCIETY  
TFRC

Task Force on Rebooting Computing




INTERNATIONAL ROADMAP FOR DEVICES AND SYSTEMS

© Vladimir Getov - all rights reserved

# Roadmap Visions Through the Years

The International Roadmap for Semiconductors (ITRS) and its evolution to the International Roadmap for Devices and Systems (IRDS) provided leadership and continues to play a key role in guiding design and implementation of devices and systems.



IEEE  
COMPUTER  
SOCIETY



IEEE

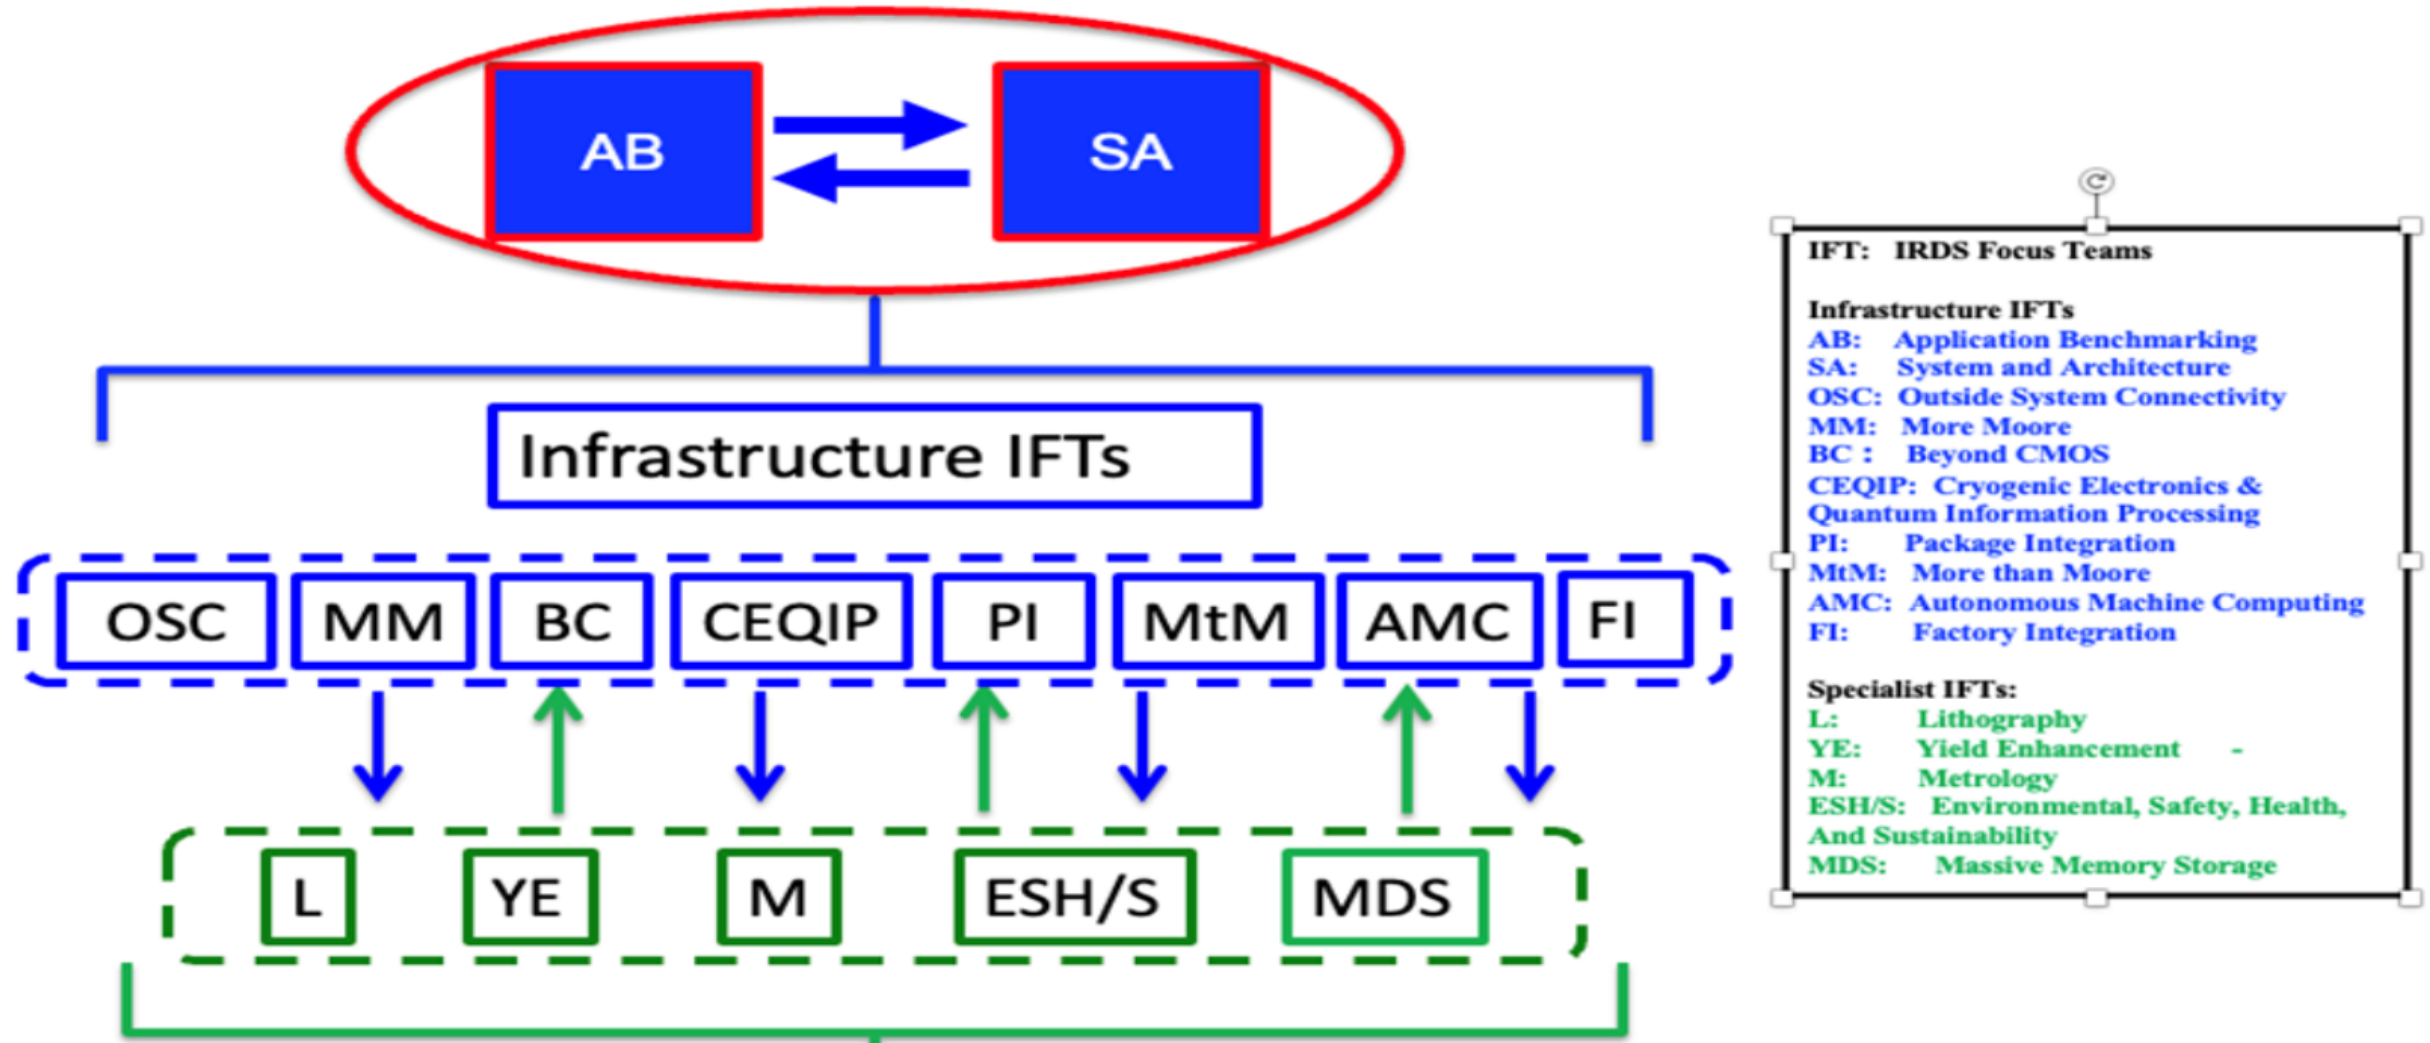


IEEE COMPUTER SOCIETY  
TFRC  
Task Force on Rebooting Computing



© Vladimir Getov - all rights reserved

# IFT structure of the International Roadmap for Devices and Systems (IRDS)





# IRDS: Key Roadmap Elements - <https://irds.ieee.org/>

To achieve its goals, each IFT assesses how their technology could evolve and identifies the following:

## Difficult Challenges and Showstoppers

- Top 5 Challenges for Near-term and Top 5 challenges for Long-term

## Technology Requirements Guidelines

- Annualized tables show technology needs and level of difficulty and where gaps in solutions may occur

## Potential Solutions Guidelines

- IFTs review solutions for assessed needs
- The potential solutions chart indicates maturity of a particular solution
- IFTs do not select only one solution; but instead include areas for innovative answers



IEEE  
COMPUTER  
SOCIETY



IEEE



IEEE COMPUTER SOCIETY  
TFRC

Task Force on Rebooting Computing



© Vladimir Getov - all rights reserved




# Interaction Between AB and SA IFTs

- **Applications Benchmarking (AB)** brings the critical understanding of “what do we now and what will we need to be able to compute?”
- **Systems and Architectures (SA)** brings the boundary conditions of “what are the space, weight, power, privacy, security and sustainability criteria at each design envelope from the edge device to the Exascale data center?”
  - Four System Categories:
    - Data Center (Hyperscale and HPC)
    - CPS
    - Personal Augmentation
    - IoTe
  - Eleven Design Envelopes
    - $\mu$ W to MW
    - mm to km
    - mg to metric tons

# Applications Benchmarking Highlights and Emerging Applications

- The electronics industry in general, and the computer industry in particular is driven by *application domains*
- *Emerging* applications require new technologies thus defining roadmap trajectories – e.g. personal augmentation, AI, autonomy



# Systems and Architectures Scope – Market Drivers

- **Cloud:** This category is for server devices deployed in data centers. The term “cloud” refers to the engineering of data center scale computing operations: compute, storage, networking engineered for scale and for continuous resource redeployment and reconfiguration via APIs.
- **Internet-of-Things edge devices (IoT-e):** Although IoT is a broad class of computing applications spanning the server to the ultimate sensors and actuators, an IoT edge (IoTe) device is a wireless device with computation, sensing, communication, and possibly storage.
- **Cyber-Physical Systems (CPS):** This category encompasses computer-based control of physical devices characterized by real-time processing and used primarily in industrial control. Many cyber-physical systems are safety-critical.
- **Personal Augmentation (PA):** Personal augmentation devices provide multiple use cases: telephony and video telephony; multimedia viewing; photography and videography; email and electronic communication; positioning and mapping, authenticated financial transactions, health and fitness monitoring, personal safety and environmental warning.



IEEE  
COMPUTER  
SOCIETY



IEEE



IEEE COMPUTER SOCIETY  
TFRC  
Task Force on Rebooting Computing



INTERNATIONAL ROADMAP FOR DEVICES AND SYSTEMS

© Vladimir Getov - all rights reserved

# Technology and/or Research Breakthroughs: Application Areas

		<i>Application Area</i>	<i>Description</i>
		Graph Analytics	Applications of graph algorithms to large static or dynamic graphs
		Artificial Intelligence	Modern artificial intelligence applications with emphasis on machine learning approaches. Graphical dynamic moving image (movie) recognition of a class of targets (e.g., face, car). This can include neuromorphic / deep learning approaches such as DNNs.
Rep-?	{	Discrete Event Simulation	Large discrete event simulation of a discretized-time system. (e.g., large computer system simulation) Generally used to model engineered systems. Computation is integer-based.
		Physical System Simulation	Simulation of physical real-world phenomena. Typically, finite-element based. Examples are fluid flow, weather prediction, thermo-evolution. Computation is floating-point-based, includes mixed precision.
Rep-?	{	Optimization	Integer NP-hard optimization problems, often solved with near-optimal approximation techniques.
		Graphics/VR/AR	Large scale, real-time photorealistic rendering driven by physical world models. Examples include interactive gaming, Augmented Reality, Virtual Reality.
		Cryptographic codec	Cryptographic encoding and decoding, including specialized hardware acceleration
NEW	{	Video codec	Encoding, transcoding of video, AOM AV-1
		Autonomy	Autonomous route planning, motion planning, navigation, end-to-end control
		ML for Science	Use of machine learning techniques for scientific exploration (eg, graph NNs)
		IoT	Applications that run at IoT edge or in the fog

# IRDS AB IFT Cross-team Alignments

<https://irds.ieee.org/editions/2023>

- Since the inception of IRDS in 2016, the Applications Benchmarking effort has been positioned as a “top end” to the roadmap. Its mission is to identify key application drivers and to track and roadmap their performance.
- The mission of the Applications Benchmarking IFT is to identify key application drivers, and to track and roadmap the performance of these applications for the next 15 years. Given a list of market drivers from the Systems and Architectures International Focus Team (SA IFT), AB IFT generates a cross matrix map showing which application(s) are important or critical (gating) for each market.



# Cross Matrix: Application Areas vs. Market Drivers

Application Area	Cloud	IoTe	CPS	PA
Graph Analytics	Y	Y	Y	Y
Artificial Intelligence	Y	Y	Y	Y
Physical System Simulation	Y	Y	Y	
Cryptography	Y	Y	Y	Y
Video codec	Y	Y		Y
Machine Learning for Science	Y			
Internet of Things Applications		Y	Y	Y

# Current Critical Issues


- Technology constraints on applications have remained relatively constant since last edition of roadmap
- Software ecosystem changes impact application performance
  - Ex: changes in software stack and application profiles for Artificial Intelligence
- Memory performance is critical for all application areas
- Stall in technology power efficiency leading to specialization
- Continuing challenge: some application areas, although highly important, remain difficult to track quantitatively



# Background: Moore's 'Law'

- At the inaugural Int. Solid-State Circuits Conference at the University of Pennsylvania in 1960, a young computer engineer named Douglas Engelbart introduced the computing industry to the remarkably simple but ground-breaking concept of “scaling.”
- Another young engineer – Gordon Moore – was in the audience. In 1965, Dr Moore sketched out his prediction of the development pace of silicon technology. Moore's law describes a long-term trend in computing hardware.


# Gordon Moore's Original Sketch



It states that “the number of components that can be placed inexpensively on an integrated circuit doubles approximately every two years.”

G. E. Moore, “Cramming more components onto integrated circuits”, Electronics, vol. 38(8), pp. 114-117, April 1965.

# Moore's Law for Clock Rate



# 26 July 2000: Intel Predicts 10GHz Chips by 2011 - ?!

“Intel is predicting that its microprocessors will hit 10GHz clock rate by 2011.”

ZDNet

Has this prediction worked as before?

NO!

# Is Moore's 'Law' going to end and what are we going to do then?

- It sounds scary – when?!?
- What does it currently state?
- “The number of people predicting the death of Moore's law doubles every two years!”
  - Peter Lee, VP Microsoft Research, March 2016.
- Robert H. Dennard's scaling (1974) stopped working nearly 20 years ago.
- Subsequently, we entered a new era of dark silicon, multithreading and energy consumption challenges.



# Dennard's Scaling

- “As the dimensions of a device go down, so does power consumption. Smaller transistors run faster, use less power, and cost less.”
- Is this true nowadays? – NO!
- Why? – leakage, negligible before, approached the same order of magnitude as the chip's dynamic power.
- While feature sizes have continued to shrink, threshold voltage has not since switching a transistor at a lower threshold voltage needs a thinner gate dielectric, but leakage places a lower bound on dielectric thickness.

# Dark Silicon

- The failure of Dennard's scaling has introduced a new era of plenty of transistors, but not enough power.
- If the number of transistors doubles, the power available for each transistor is cut in half – e.g. only half the number of transistors can operate at the same time.
- Calculating the energy consumption of a generic chip is difficult. It depends on a wide range of factors.



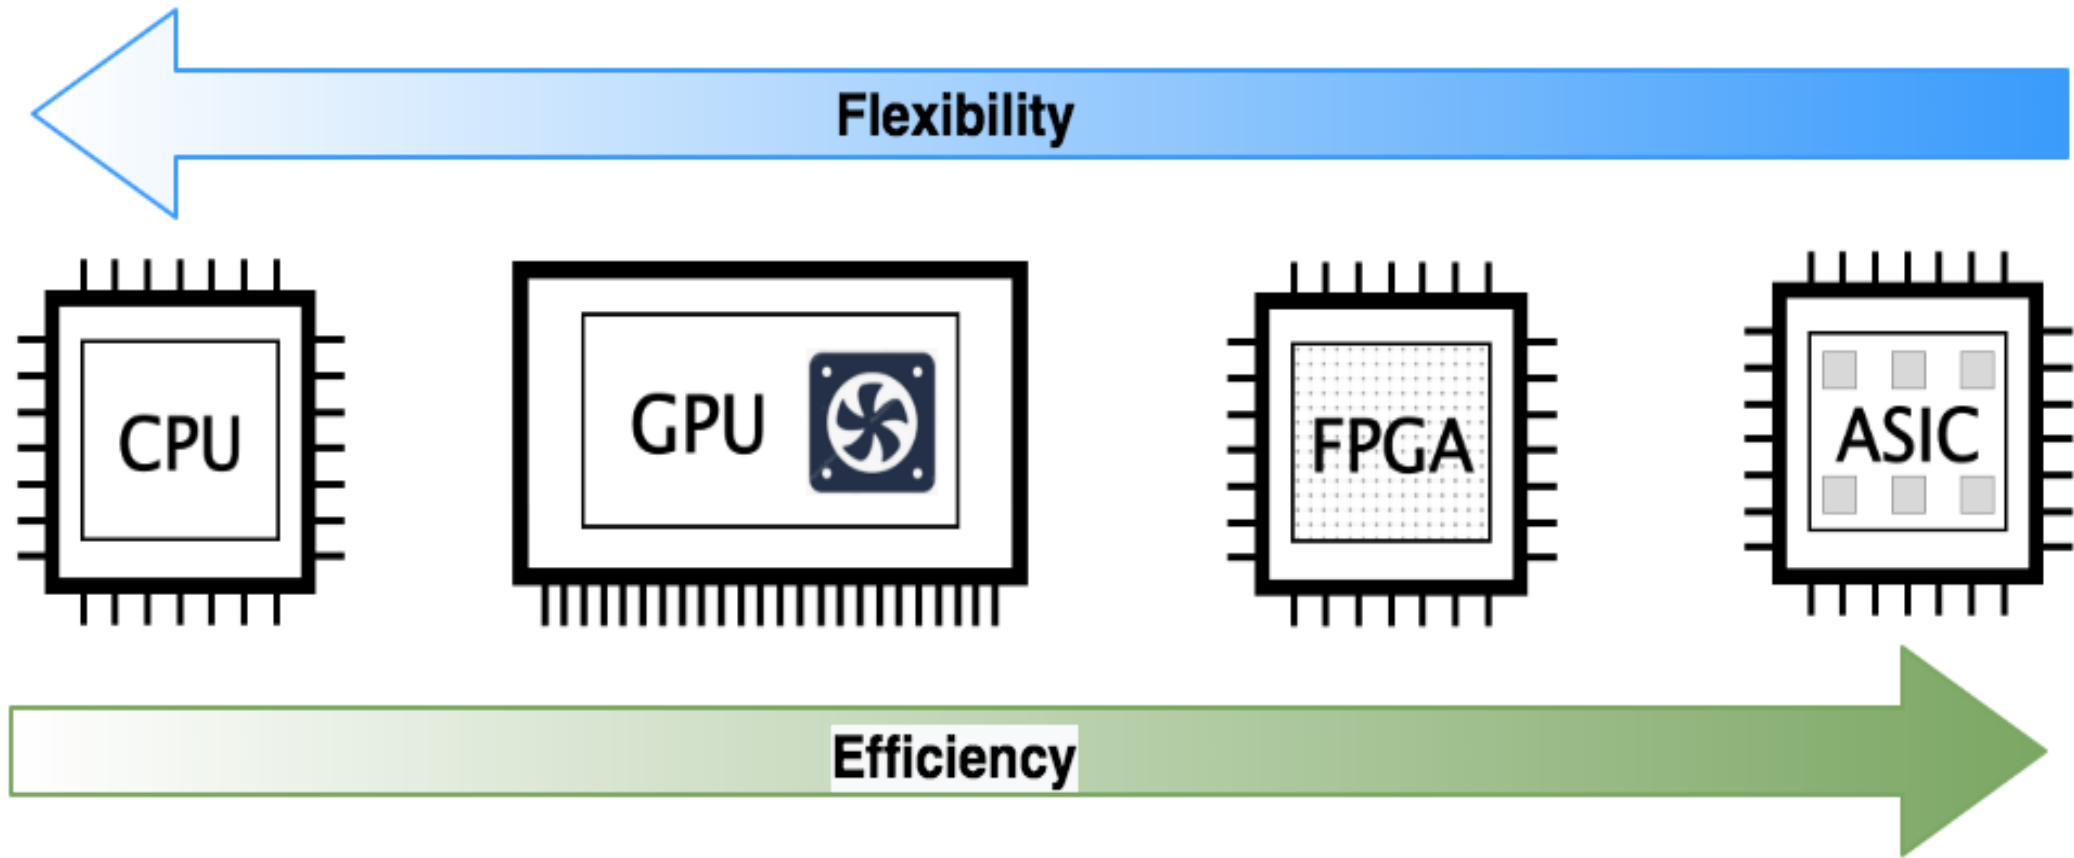
# When Do We Need New Architectures?

- When we hit a “wall” for some important class of applications
- 1st Wall – Mid 90s: the Memory Wall – much faster processor speed
- 2nd Wall – 2004: the Power Wall - the failure of Dennard scaling, emergence of dark silicon
- 3rd Wall – Now: the Locality Wall - The data locality we expect from our apps is disappearing

# Moore's 'Law' Nowadays

- Still **WORKS** according to the IEEE IRDS roadmap
- Limits – “power wall” challenges
  - Power dissipation of air-cooled chip
  - Little ILP left to exploit efficiently
  - Almost unchanged memory latency
- In 2004 Intel cancelled HPC uniprocessor projects
  - Thread and data level parallelism
  - In comparison to ILP, need programmer input
- Modern processors - multicore and manycore, accelerators, GPUs and FPGAs, Air vs liquid cooling

# Critical Tradeoffs Among Computer Architectures



# Current Key Messages

- Reduced *device efficiencies* have resulted in plateaus in application power vs. performance
- Specialization (e.g., FPGA, ASIC) is the preferred solution for increased efficiency for most application domains
- Memory bandwidth is the most critical technology need across all domains
- IRDS / IEEE relies on external parties for benchmarking, and benchmarking changes impact our ability to track performance and efficiency
  - In the future: IRDS / IEEE standardization and control of benchmarking for the roadmap



IEEE  
COMPUTER  
SOCIETY



IEEE



IEEE COMPUTER SOCIETY  
TFRC  
Task Force on Rebooting Computing



INTERNATIONAL ROADMAP FOR DEVICES AND SYSTEMS

© Vladimir Getov - all rights reserved

# New or Expanded Technology Requirements

<i>Application Area</i>	<i>Improvement Paths</i>				
	<b>Algorithmic improvement</b>	<b>Memory bandwidth</b>	<b>Memory latency</b>	<b>Network bandwidth</b>	<b>Fixed-function acceleration</b>
Graph Analytics		X	X	X	
Artificial Intelligence	X	X			X
Discrete Event Simulation		X	X		
Physical System Simulation	X	X	X	X	X
Optimization	X	X	X		
Graphics/VR/AR	X	X			X
Cryptographic codec		X			X
Video codec	X	X	X		X
Autonomy	X	X	X	X	X

# Example: Physical System Simulations

- Computer simulation of physical real-world phenomena emerged with the invention of electronic digital computing
- Led to the creation of supercomputers; large distributed systems; access to huge data sets, and high throughput communications
- Subsequently, the term 'e-science' was adopted to capture these new revolutionary methods for scientific discovery



IEEE  
COMPUTER  
SOCIETY



IEEE



IEEE COMPUTER SOCIETY  
TFRC  
Task Force on Rebooting Computing



INTERNATIONAL ROADMAP FOR DEVICES AND SYSTEMS

© Vladimir Getov - all rights reserved

# Major Benchmarking Efforts in the Last 35 Years

- The NAS Parallel Benchmarks (NPB)
- The GENESIS Distributed-Memory Benchmarks
- The PARKBENCH International Benchmarks
- SPEComp2001 - All major machine vendors
- Another more recent “pencil and paper” parallel benchmark suite is the Dwarfs Mine based on the initial “Seven Dwarfs” proposal



# Benchmarking Methodology


- These benchmarking projects mentioned above cover predominantly legacy dense applications with high computational intensity
- Current ***application domains are different*** and cover a wider spectrum
- The ***hierarchical benchmarking approach*** has been attractive, but we now know that it is ***practically not achievable***

# Selected Benchmarks


- Two popular codes with good regularity of results covering different types of systems
- HPL – dense systems
- HPCG – sparse systems
- The earliest results for HPCG are from June 2014 while HPL results available for 25 years
- Updates published twice per year – Jun and Nov
- Using the average 10 best performance results

# PHYSICAL SYSTEM SIMULATION

HPL (dense) vs HPCG (sparse) results



HPL (dense) vs HPCG (sparse) fraction of peak performance




Gap between HPL and HPCG relatively constant over time, result of inadequate memory latency

# PHYSICAL SYSTEM SIMULATION

Technology needs

- Reduce data movement or improve memory access costs
- Improve FP arithmetic efficiency
- Reduce power consumption: energy monitoring and tuning, increase instrumentation

HPL (dense) vs HPCG (sparse) energy efficiency results in Gflop/s/W



# Analysis and Technology Needs: Memory

- Higher bandwidth and lower latency for accessing and moving data – both locally (memory systems) and remotely (interconnection networks).
- High Bandwidth Memory - HBM3+ and HBM4 expected to be released between 2022 and 2024 is likely to change substantially the application performance landscape for future supercomputers.

# Analysis and Technology Needs: Floating-Point Arithmetic

- The IEEE 754 Standard was simply renewed in July 2019
- Important aspects have been criticized: wasted cycles, energy inefficiencies, and accuracy.
- Several efforts to address these problems follow two main approaches:
  - Analysis of specific algorithms and re-writing of existing codes using mixed precision
  - More radical approaches proposing new solutions – e.g. the Posit Arithmetic proposal

# Summary

- The “Physical System Simulation” application area urgently needs novel and innovative architectures that can help address the 3<sup>rd</sup> Locality Wall.
- Energy efficiency indicators need urgent improvements by at least an order of magnitude. This is equally valid for both homogeneous vs. heterogeneous architectures (including accelerators and FPGAs).
- Since this application area is based predominantly on floating-point arithmetic, novel architecture proposals that address floating-point processing challenges can also be expected to have substantial impact, particularly for dense system computation.



IEEE  
COMPUTER  
SOCIETY



IEEE



IEEE COMPUTER SOCIETY  
TFRC

Task Force on Rebooting Computing



INTERNATIONAL ROADMAP FOR DEVICES AND SYSTEMS

© Vladimir Getov - all rights reserved



# APPENDIX: AB IFT Team Members

Name	Representing	Region
Tom Conte [co-chair]	Georgia Institute of Technology, USA	US
Natesh Ganesh	University of Colorado / NIST, USA	US
Vladimir Getov [co-chair]	University of Westminster, UK	Europe
Yoshihiro Hayashi	Keiko University, JAPAN	Japan
Masatoshi Ishii	IBM Tokyo Research Laboratory, JAPAN	Japan
Takeshi Iwashita	Hokkaido University, JAPAN	Japan
Siva Rajamanickam	Sandia National Laboratories, USA	US
Vijay Janapa Reddi	Harvard University, USA [MLperf representative]	US
Masaaki Kondo	Keio University, JAPAN	Japan
Tushar Krishna	Georgia Institute of Technology, USA	US
Peter M. Kogge	University of Notre Dame, USA	US
Scott Koziol	Baylor University, USA	US
Dam Sunwoo	Arm Research	US
Josep Torrellas	University of Illinois at Urbana-Champaign, USA	US
Peter Torelli	Chair, EEMBC [EEMBC representative]	US
Rio Yokoda	Tokyo Tech, JAPAN [SDRJ representative]	Japan

# AB IFT External Roadmap Collaborations

- EEMBC (The Embedded Microprocessor Benchmark Consortium):  
<https://www.eembc.org/>
- SDRJ (The System Device Roadmap Committee of Japan):  
<https://www.sdrj.jp/>
- MLPerf (Machine Learning Performance) Benchmarks as part of the MLCommons collaborative engineering organization :  
<https://mlcommons.org/en/>
- NIST (The National Institute of Standards and Technology):  
<https://www.nist.gov/>