OAK RIDGE National Laboratory

# Striving for SDR Performance Portability in the Era of Heterogeneous SoCs

Jeffrey S. Vetter
Seyong Lee
Mehmet Belviranli
Roberto Gioiosa
Richard Glassbrook
Abdel-Kareem Moadi

GNU Radio Conference 2019
Huntsville, Alabama
16 Sep 2019

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

Future Technologies Group ornl

http://ft.ornl.gov vetter@computer.org

# Highlights

- Motivation: Recent trends in computing paint an ambiguous future
  - Contemporary systems provide evidence that power constraints are driving architectures to change rapidly
  - Multiple architectural dimensions are being (dramatically) redesigned: Processors, node design, memory systems, I/O
  - Complexity is our main challenge
- Applications and software systems across many areas are all reaching a state of crisis
  - Applications will not be functionally or performance portable across architectures
  - Programming and operating systems need major redesign to address these architectural changes
  - Procurements, acceptance testing, and operations of today's new platforms depend on performance prediction and benchmarking.
- ORNL Cosmic project investigating design and programming challenges for these trends in SDR
  - Performance modeling and ontologies
  - Performance portable compilation to many different heterogeneous architectures/SoCs
  - Intelligent scheduling system to automate discovery, device selection, and data movement
  - Targeting wide variety of existing and future architectures (DSSoC and others)

# Motivating Trends

# Contemporary devices are approaching fundamental limits

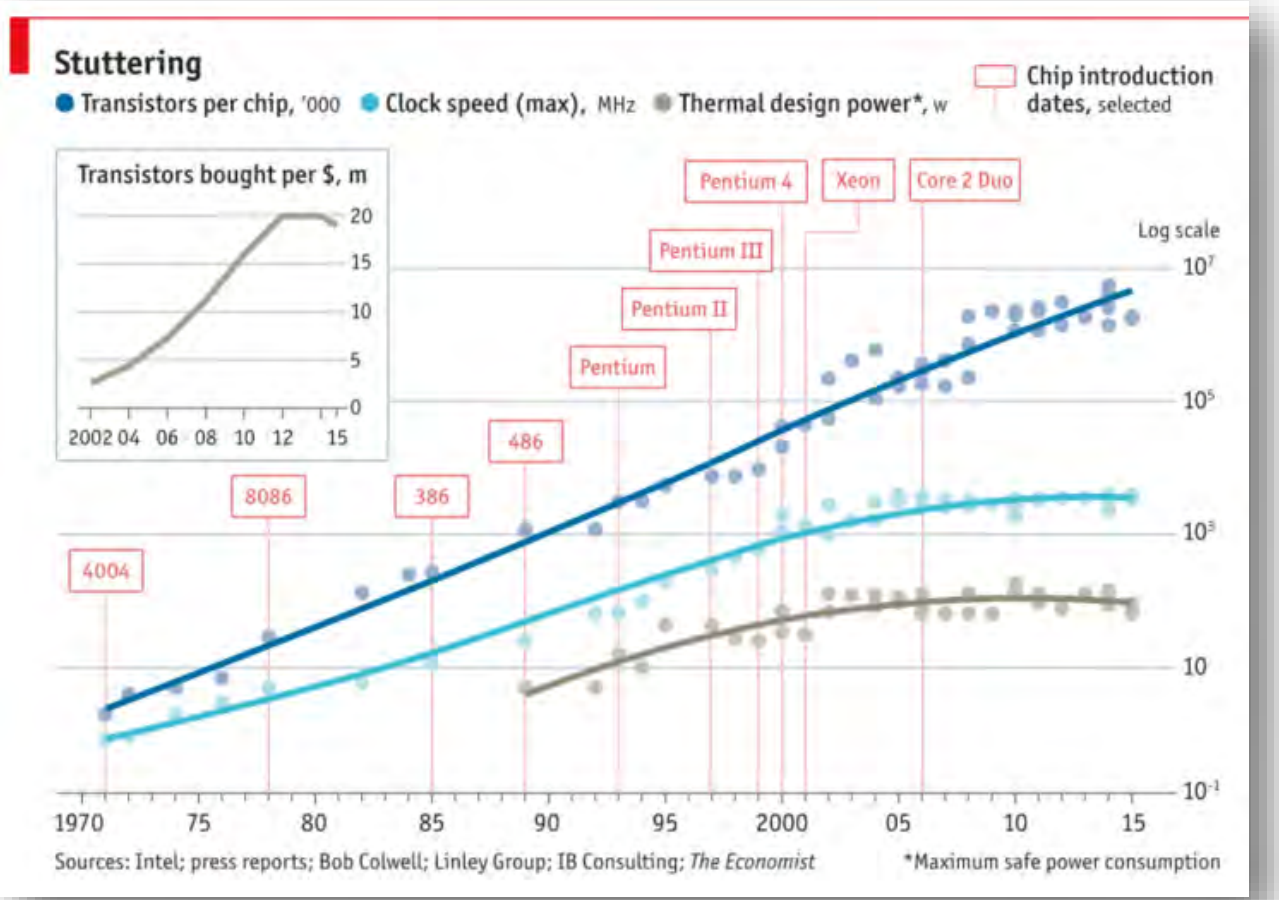


Economist, Mar 2016

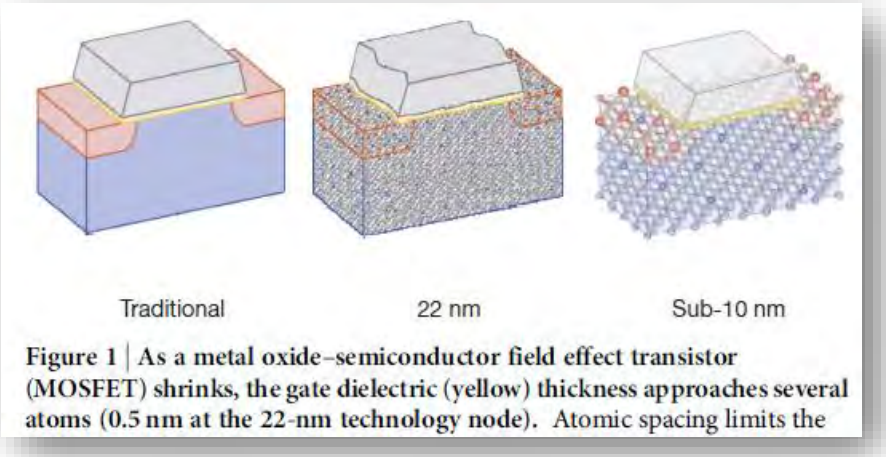


**Figure 1 | As a metal oxide–semiconductor field effect transistor (MOSFET) shrinks, the gate dielectric (yellow) thickness approaches several atoms (0.5 nm at the 22-nm technology node). Atomic spacing limits the**

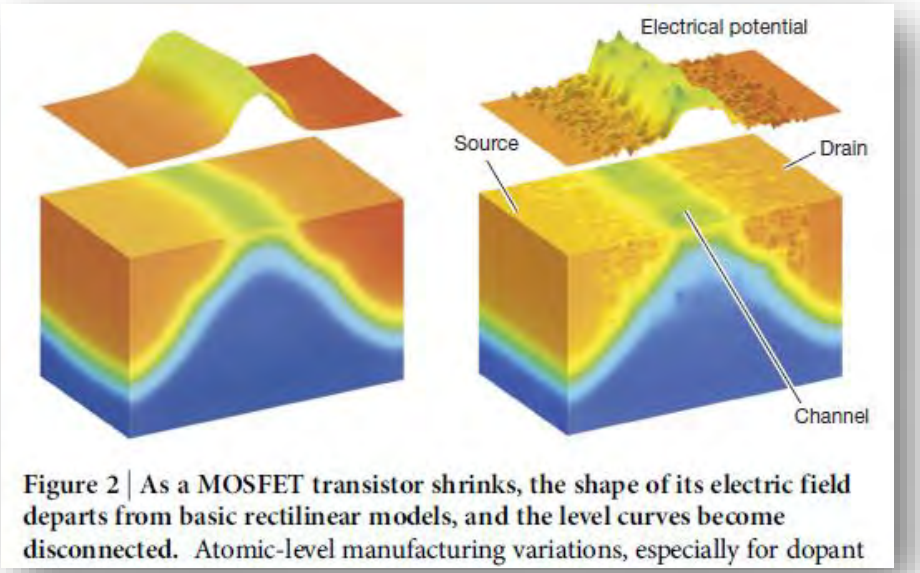


**Figure 2 | As a MOSFET transistor shrinks, the shape of its electric field departs from basic rectilinear models, and the level curves become disconnected. Atomic-level manufacturing variations, especially for dopant**

Dennard scaling has already ended. Dennard observed that voltage and current should be proportional to the linear dimensions of a transistor: 2x transistor count implies 40% faster and 50% more efficient.

R.H. Dennard, F.H. Gaensslen, V.L. Rideout, E. Bassous, and A.R. LeBlanc, "Design of ion-implanted MOSFET's with very small physical dimensions," *IEEE Journal of Solid-State Circuits, 9(5):256-68, 1974,*

I.L. Markov, "Limits on fundamental limits to computation," *Nature, 512(7513):147-54, 2014, doi:10.1038/nature13570.*

designlines AUTOMOTIVE

News & Analysis

# Foundries' Sales Show Hard Times Continuing

Peter Clarke

5/23/2016 09:33 PM EDT

2 comments

Like 6 Tweet Share 43

SEMICONDUCTOR ENGINEERING

Home > Manufacturing, Design & Test > Uncertainty Grows For 5nm, 3nm

MANUFACTURING, DESIGN & TEST

## Uncertainty Grows For 5nm, 3nm

Nanosheets and nanowire FETs under development, but costs are skyrocketing. New packaging options could

eetasia.com

## GlobalFoundries Forfeit 7nm Manufacturing - EE Times Asia

## Samsung to Invest $115 Billion in Foundry & Chip Businesses by 2030

# Intel's 10nm Is Broken, Delayed Until 2019

37 COMMENTS

by Paul Alcorn April 26, 2018 at 6:30 PM

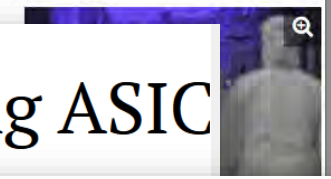

DESIGNLINES | WIRELESS AND NETWORKING DESIGNLINE

# GlobalFoundries Selling ASIC Business to Marvell

By Dylan McGrath, 05.20.19 1

Share Post

## Another Step Toward the End of Moore's Law

Samsung and TSMC move to 5-nanometer manufacturing

**Number of Foundries with a Cutting Edge Logic Fab**

| | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| SilTerra | | | | | | | | | | |
| X-FAB | | | | | | | | | | |
| Dongbu HiTek | | | | | | | | | | |
| ADI | ADI | | | | | | | | | |
| Atmel | Atmel | | | | | | | | | |
| Rohm | Rohm | | | | | | | | | |
| Sanyo | Sanyo | | | | | | | | | |
| Mitsubishi | Mitsubishi | | | | | | | | | |
| ON | ON | | | | | | | | | |
| Hitachi | Hitachi | | | | | | | | | |
| Cypress | Cypress | Cypress | | | | | | | | |
| Sony | Sony | Sony | | | | | | | | |
| Infineon | Infineon | Infineon | | | | | | | | |
| Sharp | Sharp | Sharp | | | | | | | | |
| Freescale | Freescale | Freescale | | | | | | | | |
| Renesas (NEC) | Renesas | Renesas | Renesas | Renesas | | | | | | |
| SMIC | SMIC | SMIC | SMIC | SMIC | | | | | | |
| Toshiba | Toshiba | Toshiba | Toshiba | Toshiba | | | | | | |
| Fujitsu | Fujitsu | Fujitsu | Fujitsu | Fujitsu | | | | | | |
| TI | TI | TI | TI | TI | | | | | | |
| Panasonic | Panasonic | Panasonic | Panasonic | Panasonic | Panasonic | | | | | |
| STMicroelectronics | STM | STM | STM | STM | STM | | | | | |
| UMC | UMC | UMC | UMC | UMC | UMC | | | | | |
| IBM | IBM | IBM | IBM | IBM | IBM | IBM | | | | |
| AMD | AMD | AMD | GlobalFoundries | GF | GF | GF | GF | | | |
| Samsung | Samsung | Samsung | Samsung | Samsung | Samsung | Samsung | Samsung | Samsung | Samsung | |
| TSMC | TSMC | TSMC | TSMC | TSMC | TSMC | TSMC | TSMC | TSMC | TSMC | |
| Intel | Intel | Intel | Intel | Intel | Intel | Intel | Intel | Intel | Intel | Future |
| 180 nm | 130 nm | 90 nm | 65 nm | 45 nm/40 nm | 32 nm/28 nm | 22 nm/20 nm | 16 nm/14 nm | 10 nm | 7 nm | 5 nm |

# Business climate reflects this uncertainty, cost, complexity, consolidation

## NVIDIA Buys Mellanox To Bring HPC Scaling To Data Centers

**Kevin Krewell** Contributor
**Tirias Research** Contributor Group
Enterprise & Cloud

The 2019 semiconductor merger and acquisition season has officially been kicked
technol
offer aft
and Xili
said in a
on the c
center p
perform

nytimes.com

## Hewlett Packard Enterprise to Acquire Supercomputer Pioneer Cray

5-6 minutes

Technology|Hewlett Packard Enterprise to Acquire Supercomputer Pioneer Cray

Hewlett Packard Enterprise will pay about $1.4 billion to acquire Cray, which has designed some of the most powerful computer systems in use.CreditPaco Freire/SOPA Images, via LightRocket and Getty Images


## Intel to acquire Altera for $54 a share

Monday, 1 Jun 2015

REUTE

## Broadcom acquires Brocade in $5.9 billion deal

Posted 1 hour ago by Ron Miller (@ron_miller)

## Avago Agrees to Buy Broadcom for $37 Billion

By MICHAEL J. de la MERCED and CHAD BRAY MAY 28, 2015



## Tech giant ARM Holdings sold to Japanese firm for £24bn

Britain's large
deal includin

## SoftBank to sell 25% of Arm to Saudi-backed fund

Son puts stake worth $8bn in UK's largest tech company into $100bn Vision Fund


EXCLUSIVE

## Amazon Is Becoming an AI Chip Maker, Speeding Alexa Responses

By Aaron Tilley Feb. 12, 2018 7:00 AM PST · Comments by Yonatan Raz-Fridman and Mohammad Musa

Subscribe now

Amazon.com is developing a chip designed for artificial intelligence to work on the Echo and other hardware powered by Amazon's Alexa virtual assistant, says a person familiar with Amazon's plans. The chip should allow Alexa-powered devices to respond more quickly to commands, by allowing more data processing to be handled on the device than in the cloud.

The effort makes Amazon the latest major tech company, after Google and Apple, to design its own AI chips, in hopes of differentiating their products from those of rivals. That strategy has major ramifications for chip companies like Intel and Nvidia, which are now competing with companies that previously

Britain's bigge
ARM Holding

purchased only six months ago, placing 25 per cent of Britain's largest technology company into a new, Saudi-backed $100bn investment fund.

## SANDISK COMPLETES ACQUISITION OF FUSION IO

JUL
AC

MILP
annou
hardv


lash storage solutions, today
veloper of flash-based PCIe
erscale datacenters.

"I am
the F
and c
the in

ring and go-to-market talent of
id Sanjay Mehrotra, president
t of enterprise flash solutions in

## Toshiba to sell 'minority stake' in chip business to Western Digital

In April-June 2016, Toshib

The Dollar Business B

Japanese conglomerat
performing flash mem
funds, according to a s

Earlier, in a statement
memory business, ever

The laptops-to-engine
surprised its investors
acquired in 2015, whic

## Western Digital Now A Storage Powerhouse With SanDisk Acquisition


DESIGNLINES | MEMORY DESIGNLINE

## Q1 Chip Sales Drop Among Largest on Record

By Dylan McGrath, 05.01.19 0

Share Post

SAN FRANCISCO — Global chip sales sank by 15.5% sequentially in the first quart
among the largest quarter-to-quarter declines for the industry in the last 35 years

Chip sales totaled $96.8 billion in the first quarter, down from $114.7 billion last y
according to the World Semiconductor Trade Statistics (WSTS) organization, whic
sales data from chipmaker member companies. On a year-over-year basis, first-qu

# Sixth Wave of Computing

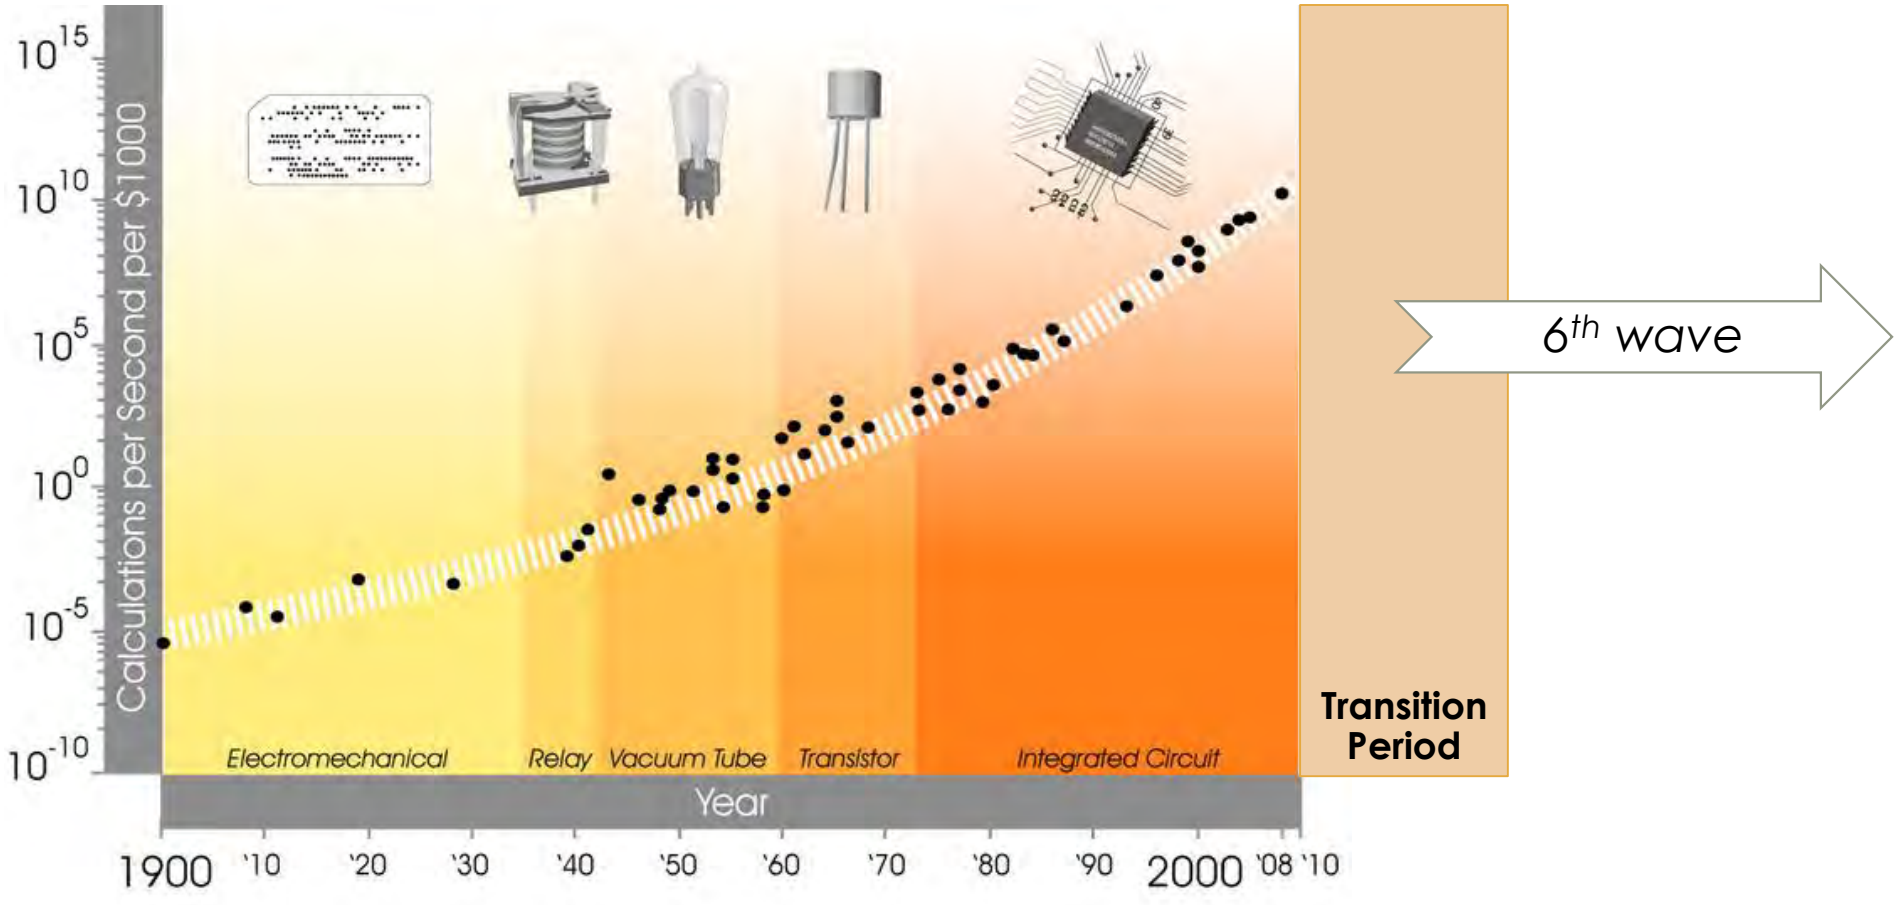


http://www.kurzweilai.net/exponential-growth-of-computing

# Predictions for Transition Period

## Optimize Software and Expose New Hierarchical Parallelism

- Redesign software to boost performance on upcoming architectures
- Exploit new levels of parallelism and efficient data movement

## Architectural Specialization and Integration

- Use CMOS more effectively for specific workloads
- Integrate components to boost performance and eliminate inefficiencies
- Workload specific memory+storage system design

## Emerging Technologies

- Investigate new computational paradigms
  - Quantum
  - Neuromorphic
  - Advanced Digital
  - Emerging Memory Devices

# Predictions for Transition Period

## Optimize Software and Expose New Hierarchical Parallelism

- Redesign software to boost performance on upcoming architectures
- Exploit new levels of parallelism and efficient data movement

## Architectural Specialization and Integration

- Use CMOS more effectively for specific workloads
- Integrate components to boost performance and eliminate inefficiencies
- Workload specific memory+storage system design

## Emerging Technologies

- Investigate new computational paradigms
  - Quantum
  - Neuromorphic
  - Advanced Digital
  - Emerging Memory Devices

# Predictions for Transition Period

## Optimize Software and Expose New Hierarchical Parallelism

- Redesign software to boost performance on upcoming architectures
- Exploit new levels of parallelism and efficient data movement

## Architectural Specialization and Integration

- Use CMOS more effectively for specific workloads
- Integrate components to boost performance and eliminate inefficiencies
- Workload specific memory+storage system design

## Emerging Technologies

- Investigate new computational paradigms
  - Quantum
  - Neuromorphic
  - Advanced Digital
  - Emerging Memory Devices

# Predictions for Transition Period

## Optimize Software and Expose New Hierarchical Parallelism

- Redesign software to boost performance on upcoming architectures
- Exploit new levels of parallelism and efficient data movement

## Architectural Specialization and Integration

- Use CMOS more effectively for specific workloads
- Integrate components to boost performance and eliminate inefficiencies
- Workload specific memory+storage system design

## Emerging Technologies

- Investigate new computational paradigms
  - Quantum
  - Neuromorphic
  - Advanced Digital
  - Emerging Memory Devices

# Pace of Architectural Specialization is Quickening

- Industry, lacking Moore's Law, will need to continue to differentiate products (to stay in business)
  - Use the same transistors differently to enhance performance
- Architectural design will become extremely important, critical
  - Dark Silicon
  - Address new parameters for benefits/curse of Moore's Law
- 50+ new companies focusing on hardware for Machine Learning


http://www.theinquirer.net/inquirer/news/2477796/intels-nervana-ai-platform-takes-aim-at-nvidias-gpu-techology


http://www.wired.com/2016/05/google-tpu-custom-chips/

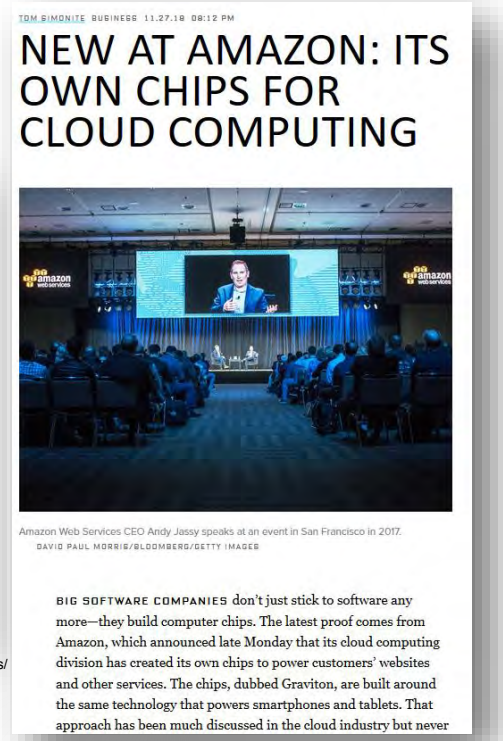



D.E. Shaw, M.M. Deneroff, R.O. Dror *et al., "Anton, a special-purpose machine for molecular dynamics simulation," Communications of the ACM, 51(7):91-7, 2008.*

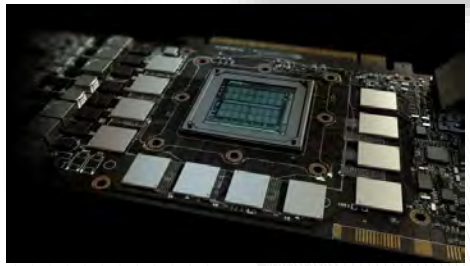

https://fossbytes.com/nvidia-volta-gddr6-2018/

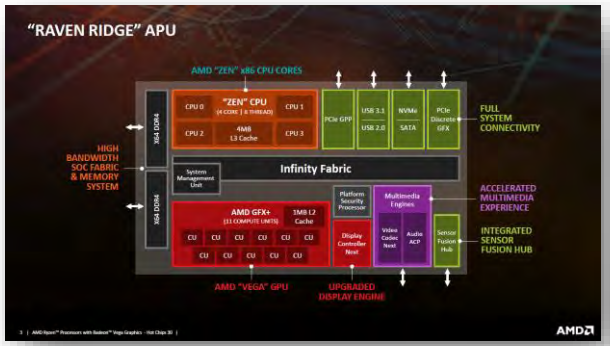


HotChips 2018


HotChips 2018

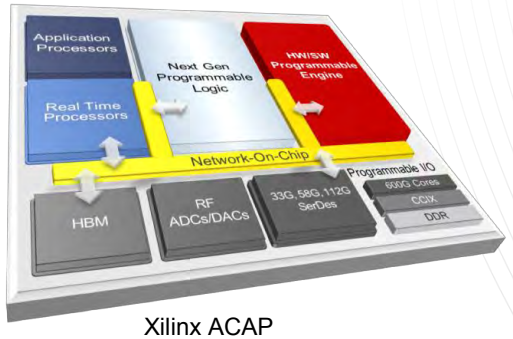


Xilinx ACAP


https://www.thebroadcastbridge.com/content/entry/1094/altera-announces-arria-10-gx-sx-soc-fpgas

# Analysis of Apple A-* SoCs

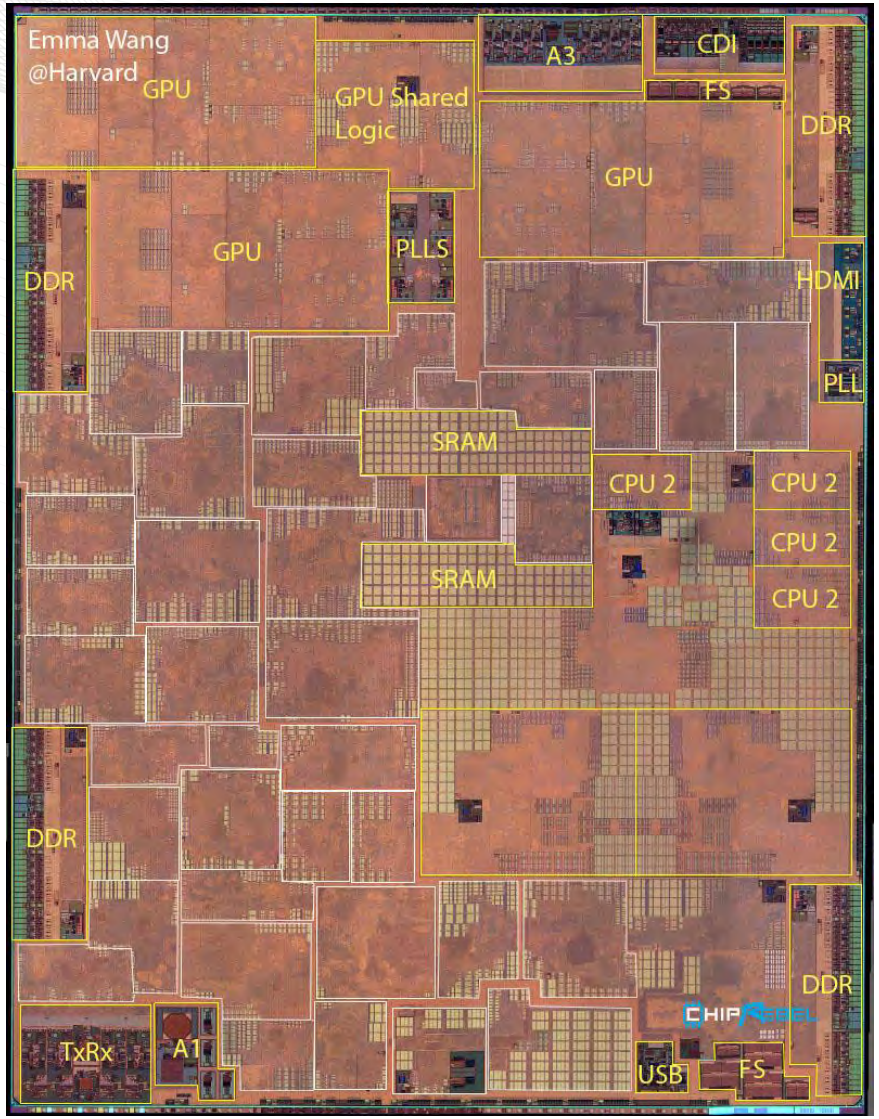


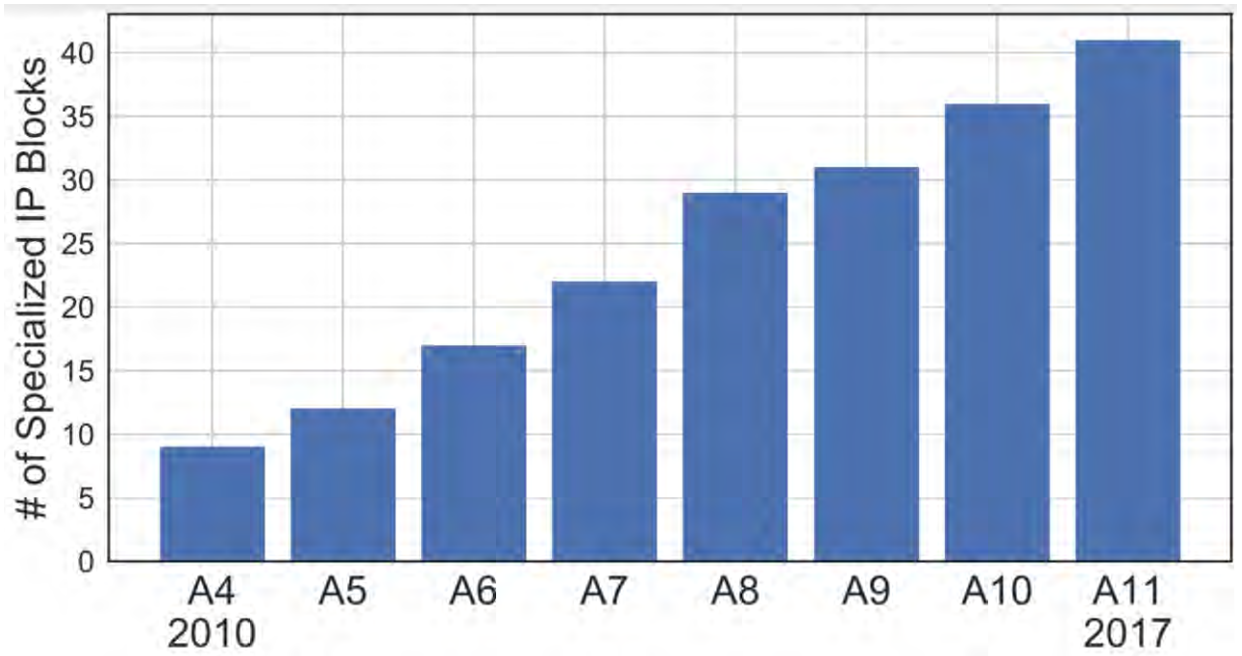


http://vlsiarch.eecs.harvard.edu/accelerators/die-photo-analysis

# Growing Open Source Hardware Movement Enables Rapid Chip Design

## RISC-V Ecosystem

**Software**

**Open-source software:**
Gcc, binutils, glibc, Linux, BSD, LLVM, QEMU, FreeRTOS, ZephyrOS, LiteOS, SylixOS, …

**Commercial software:**
Lauterbach, Segger, Micrium, ExpressLogic, …

**RISC-V Foundation**

ISA specification | Golden Model | Compliance

**Hardware**

**Open-source cores:**
Rocket, BOOM, RI5CY, Ariane, PicoRV32, Piccolo, SCR1, Hummingbird, …

**Commercial core providers:**
Andes, Bluespec, Cloudbear, Codasip, Cortus, C-Sky, Nuclei, SiFive, Syntacore, …

**Inhouse cores:**
Nvidia, +others

7

# Summary:
# Transition Period will be Disruptive – Opportunities and Pitfalls Abound

- New devices and architectures may not be hidden in traditional levels of abstraction

- Examples
  - A new type of CNT transistor may be completely hidden from higher levels
  - A new paradigm like quantum may require new architectures, programming models, and algorithmic approaches

| Layer | Switch, 3D | NVM | Approximate | Neuro | Quantum |
|---|---|---|---|---|---|
| *Application* | 1 | 1 | 2 | 2 | 3 |
| *Algorithm* | 1 | 1 | 2 | 3 | 3 |
| *Language* | 1 | 2 | 2 | 3 | 3 |
| *API* | 1 | 2 | 2 | 3 | 3 |
| *Arch* | 1 | 2 | 2 | 3 | 3 |
| *ISA* | 1 | 2 | 2 | 3 | 3 |
| *Microarch* | 2 | 3 | 2 | 3 | 3 |
| *FU* | 2 | 3 | 2 | 3 | 3 |
| *Logic* | 3 | 3 | 2 | 3 | 3 |
| *Device* | 3 | 3 | 2 | 3 | 3 |

Adapted from IEEE Rebooting Computing Chart

# The Summit System @ ORNL

#1 on Top 500 since June 2018

## System Performance

- Peak of 200 Petaflops ($FP_{64}$) for modeling & simulation
- Peak of 3.3 ExaOps ($FP_{16}$) for data analytics and artificial intelligence
- Max power 13 MW

## The system includes

- 4,608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM file system transferring data at 2.5 TB/s

## Each node has

- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory
  (96 GB HBM2 + 512 GB DDR4)
- 1.6 TB of NV memory

# Frontier | 7 May 2019

## U.S. Department of Energy and Cray to Deliver Record-Setting Frontier Supercomputer at ORNL

### Exascale system expected to be world's most powerful computer for science and innovation

**Topic:** Supercomputing

May 7, 2019


OAK RIDGE, Tenn., May 7, 2019—The U.S. Department of Energy today anr contract with Cray Inc. to build the Frontier supercomputer at Oak Ridge Nati which is anticipated to debut in 2021 as the world's most powerful computer v performance of greater than 1.5 exaflops.

Scheduled for delivery in 2021, Frontier will accelerate innovation in science and technology and maintain U.S. leadership in high-performance computing and artificial intelligence. The total contract award is valued at more than $600 million for the system and technology development. The system will be based on Cray's new Shasta architecture and Slingshot interconnect and will feature high-performance AMD EPYC CPU and AMD Radeon Instinct GPU technology.

| | |
|---|---|
| Peak Performance | >1.5 EF |
| Footprint | > 100 cabinets |
| Node | 1 HPC and AI Optimized AMD EPYC CPU<br>4 Purpose Built AMD Radeon Instinct GPU |
| CPU-GPU Interconnect | AMD Infinity Fabric<br>Coherent memory across the node |
| System Interconnect | Multiple Slingshot NICs providing 100 GB/s network bandwidth<br>Slingshot dragonfly network which provides adaptive routing, congestion management and quality of service. |
| Storage | 2-4x performance and capacity of Summit's I/O subsystem. Frontier will have near node storage like Summit. |

# Department of Energy (DOE) Roadmap to Exascale Systems

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission

**Pre-Exascale Systems** [Aggregate Linpack (Rmax) = 323 PF!]

**First U.S. Exascale Systems**

| 2012 | 2016 | 2018 | 2020 | 2021-2023 |
|---|---|---|---|---|


**ORNL**
Cray/AMD/NVIDIA


**ORNL**
IBM/NVIDIA

**Heterogeneous Cores**


**ORNL**
AMD/Cray


**ANL**
IBM BG/Q


**ANL**
Cray/Intel KNL

**Deep Memory incl NVM**


**ANL**
Intel/Cray


**LBNL**
Cray/Intel Xeon/KNL


**LBNL**
Cray/AMD/NVIDIA

**Plateauing I/O Performance**


**LLNL**
IBM BG/Q


**LANL/SNL**
Cray/Intel Xeon/KNL


**LLNL**
IBM/NVIDIA


**LANL/SNL**
TBD


**LLNL**
TBD

OAK RIDGE National Laboratory

Jan 2018

# Domain Specific System on Chip (DSSoC) Program to address these challenges
## Performer domains and applications

**IBM T. J. Watson Research Center**
Pradip Bose
Columbia University, Harvard University, Univ. of Illinois at Urbana-Champaign

CV+SDR
- Multi-domain application
- Multi-spectral processing
- Communications

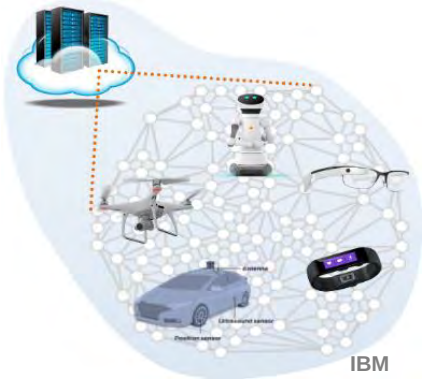

IBM

**Arizona State University**
Daniel W. Bliss
Univ. of Michigan, Carnegie Mellon University, General Dynamic Mission Systems, Arm Ltd., EpiSys Science

SDR
- Unmanned aerial
- Small robotic & leave-behind
- Universal soldier systems
- Multifunction systems

PlastyForma

**Raytheon/Xilinx**
Tom Kazior

SDR
- Xilinx ACAP
- Visual system integrator
- Improved reconfigurability of processing

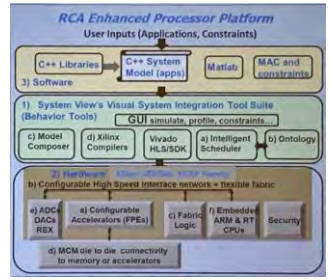


Raytheon

**Stanford University**
Mark Horowitz
Clark Barrett, Kayvon Fatahalian, Pat Hanrahan, Priyanka Raina

Computer Vision
- Still image and video processing
- Autonomous navigation
- Continuous surveillance
- Augmented reality

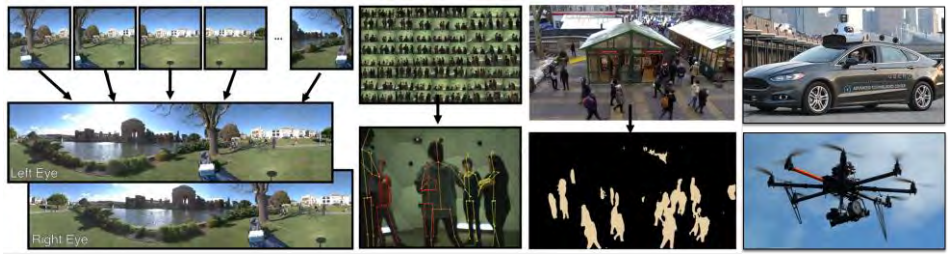


Stanford



Google/YouTube

**Oak Ridge National Laboratory**
Jeffrey Vetter

SDR
- Communications and signal processing focused
- Up-front processing / data cutdown
- Improving understanding of processing systems

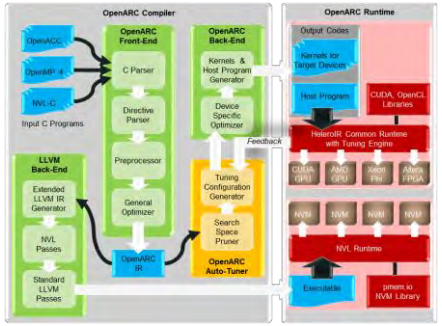


ORNL



*Slide courtesy of Dr. Tom Rondeau, DARPA MTO*

# ORNL Cosmic Project

# Cosmic Castle | Development Lifecycle

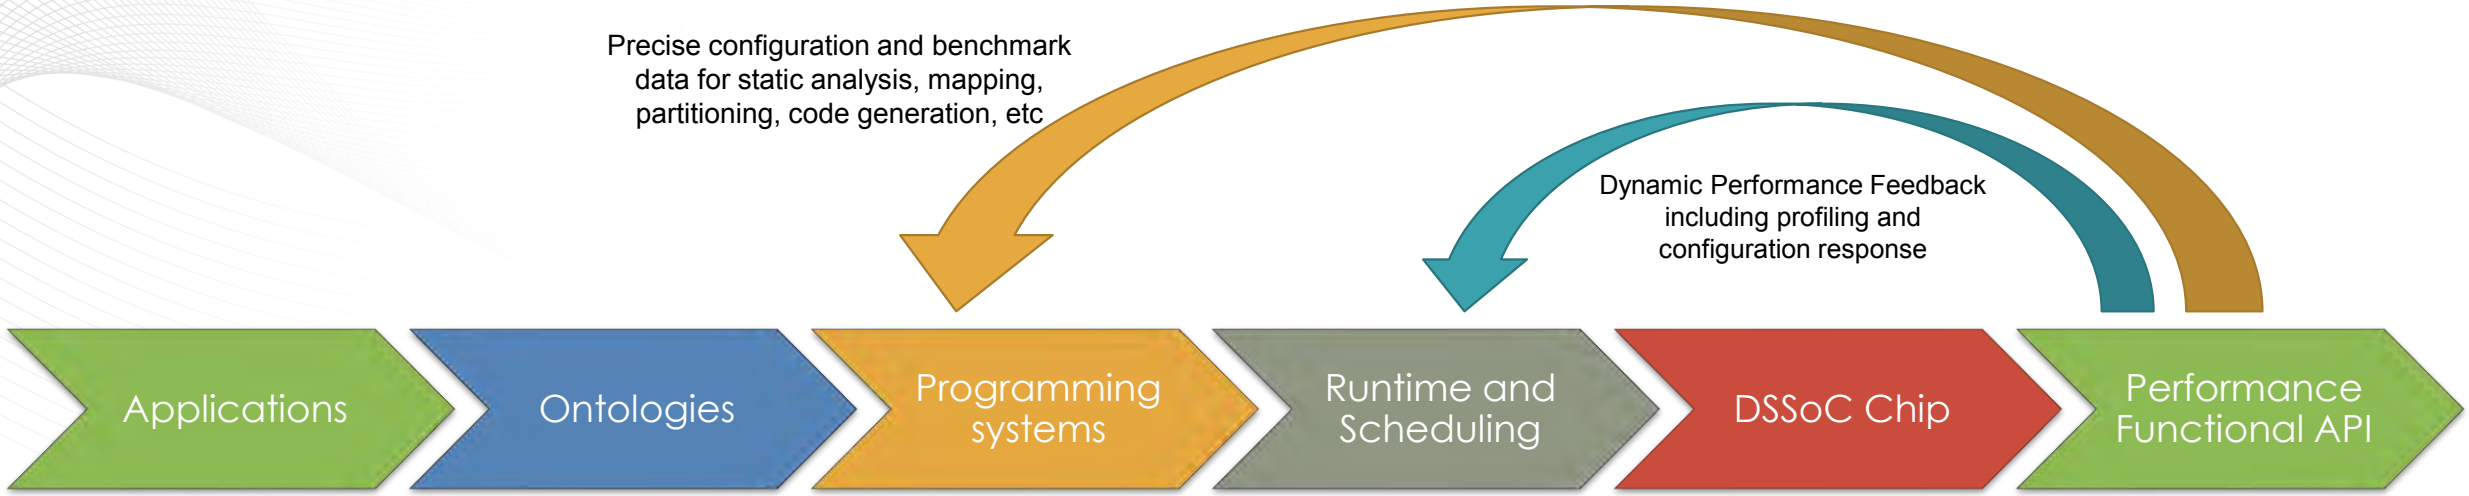


- Create scalable application Aspen models manually, with static or dynamic analysis, or using historical information
- Ontologies based on Aspen models using statistical and machine learning techniques
- Programming systems built to support ontologies
- Query Aspen models and PFU for automatic code generation, optimization, etc.
- Intelligent runtime scheduling uses models and PFU to inform dynamic decisions
- Dynamic resource discovery and monitoring
- DSSoC design quantitatively derived from application Aspen models
- Early design space exploration with Aspen
- As feature of DSSoC, PFU API provides dynamic performance response of deployed DSSoC to intelligent runtime and programming system.


- DSSoC design quantitatively derived from application Aspen models
- Early design space exploration with Aspen

# Cosmic Castle | Project Overview

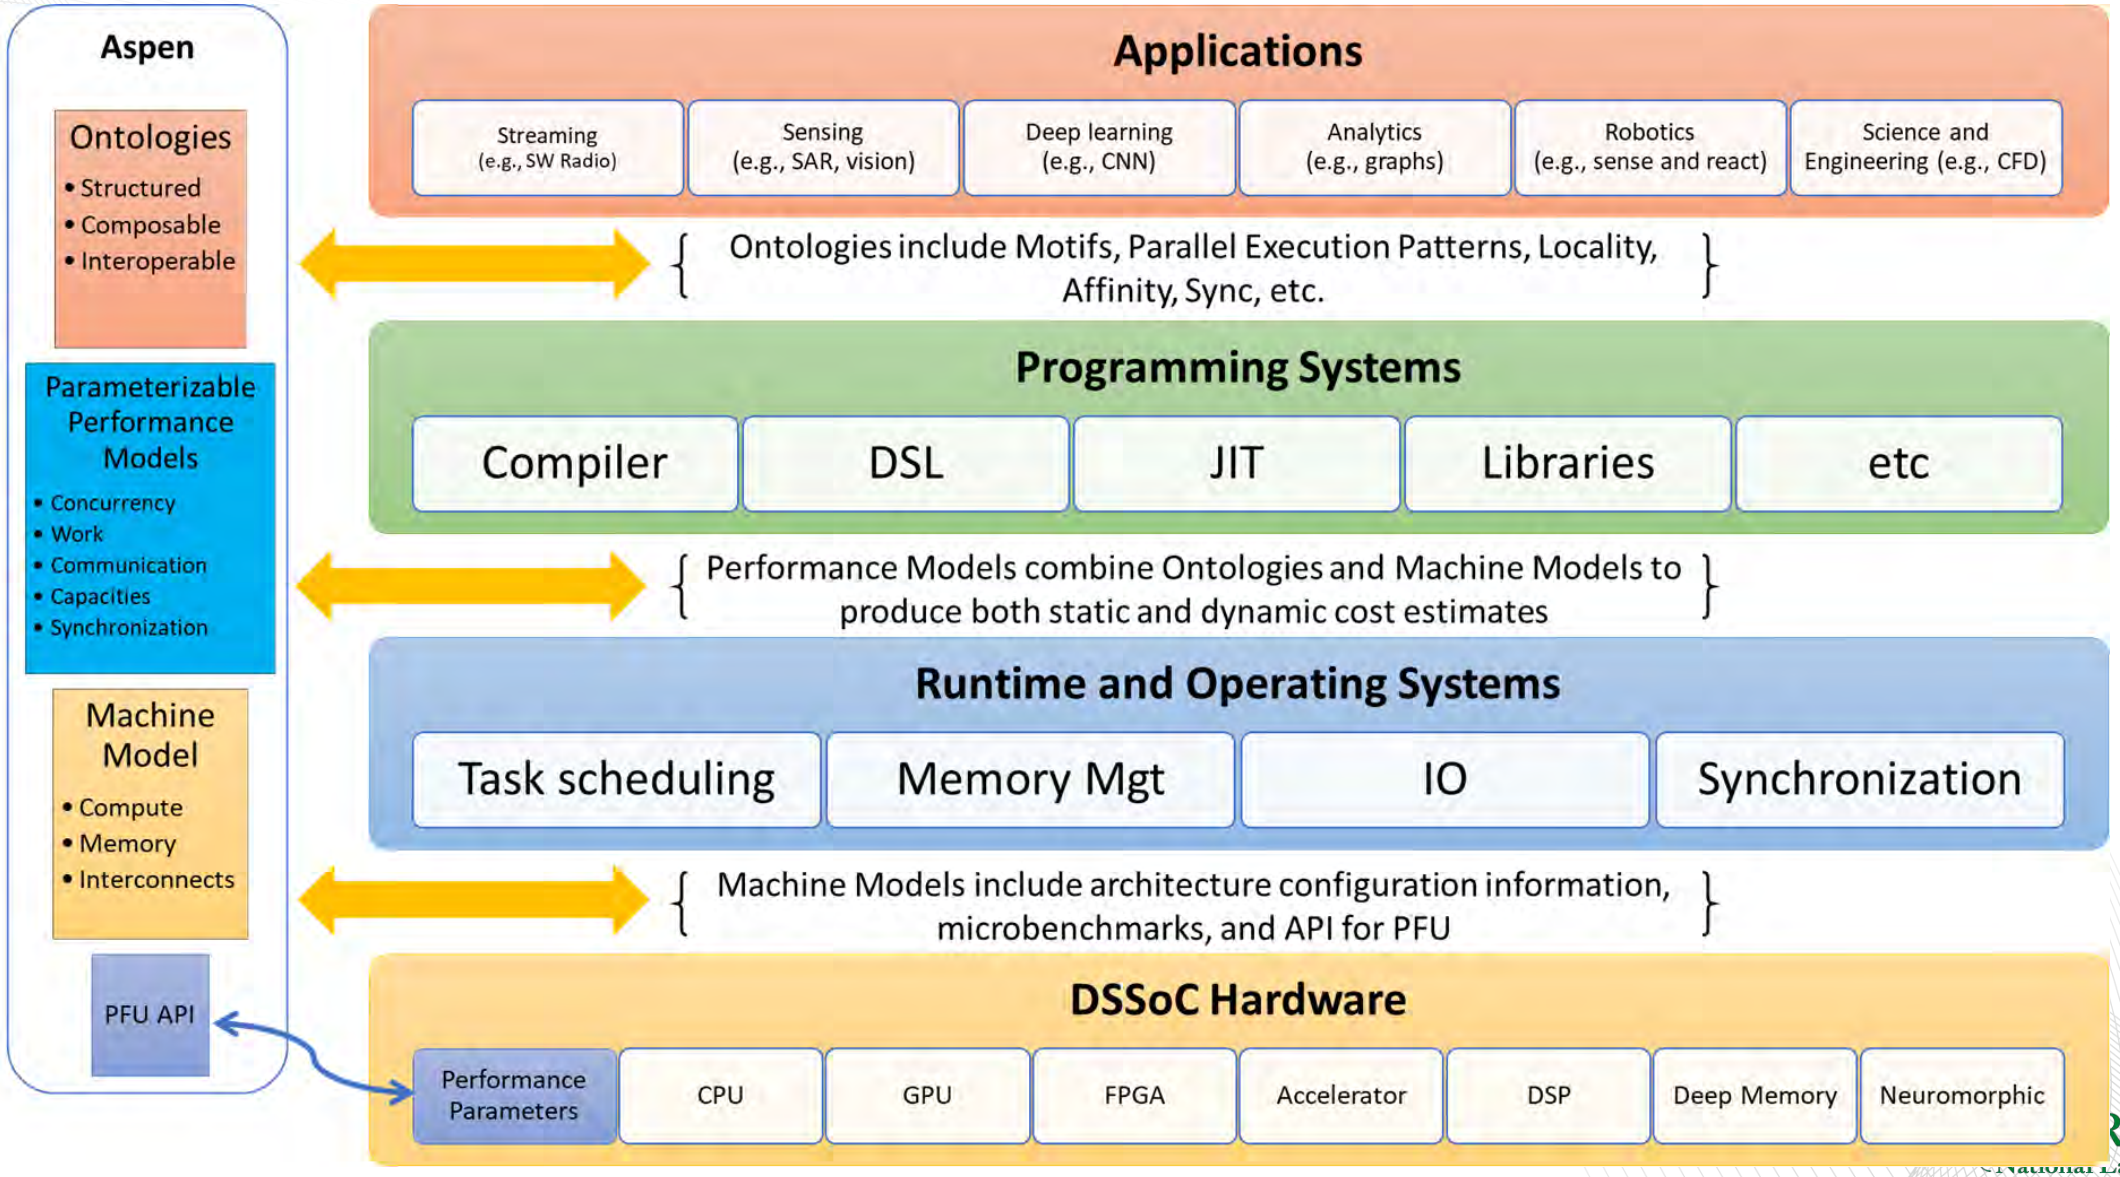

# Cosmic Castle | Project Overview

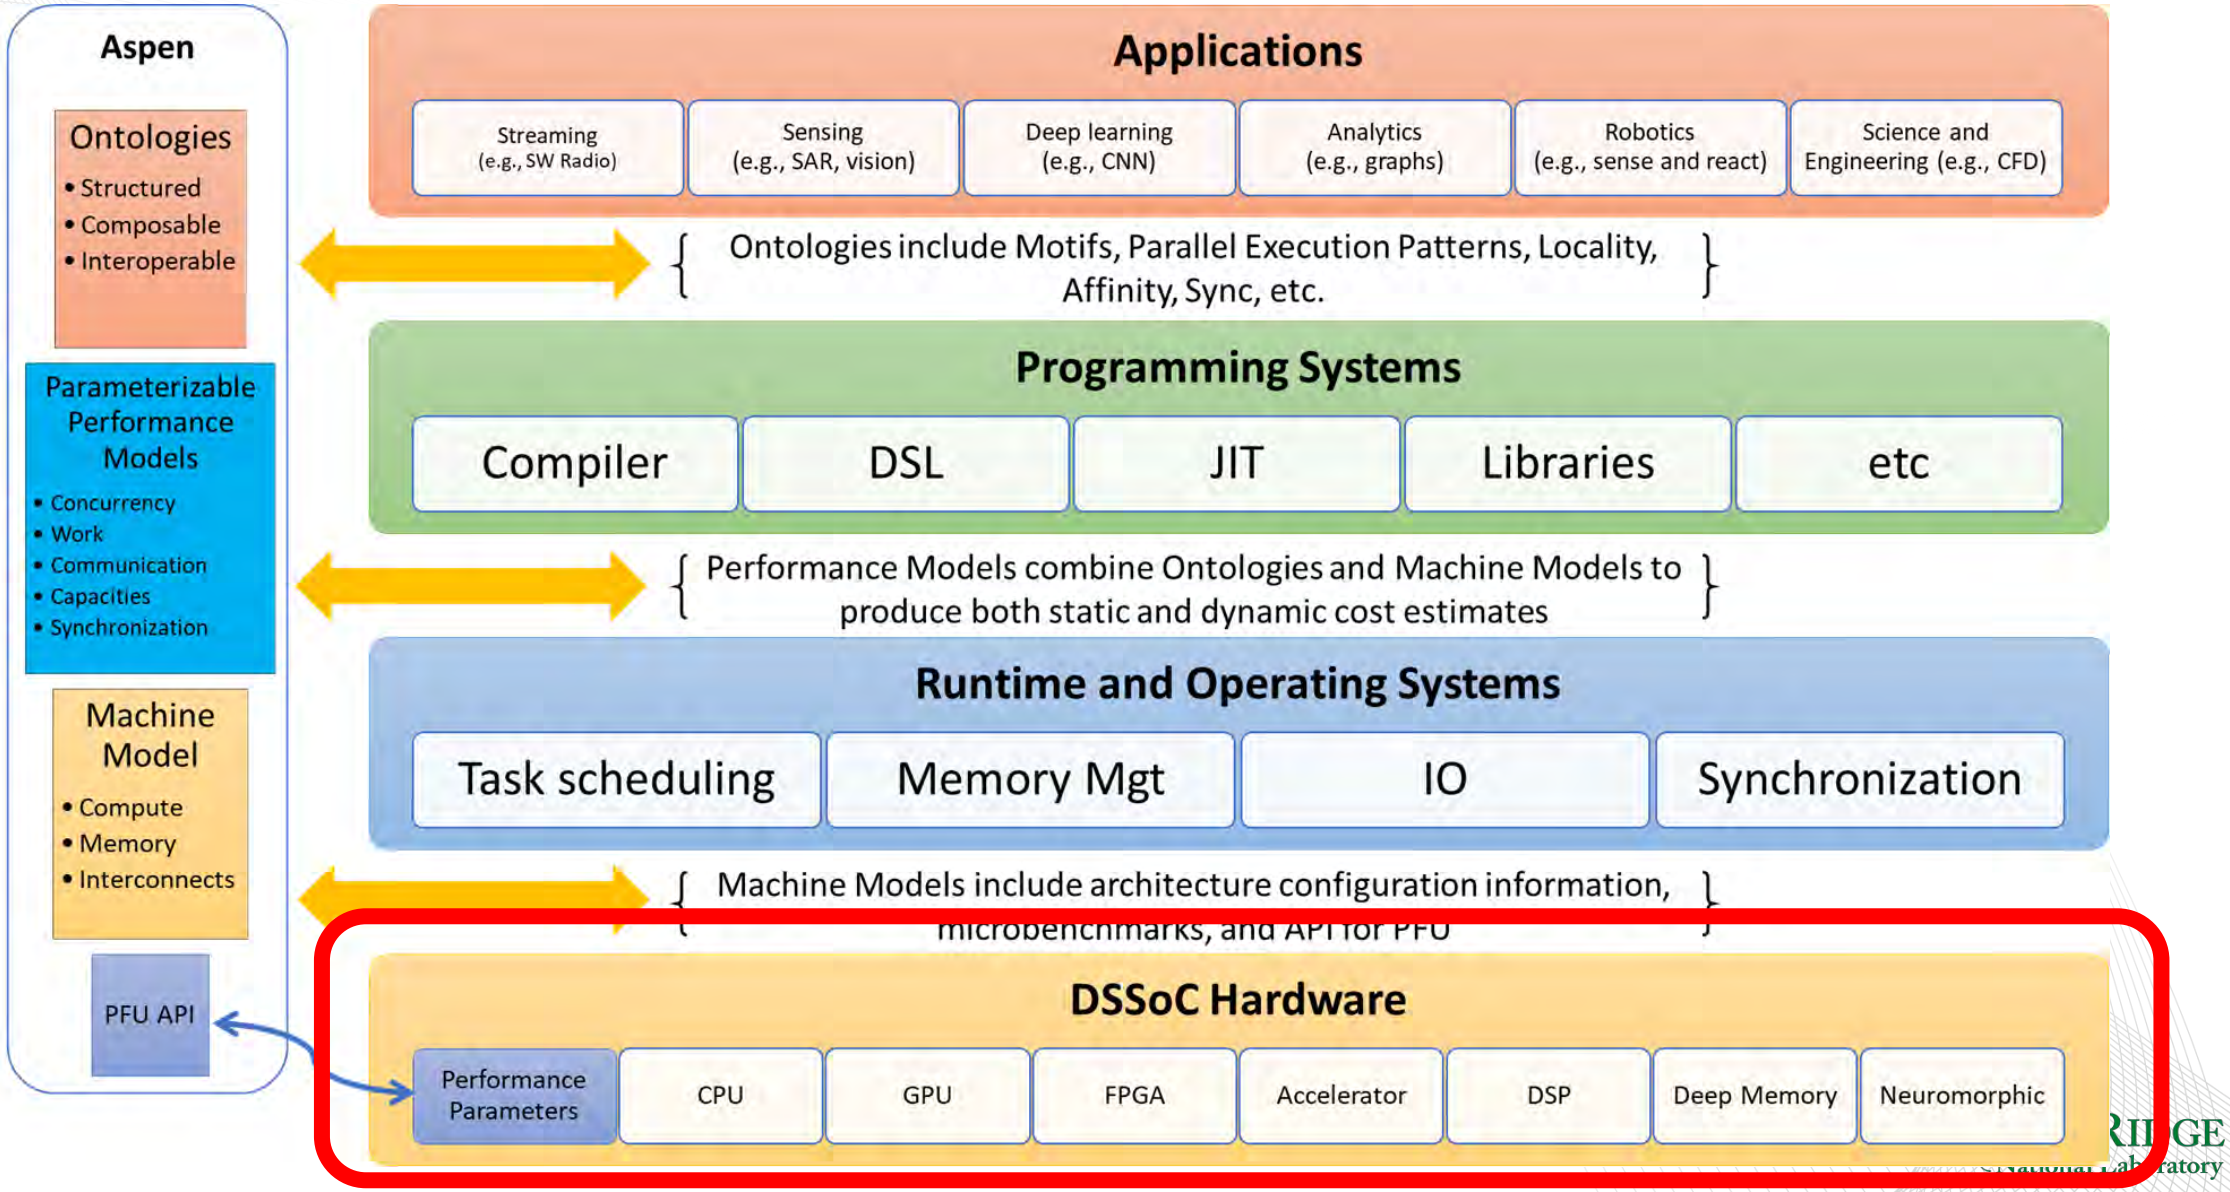

## Summit (IBM POWER9+NVIDIA Volta) Node installed

**Experimental Computing Lab (ExCL) managed by the ORNL Future Technologies Group**

IBM Summit Node with 6 Nvidia Tesla V100 GPUs (8335-GTX)

- Same CPU/GPU/Memory as nodes in OLCF Summit
  - 2 Power9 CPUs (IBM 02CY209)
    - 22 Cores each, 4 threads/core
  - 606GiB main memory
  - 6 Tesla V100 SXM2 16GB GPUs
- Provides a development and evaluation environment for Power9/V100 GPUs
- Tracks (as closely as possible) the software stack in use on Summit
- Shared / Queued / Single User availability modes will be available

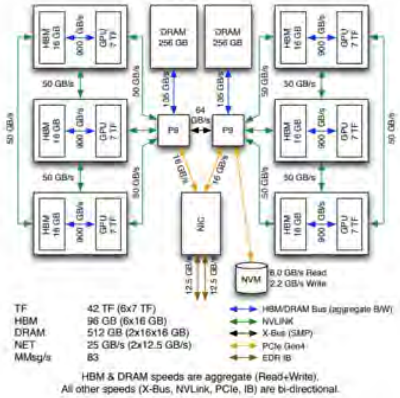


## NVIDIA DGX Workstation Available

**Experimental Computing Lab (ExCL)** managed by the ORNL Future Technologies Group

- 4X Tesla V100 GPUs
- TFLOPS (Mixed precision) 500
- GPU Memory 128 GB total system
- NVIDIA Tensor Cores 2,560
- NVIDIA CUDA® Cores 20,480
- CPU Intel Xeon E5-2698 v4 2.2 GHz (20-Core)
- System Memory 256 GB RDIMM DDR4
- Full NVIDIA stack
- Other compilers/tools installable on request



## AMD Radeon VII Available

**Experimental Computing Lab (ExCL)** managed by the ORNL Future Technologies Group

- AMD Radeon VII, Vega 20 Architecture
  - GCN 5 on TSMC 7FF process, 13.2B transistors
  - 60 Compute Units with 3.4 DP peak TF
  - 16 GB HBM2 with 4096-bit width for ~1TBps bandwidth
  - TBP 300W
  - PCIe 3.0 x16
- Intel Xeon Skylake Host
  - HP Z4 G4 Workstation w/ PCIe 3.0 x16
  - W-2123 / 64Gb host
  - 1 CPU * 4 cores * 2 threads/core
  - 512 GB SSD uncommitted/available
- Software
  - AMD ROCm development tools
  - HIP (Heterogeneous Compute Interface for Portability) available
  - OpenCL 2.1
- Additional Details
  - https://www.anandtech.com/show/13832/amd-radeon-vii-high-end-7nm-february-7th-for-699
  - https://en.wikipedia.org/wiki/AMD_RX_Vega_series#cite_note-anand_radeon_vii





## ARM ThunderX2 Node Available

**Experimental Computing Lab (ExCL)** managed by the ORNL Future Technologies Group

ThunderX2 Workstation

- Cavium (Marvell) ThunderX2 with ARMv8.1 instruction set.
- 2 Cpus, each with 28 Cores with 4 threads/core
- 128 GiB Main Memory
- Gigabyte MT91-FS1-00 motherboard
- Multiple access levels available to researchers investigating ARM8v1 performance
- Traditional ARM/Linux software stack available


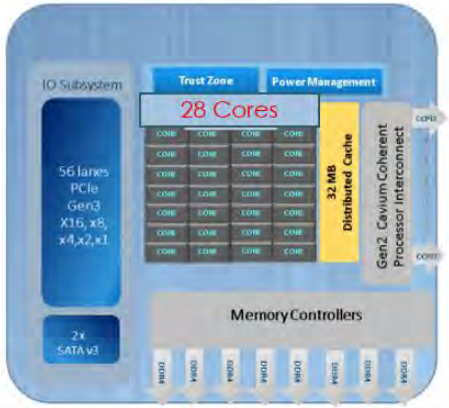

# Intel Stratix 10 FPGA available

**Experimental Computing Lab (ExCL) managed by the ORNL Future Technologies Group**

- Intel Stratix 10 FPGA and four banks of DDR4 external memory
  - Board configuration: Nallatech 520 Network Acceleration Card
- Up to 10 TFLOPS of peak single precision performance
- 25MBytes of L1 cache @ up to 94 TBytes/s peak bandwidth
- 2X Core performance gains over Arria 10
- Quartus and OpenCL software (Intel SDK v18.1) for using FPGA
- Provide researcher access to advanced FPGA/SOC environment


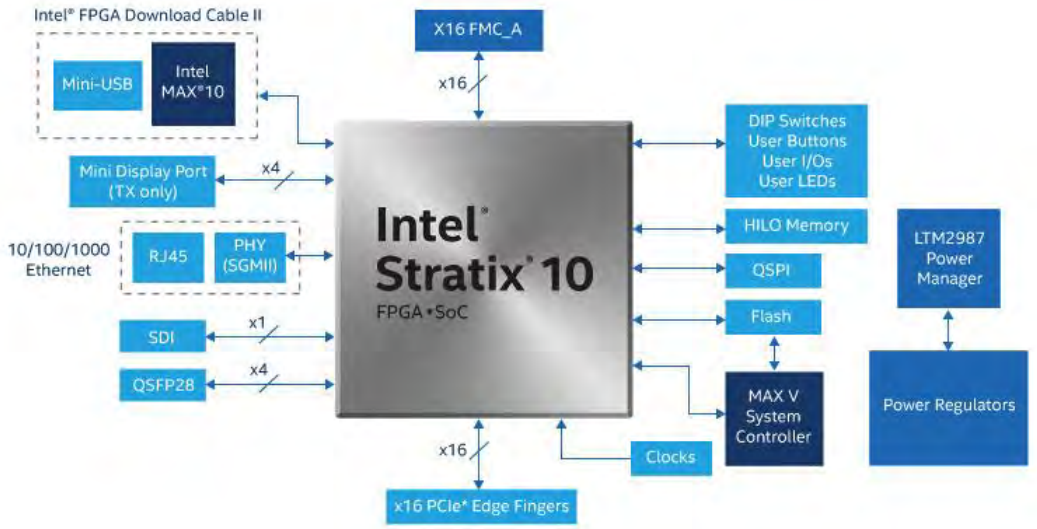




For more information or to apply for an account, visit https://excl.ornl.gov/

# NVIDIA Jetson AGX Xavier SoC available

**Experimental Computing Lab (ExCL) managed by the ORNL Future Technologies Group**

NVIDIA Jetson AGX Xavier:

- High-performance system on a chip for autonomous machines
- Heterogeneous SoC contains:
  - Eight-core 64-bit ARMv8.2 CPU cluster (Carmel)
  - 1.4 CUDA TFLOPS (FP32) GPU with additional inference optimizations (Volta)
  - 11.4 DL TOPS (INT8) Deep learning accelerator (NVDLA)
  - 1.7 CV TOPS (INT8) 7-slot VLIW dual-processor Vision accelerator (PVA)
  - A set of multimedia accelerators (stereo, LDC, optical flow)
- Provides researchers access to advanced high-performance SOC environment

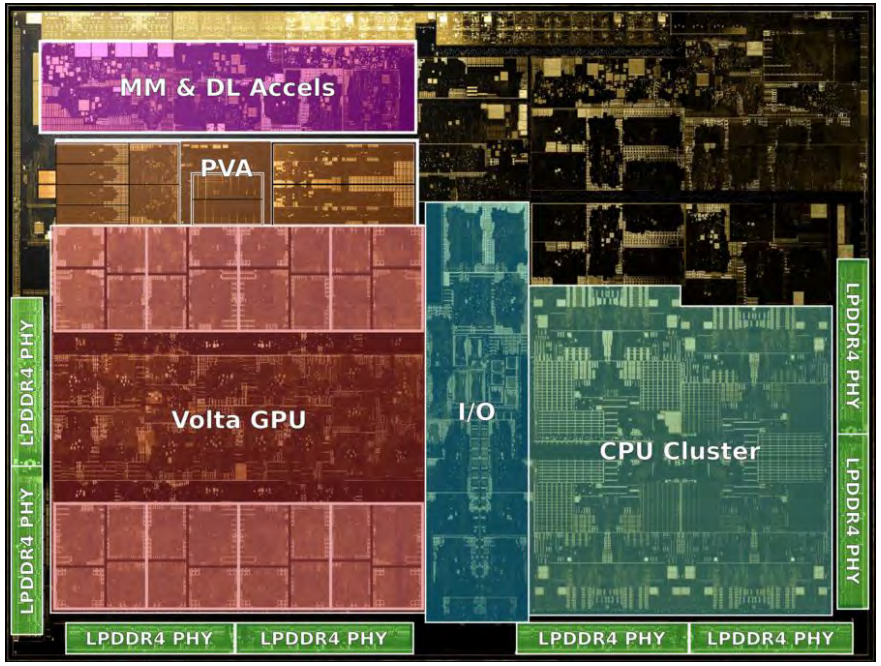


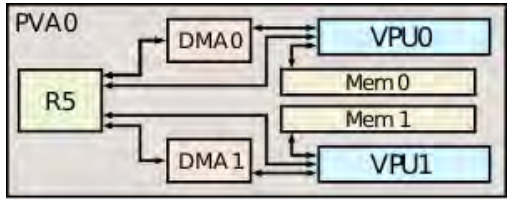


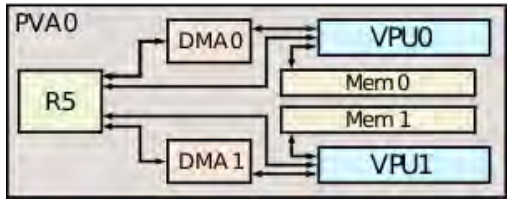




For more information or to apply for an account, visit https://excl.ornl.gov/

# Qualcomm 855 SoC (SM8510P)

**Experimental Computing Lab (ExCL) managed by the ORNL Future Technologies Group**


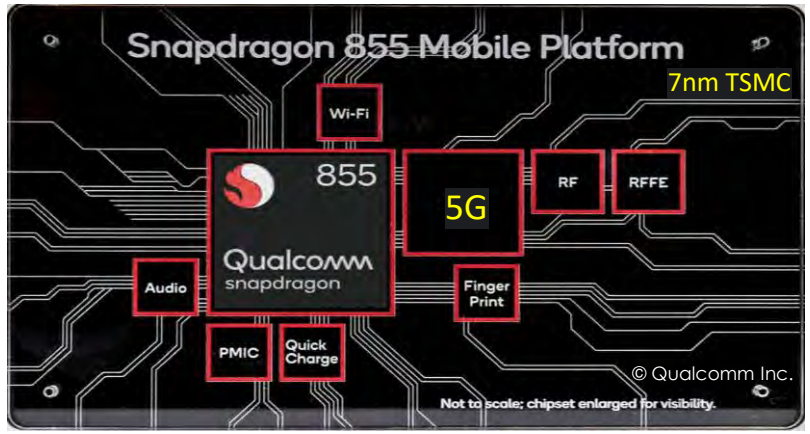


## Kyro 485 (8-ARM Prime+BigLittle Cores)


## Hexagon 690 (DSP + AI)

- Quad threaded Scalar Core
- DSP + 4 Hexagon Vector Xccelerators
- New Tensor Xccelerator for AI
- Apps: AI, Voice Assistance, AV codecs

## Connectivity (5G)

- Snapdragon X24 LTE (855 built-in) modem LTE Category 20
- Snapdragon X50 5G (external) modem (for 5G devices)
- Qualcomm Wi-Fi 6-ready mobile platform: (802.11ax-ready, 802.11ac Wave 2, 802.11ay, 802.11ad)
- Qualcomm 60 GHz Wi-Fi mobile platform: (802.11ay, 802.11ad)
- Bluetooth Version: 5.0
- Bluetooth Speed: 2 Mbps
- High accuracy location with dual-frequency GNSS.

## Adreno 640

- Vulkan, OpenCL, OpenGL ES 3.1
- Apps: HDR10+, HEVC, Dolby, etc
- Enables 8k-360° VR video playback
- 20% faster compared to Adreno 630

## Spectra 360 ISP

- New dedicated Image Signal Processor (ISP)
- Dual 14-bit CV-ISPs; 48MP @ 30fps single camera
- Hardware CV for object detection, tracking, streo depth process
- 6DoF XR Body tracking, H265, 4K60 HDR video capture, etc.

## Qualcomm Development Board connected to (mcmurdo) HPZ820

- Connected Qualcomm board to HPZ820 through USB
- Development Environment: Android SDK/NDK
- Login to mcmurdo machine
  $ ssh –Y mcmurdo
- Setup Android platform tools and development environment
  $ source /home/nqx/setup_android.source
- Run Hello-world on ARM cores
  $ git clone https://code.ornl.gov/nqx/helloworld-android
  $ make compile push run
- Run OpenCL example on GPU
  $ git clone https://code.ornl.gov/nqx/opencl-img-processing
  - Run Sobel edge detection
    $ make compile push run fetch
- Login to Qualcomm development board shell
  $ adb shell
  $ cd /data/local/tmp



For more information or to apply for an account, visit https://excl.ornl.gov/



Created by Narasinga Rao Miniskar, Steve Moulton

# RISC-V

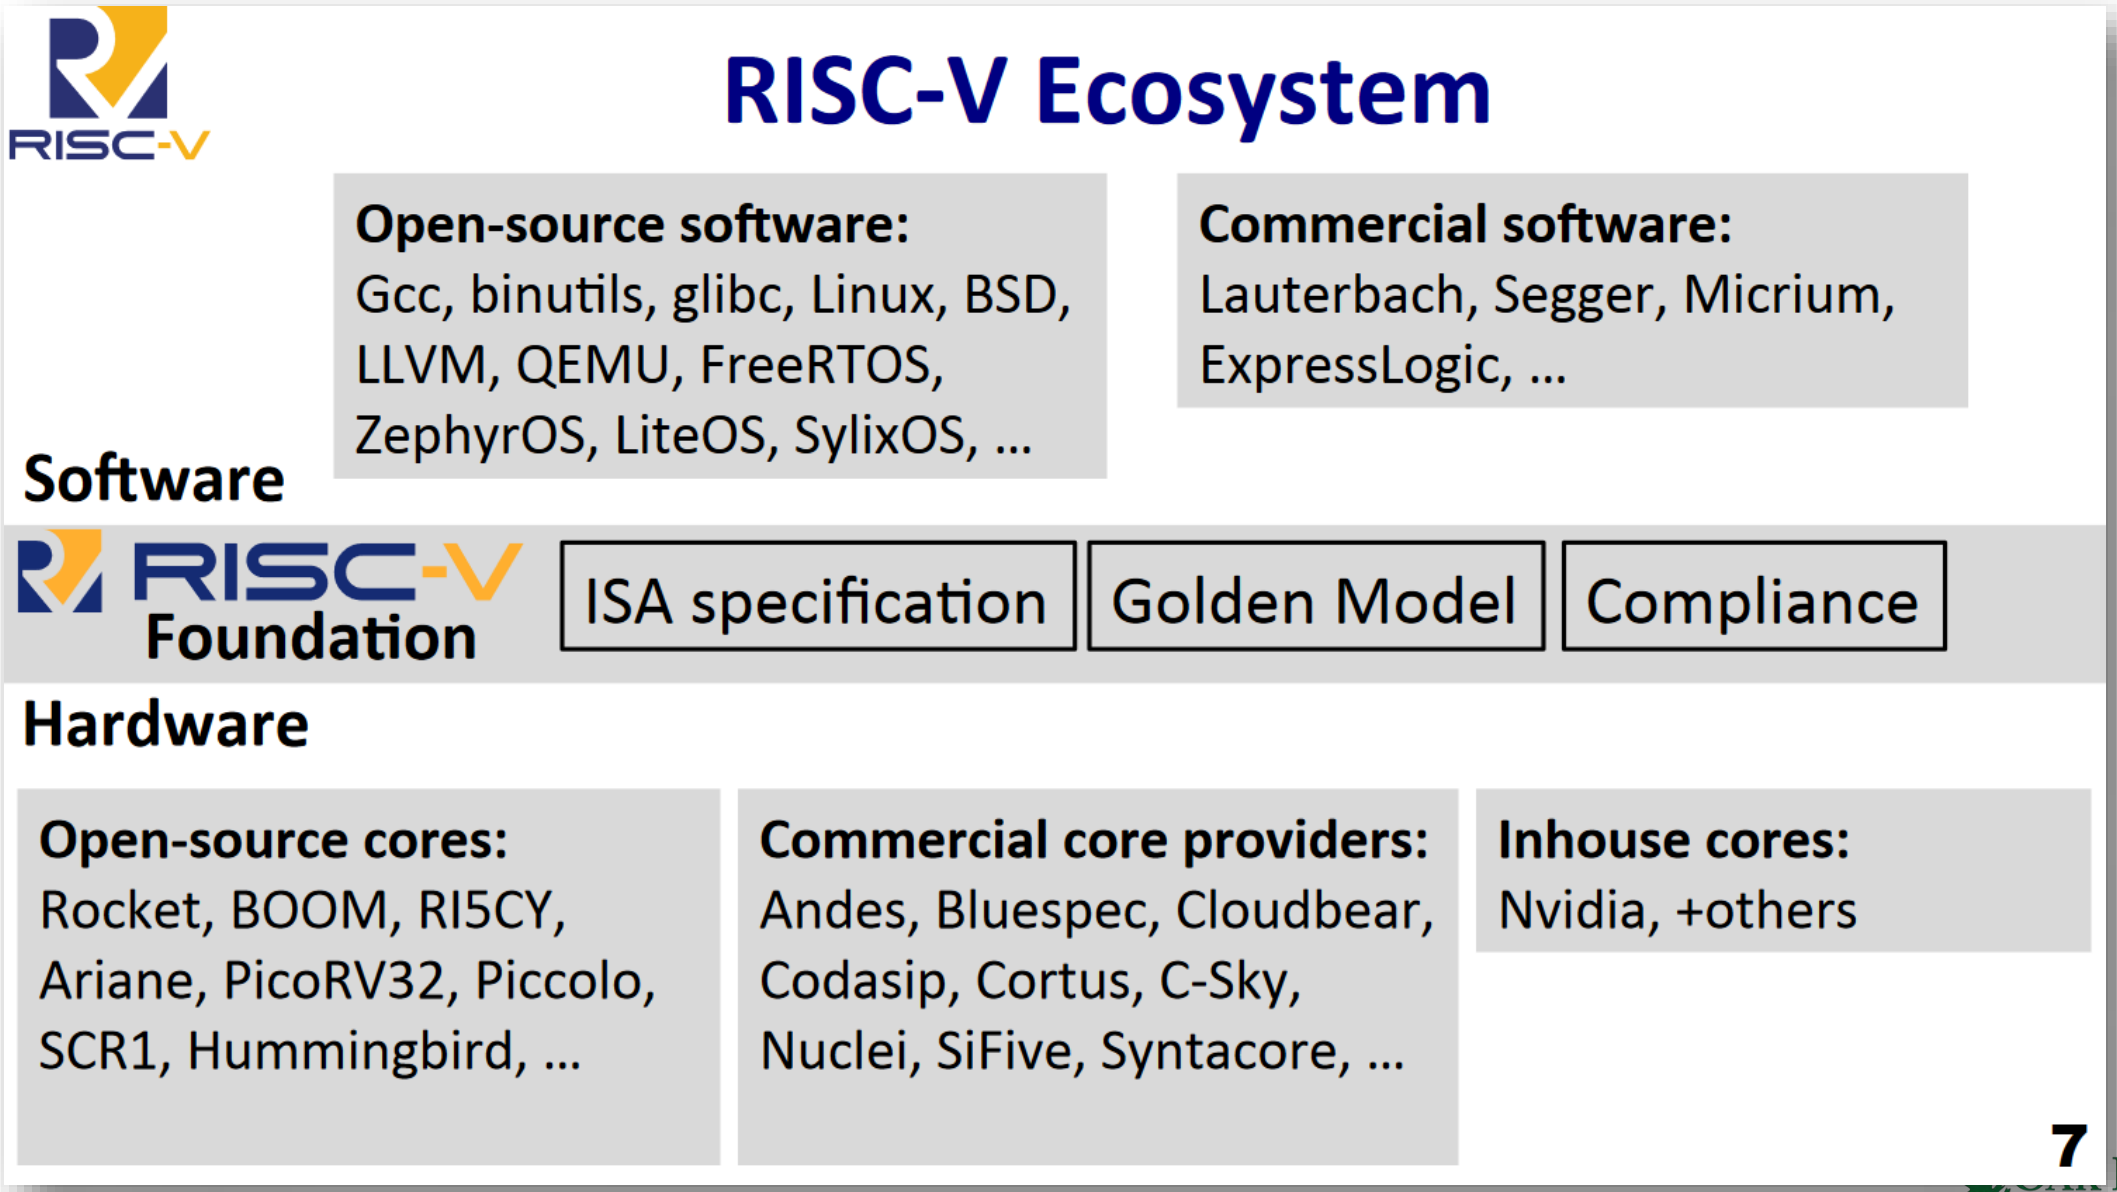
RISC-V Ecosystem
Open-source software:
Gcc, binutils, glibc, Linux, BSD, LLVM, QEMU, FreeRTOS, ZephyrOS, LiteOS, SylixOS, …
Commercial software:
Lauterbach, Segger, Micrium, ExpressLogic, …
Software
RISC-V Foundation
ISA specification
Golden Model
Compliance
Hardware
Open-source cores:
Rocket, BOOM, RI5CY, Ariane, PicoRV32, Piccolo, SCR1, Hummingbird, …
Commercial core providers:
Andes, Bluespec, Cloudbear, Codasip, Cortus, C-Sky, Nuclei, SiFive, Syntacore, …
Inhouse cores:
Nvidia, +others
7

# Applications

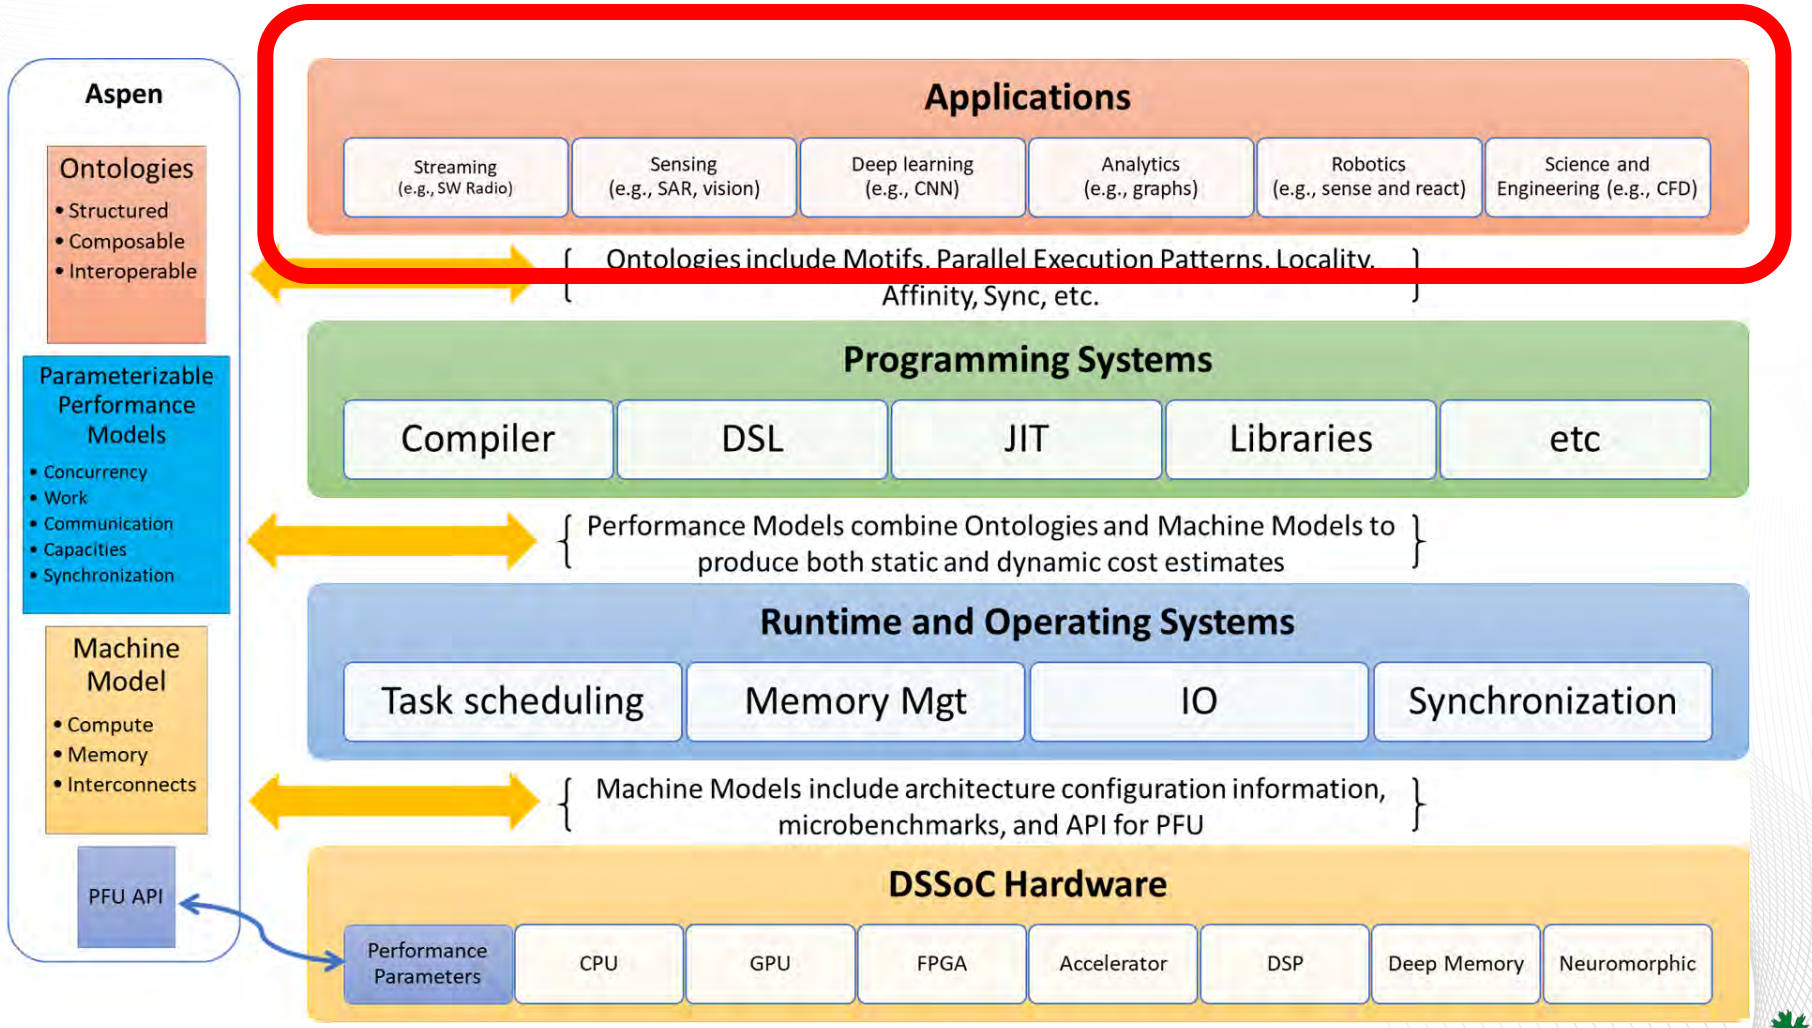

# End-to-End System: Gnu Radio for Wifi on two NVIDIA Xavier SoCs

Xavier SoC

Video/Image Files

UDP

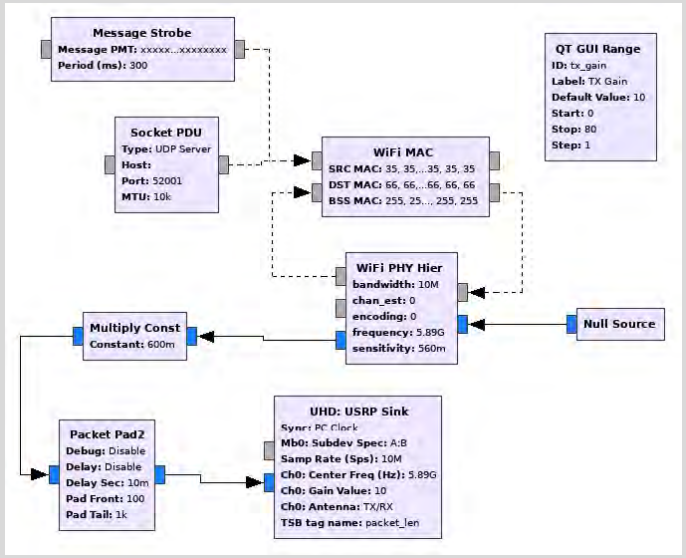


GNR IEEE-802.11 Transmit (TX)

Antenna

Xavier SoC

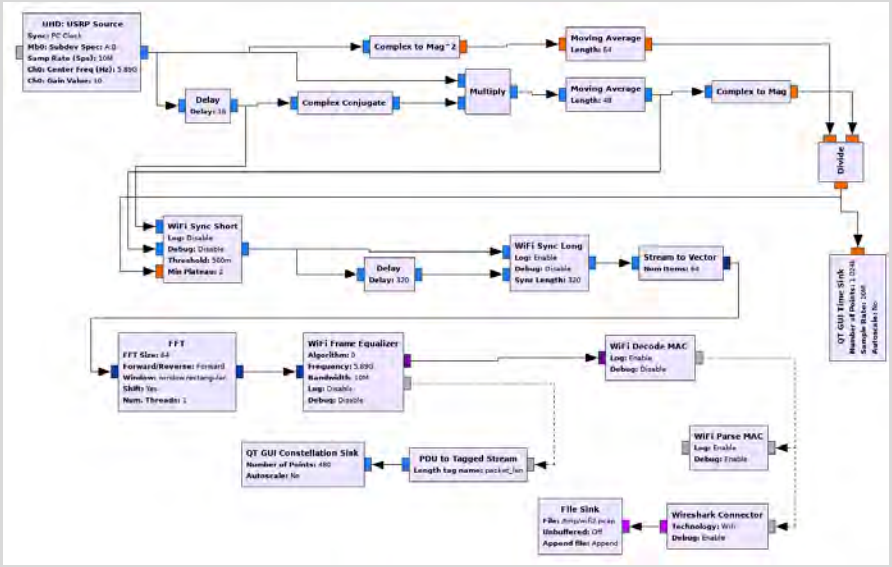


IEEE-802.11 Receive (RX)

UDP

OpenCV

OpenCL

- Signal processing: An open-source implementation of IEEE-802.11 WIFI a/b/g with GNR OOT modules.
- Input / Output file support via Socket PDU (UDP server) blocks
- Image/Video transcoding with OpenCL/OpenCV

# Applications Profiling

- **Preliminary SDR Application Profiling:**
  - **Created fully automated GRC profiling toolkit**
  - Ran each of the 89 flowgraph for 30 seconds
  - Profiled with performance counters
  - Major overheads:
    - Python glue code (libpython), O/S threading & profiling (kernel.kallsyms, libpthread), libc, ld, Qt
  - Runtime overhead:
    - Will require significant consideration when run on SoC
    - Cannot be executed in parallel
    - Hardware assisted scheduling is essential

| Library | Percentage |
|---|---|
| [kernel.kallsyms] | 27.8547 |
| libpython | 18.6281 |
| **libgnuradio** | **11.7548** |
| libc | 7.7503 |
| ld | 3.8839 |
| **libvolk** | **3.7963** |
| libperl | 3.7837 |
| [unknown] | 3.6465 |
| libQt5 | 2.9866 |
| libpthread | 2.1449 |

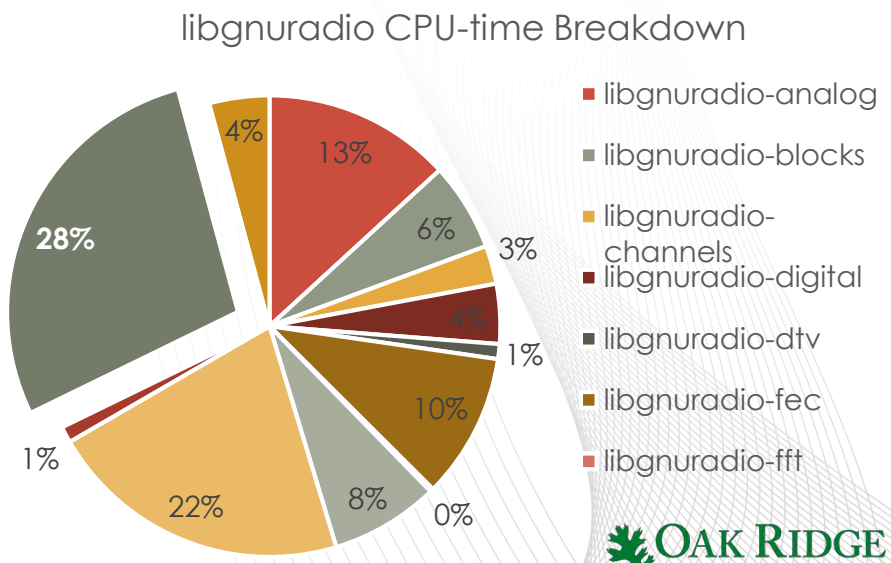


OAK RIDGE National Laboratory

# GNR-Tools

- GNR-Tools
  - PY1Q2: Three tools are released
    - Block-level Ontologies [ontologyAnalysis]
      - Following properties are extracted from a batch of block definition files: Descriptions and IDs, source and sink ports (whether input/output is scalar, vector or multi-port), allowed data types, and additional algorithm-specific parameters
    - Flowgraph Characterization [workflowAnalysis]
      - Characterization of GNR workloads at the flowgraph level.
      - Scripts automatically run for for 30 seconds and reports a breakdown of high-level library module calls
    - Design-space Exploration [designSpaceCL]
      - Script to run 13 blocks included in gr-clenabled
        - Both on a GPU and on a single CPU core
        - By using input sizes varying between 24 and 227 elements.
  - PY1Q3: Two more tools are added
    - cgran-scraper
    - GRC-analyzer

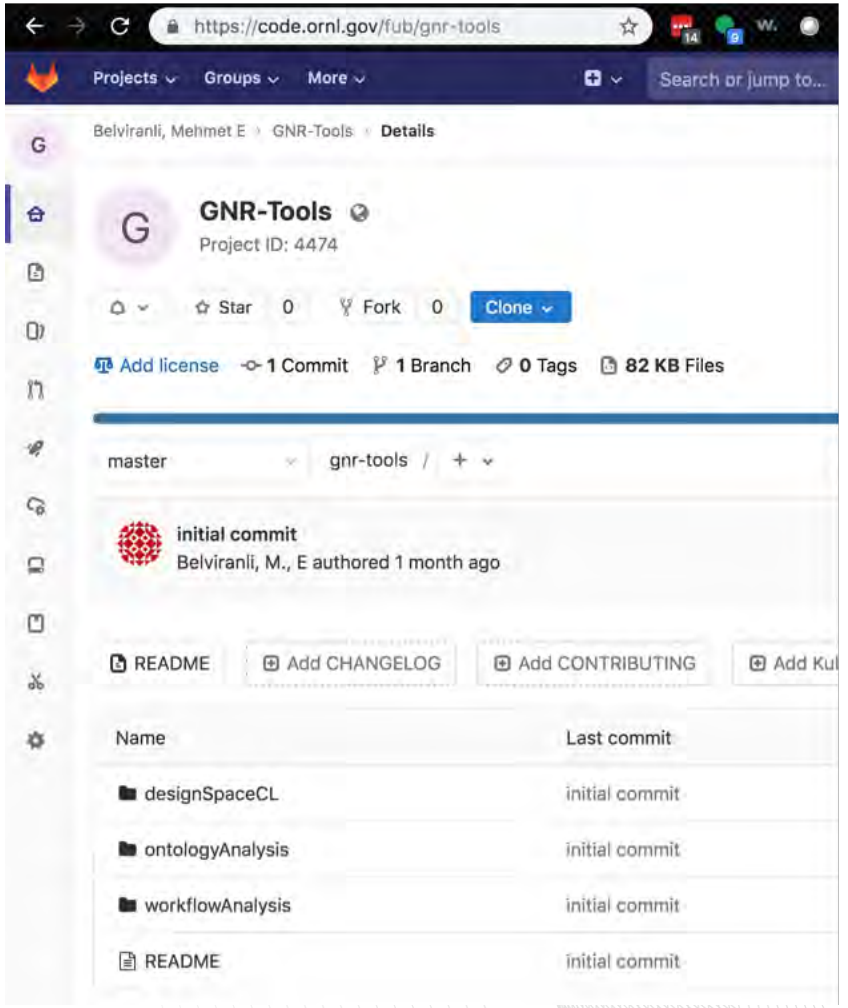


https://code.ornl.gov/fub/gnr-tools

# GRC statistics: Block Proximity Analysis

## Block proximity analysis

- Creates a graph:
  - Nodes: Unique block types
  - Edges: Blocks used in the same GRC file.
  - Every co-occurrence increases edge weight by 1.
- This example was run
  - With `--mode proximityGraph`
  - On randomly selected sub-set of GRC files

| | |
|---|---|
| borip-USRP-UHD.grc | live_signal_detection.grc |
| cdma_tx_hier1.grc | psk_burst_ldpc_tx.grc |
| cdma_tx_hier.grc | psk_burst_tx.grc |
| dsat.grc | rfnoc_digital_gain_network_host.grc |
| dsss_sim_perfekt_sync_fg_without_fec.grc | rtty_decode.grc |
| dvbt_tx_demo_8k_QPSK_rate78.grc | run_RootMUSIC_lin_array_simulation.grc |
| fbmc_frame_generator_perf_test.grc | sat_1kuns_pf.grc |
| flarm_2chan.grc | sat_3cat_2.grc |
| frontend_lilacsat1_rx_fcdpp.grc | snapshot-approach.grc |
| fsk_tx.grc | symbol_differential_filter_phases.grc |
| ieee802_15_4_OQPSK_PHY.grc | symbol_sampling.grc |
| jy1sat.grc | tx_usrp.grc |
| kr01.grc | usrp-input.grc |

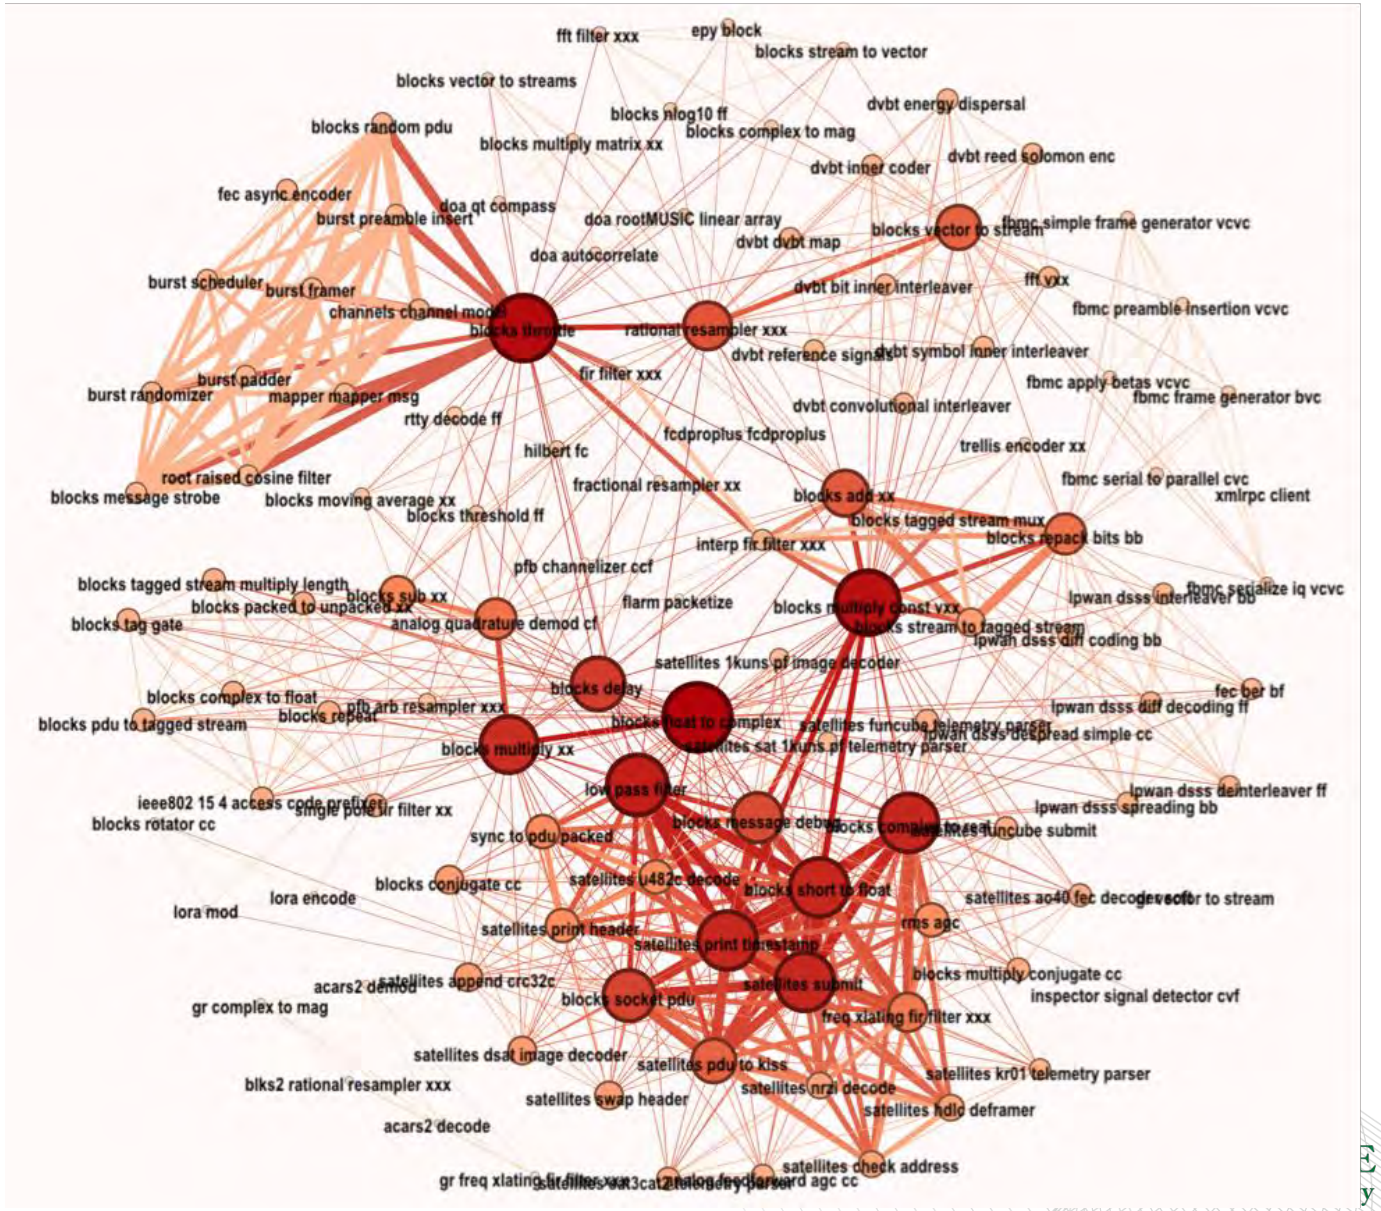

# Integrating Modeling Across the Stack with Aspen

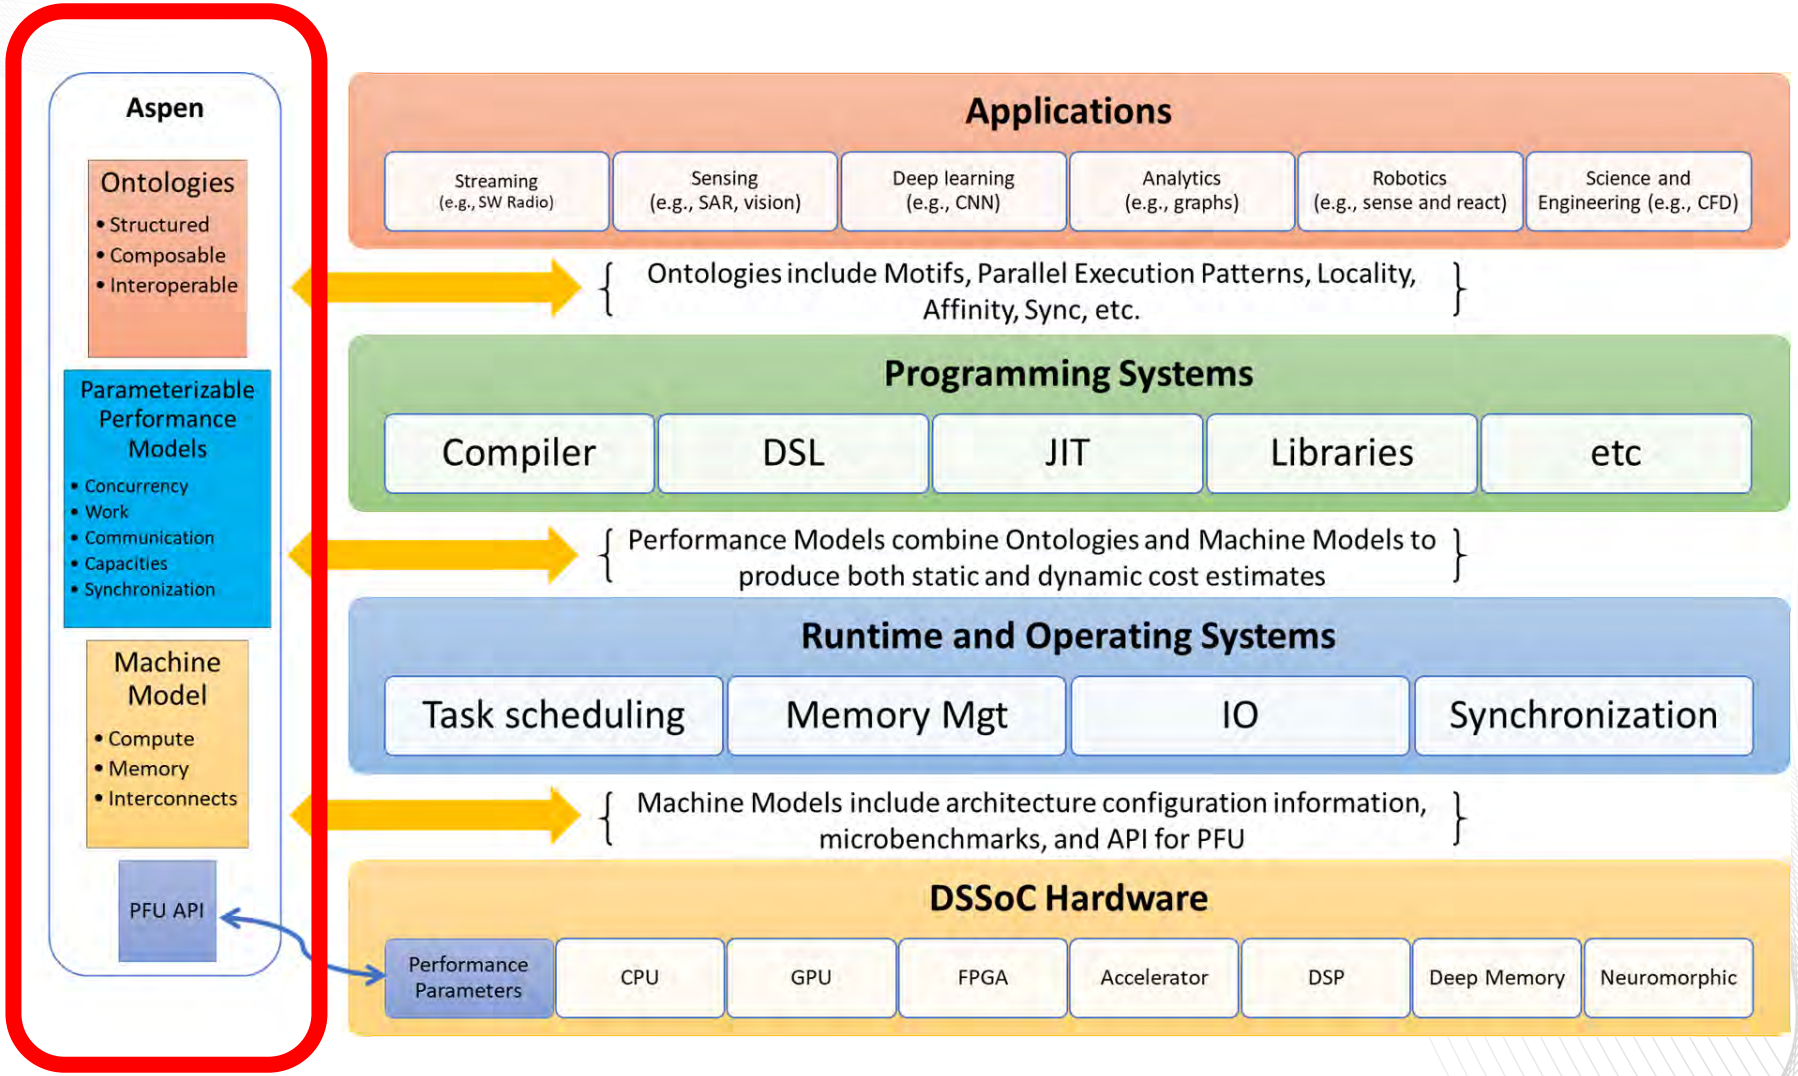

# Aspen: Abstract Scalable Performance Engineering Notation

## Model Creation

- Static analysis via compiler, tools
- Empirical, Historical
- Manual (for future applications)

**Representation in Aspen**

- **Modular**
- **Sharable**
- **Composable**
- **Reflects prog structure**

## Model Uses

- Interactive tools for graphs, queries
- Design space exploration
- Workload Generation
- Feedback to Runtime Systems

**Source code**

```
static inline
void CalcMonotonicQGradientsForElems(Index_t p_nodelist[T_NUMELEM8],
        Real_t p_x[T_NUMNODE], Real_t p_y[T_NUMNODE], Real_t p_z[T_NUMNODE],
        Real_t p_xd[T_NUMNODE], Real_t p_yd[T_NUMNODE],Real_t p_zd[T_NUMNODE],
        Real_t p_volo[T_NUMELEM], Real_t p_vnew[T_NUMELEM],
        Real_t p_delx_zeta[T_NUMELEM], Real_t p_delv_zeta[T_NUMELEM],
        Real_t p_delx_xi[T_NUMELEM], Real_t p_delv_xi[T_NUMELEM],
        Real_t p_delx_eta[T_NUMELEM], Real_t p_delv_eta[T_NUMELEM])
{
   Index_t i;
   Index_t numElem = m_numElem;
#pragma acc parallel loop independent present(p_vnew, p_nodelist, p_x, p_y, p_z, p_xd, \
p_yd, p_zd, p_volo, p_delx_xi, p_delx_eta, p_delx_zeta, p_delv_xi, p_delv_eta,\
p_delv_zeta)
   for (i = 0 ; i < numElem ; ++i ) {
      const Real_t ptiny = 1.e-36 ;
      Real_t ax,ay,az ;
      Real_t dxv,dyv,dzv ;

      const Index_t *elemToNode = &p_nodelist[8*i];
      Index_t n0 = elemToNode[0] ;
      Index_t n1 = elemToNode[1] ;
      Index_t n2 = elemToNode[2] ;
      Index_t n3 = elemToNode[3] ;
      Index_t n4 = elemToNode[4] ;
      Index_t n5 = elemToNode[5] ;
      Index_t n6 = elemToNode[6] ;
      Index_t n7 = elemToNode[7] ;

      Real_t x0 = p_x[n0] ;
```

*E.g., MD, UHPC CP 1, Lulesh, 3D FFT, CoMD, VPFFT, …*

**Aspen code**

```
kernel CalcMonotonicQGradients {
  execute [numElems]
  {
    loads [8 * indexWordSize] from nodelist
    // Load and cache position and velocity.
    loads/caching [8 * wordSize] from x
    loads/caching [8 * wordSize] from y
    loads/caching [8 * wordSize] from z

    loads/caching [8 * wordSize] from xvel
    loads/caching [8 * wordSize] from yvel
    loads/caching [8 * wordSize] from zvel

    loads [wordSize] from volo
    loads [wordSize] from vnew
    // dx, dy, etc.
    flops [90] as dp, simd
    // delvk delxk
    flops [9 + 8 + 3 + 30 + 5] as dp, simd
    stores [wordSize] to delv_xeta
    // delxi delvi
    flops [9 + 8 + 3 + 30 + 5] as dp, simd
    stores [wordSize] to delx_xi
    // delxj and delvj
    flops [9 + 8 + 3 + 30 + 5] as dp, simd
    stores [wordSize] to delv_eta
  }
}
```

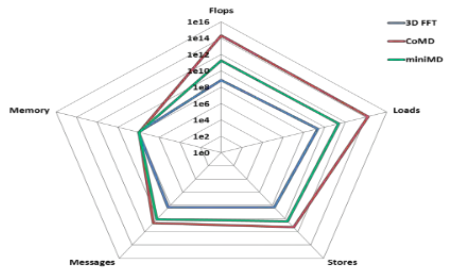


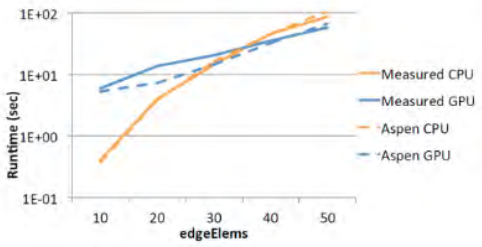


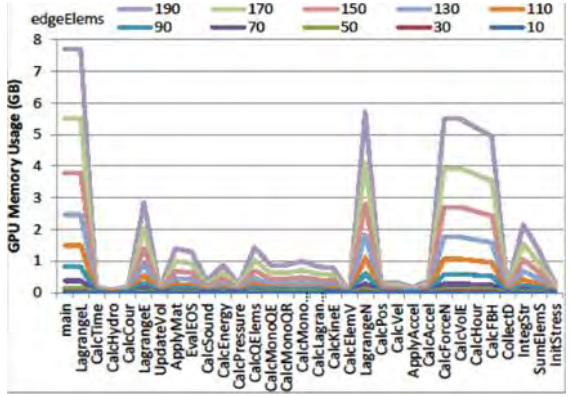


K. Spafford, et al, "Aspen: A Domain Specific Language for Performance Modeling," in Proc. *SC12.*

# GNURadio Flowgraph to Aspen Application Model Conversion

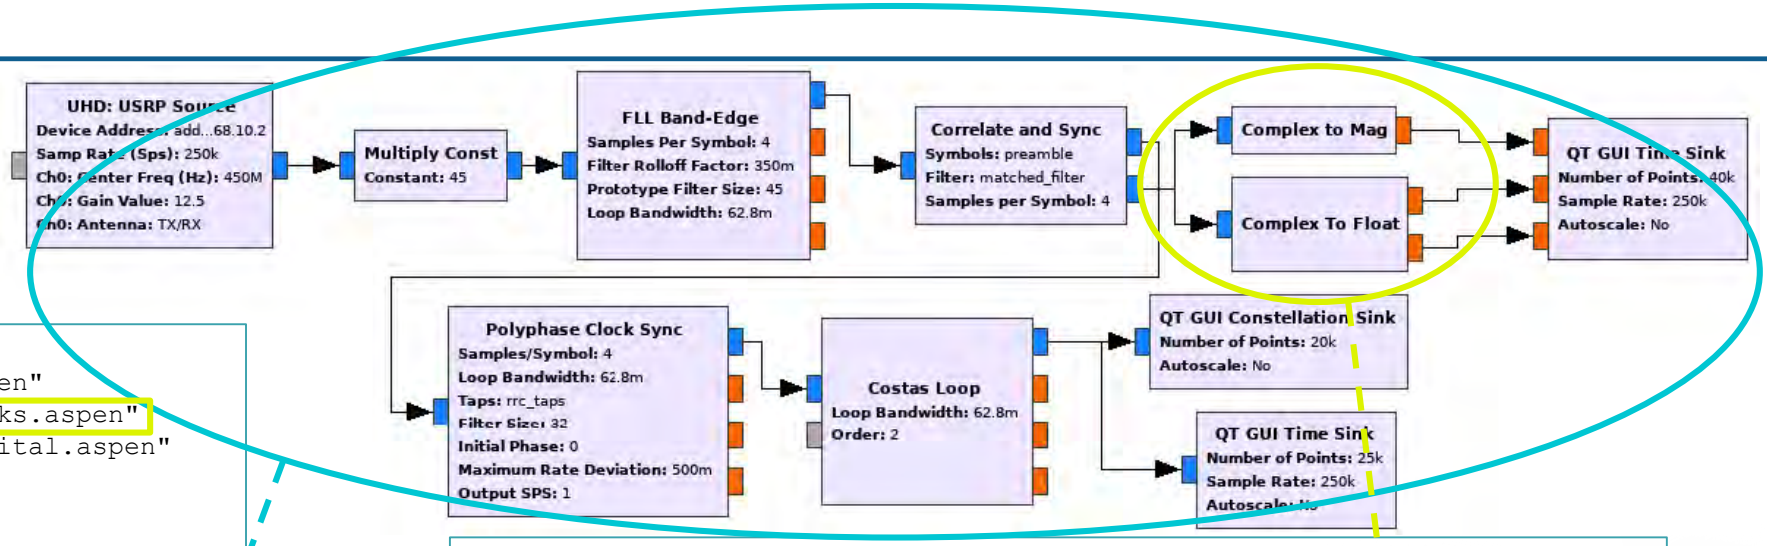


```
model demod-uhd-sync{
  import uhd from "gr-uhd.aspen"
  import blocks from "gr-blocks.aspen"
  import digital from "gr-digital.aspen"

  kernel main {
    call uhd.usrp_source()
    call blocks.multiply_const_vxx()
    call digital.fll_band_edge_cc()
    call digital.correlate_and_sync_cc()

    parallel {
      sequence{
        parallel {
          call blocks.complex_to_mag()
          call blocks.complex_to_float()
        }
        call blocks_file_sink()
      }
      sequence{
        call digital.pfb_clock_sync_xxx()
        call digital.costas_loop_cc()
        parallel {
          call blocks.file_sink()
          call blocks.file_sink()
        }
      }
    }
  }
}
```

demod-uhd-sync.aspen

```
model blocks {
  kernel complex_to_mag {
    param aspen_param_default = 1
    param aspen_param_sizeof_float = 4
    param noutput_items = 8192
    param aspen_param_sizeof_FComplex = 8
    execute [noutput_items] "block_clComplexToMag_kernel4"
      flops [1] as integer
    execute  "block_clComplexToMag_kernel4__intracommIN"
      intracomm [(aspen_param_sizeof_FComplex*noutput_items)] as copyin
    map [noutput_items] "mapblock_clComplexToMag_kernel4"
      execute  "block_clComplexToMag_kernel7" {
        loads [(1*aspen_param_sizeof_FComplex)] as stride(1)
        loads [(1*aspen_param_sizeof_FComplex)] as stride(1)
        stores [(1*aspen_param_sizeof_float)] as stride(1)
        flops [4] as dp, simd
      }
    execute  "block_clComplexToMag_kernel4__intracommOUT"
      intracomm [(aspen_param_sizeof_float*noutput_items)] as copyout
  }
  kernel complex_to_float {...}
  ...
}
```

gr-blocks.aspen

# Graph-Based Abstract Machine Model

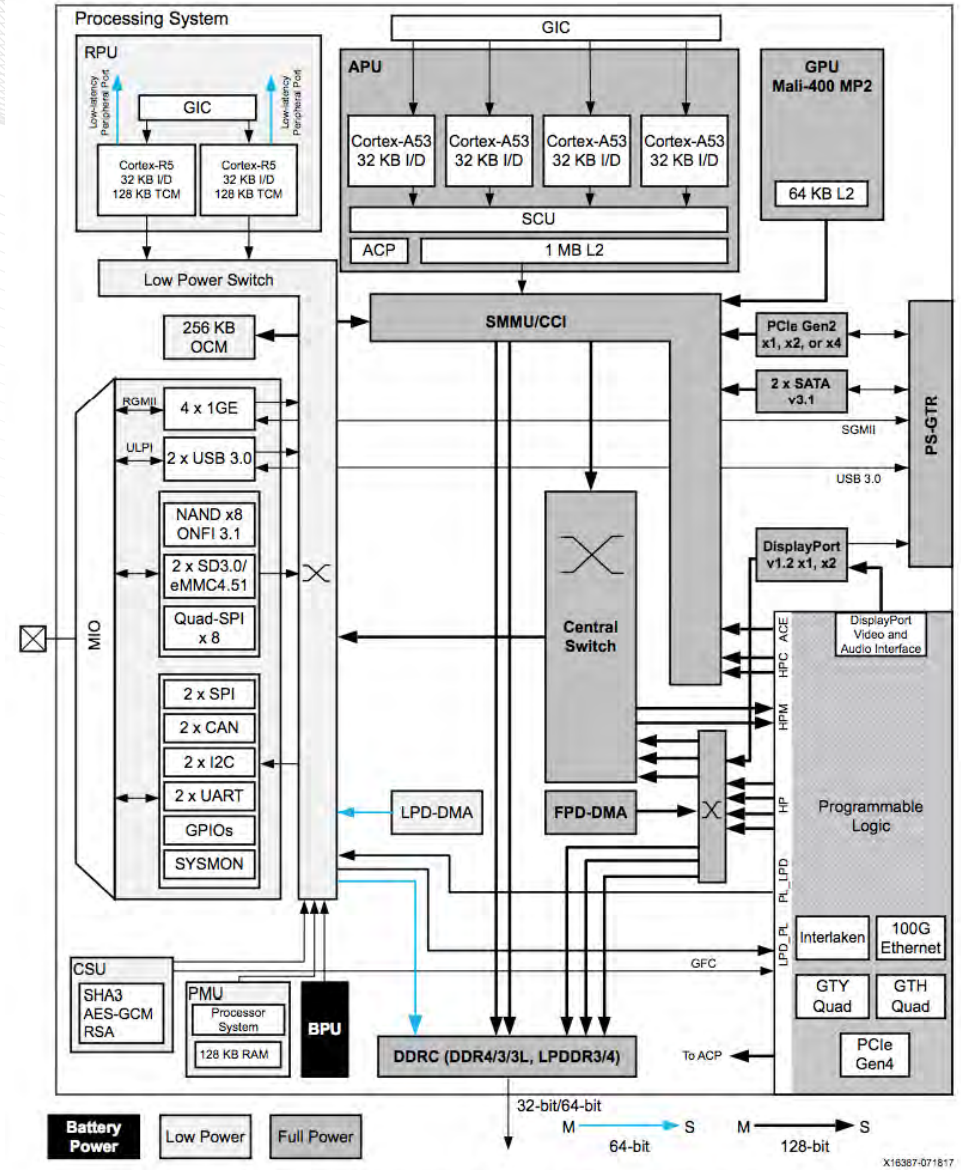


*Figure 3-1:* **Zynq UltraScale+ MPSoC Top-Level Block Diagram**

```
class Zynq::Board-ZCU102 : Aspen::CompoundNode {
// Processing units
  Zynq::APU cpu;
  ARM::Mali400MP2 gpu;
  ARM::CortexR5 rpu;
  Xilinx::UltraScale+<nFPUs=400M> fpga;

  // Memory
  Aspen::DDR3<freq=2000MHZ, CL=16> systemMemory;

  // Memory controllers, switchs, mmus
  Zynq::SMMU smmu;
  Aspen::Switch<bw=100GBs, latency= 25ns> lpSwitch;
  Aspen::Switch<bw=1TBs, latency= 35ns> centralSwitch;
  Aspen::PCIController<ver=3, totalLanes=24> pciController;

  // Define interconnects (edges)
  Aspen::Bus<bw=400GBps> cci_fp;
  Aspen::Bus<bw=100GBps> cci_lp;
  Aspen::PCIe<version=3,lane=16> pcieBus;

  @add
    cpu --cci_fp-- smmu;
    gpu --cci_fp-- smmu;
    fpga --cci_fp-- smmul
    systemMemory --cci_fp[2]-- smmu; // Multiple links

    smmu --cci_fp-- centralSwitch;
    smmu --cci_fp-- pciController;
    fpga --cci_fp[2]-- centralSwitch

    lpSwitch --cci_fp-->> smmu; // Unidirectional link
    lpSwitch <<--cci_fp-- centralSwitch
    rpu --cci_lp[2]-- lpSwitch;

    pciController --pcieBus -->> Aspen::OUTPUT
    pciController <<--pcieBus -->> Aspen::INPUT
```

Nodes

Edges

Graph



Belviranli, M E, et al, "FLAME: Graph-based Hardware Representations for Rapid and Precise Performance Modeling", DATE'19

# Programming Systems

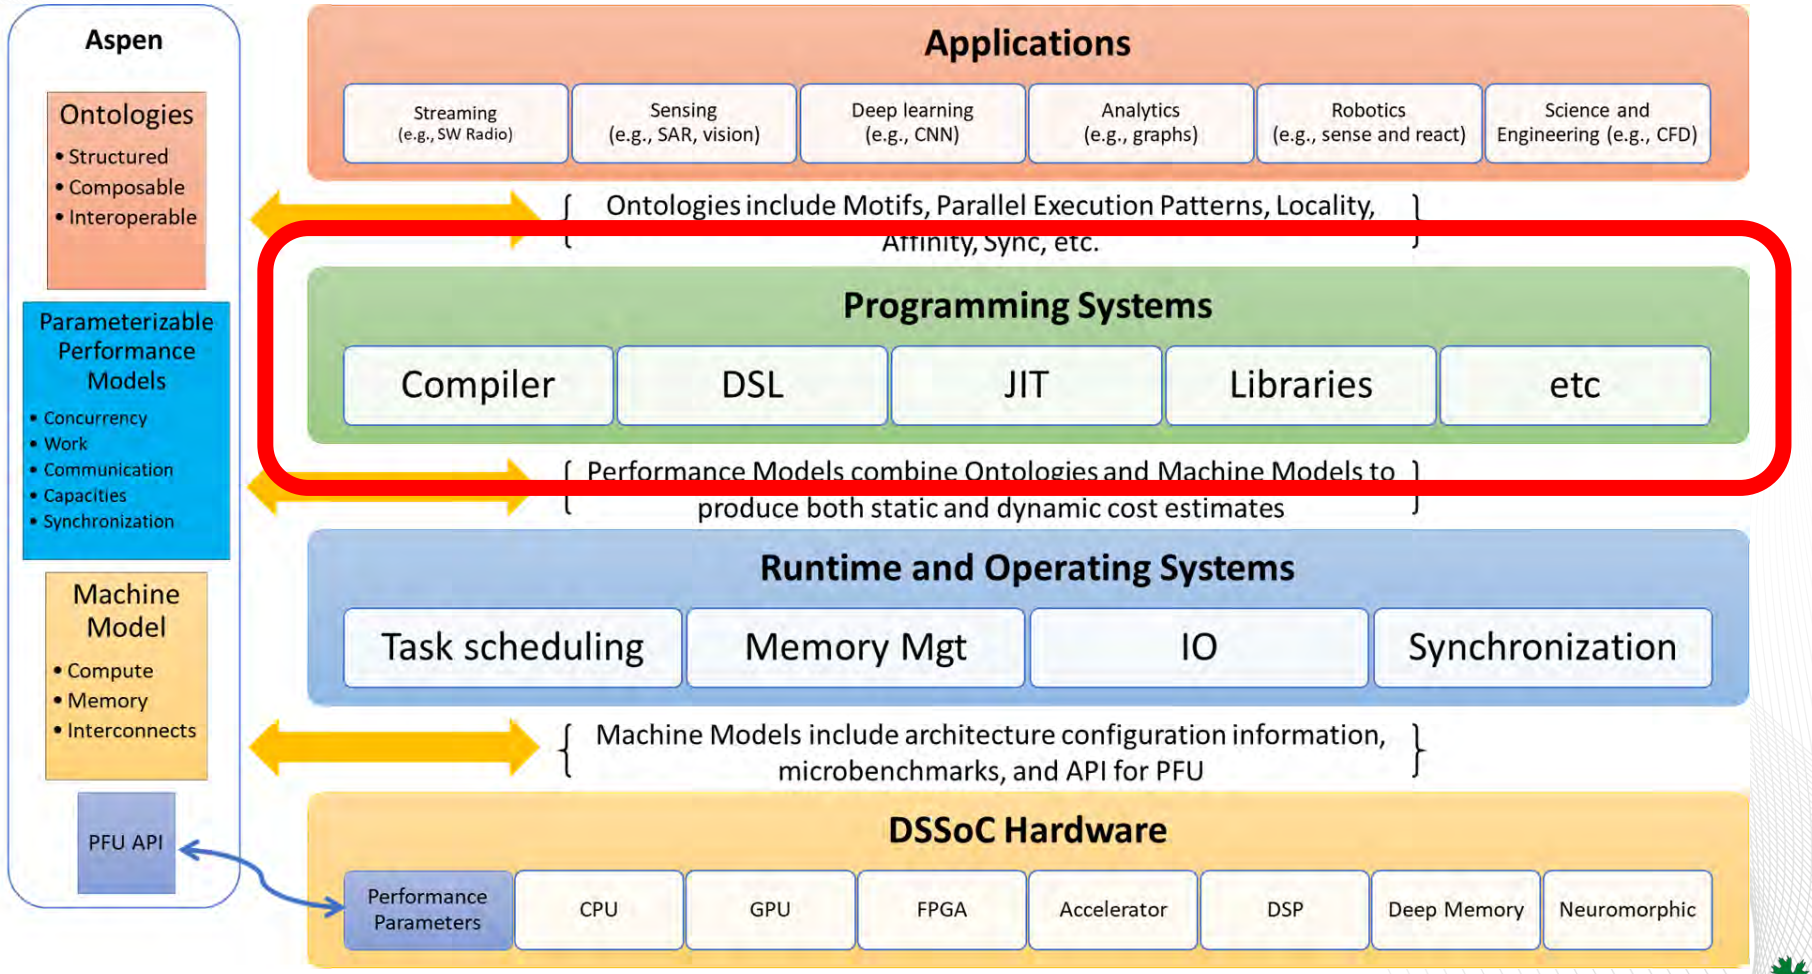

# OpenARC: Open Accelerator Research Compiler*

- OpenARC is the first open-sourced, OpenACC/OpenMP compiler supporting Altera FPGAs, in addition to NVIDIA/AMD GPUs and Intel Xeon Phis.
- OpenARC is a high-level intermediate representation based, extensible compiler framework, where various performance optimizations, traceability mechanisms, fault tolerance techniques, etc., can be built for the complex heterogeneous computing.

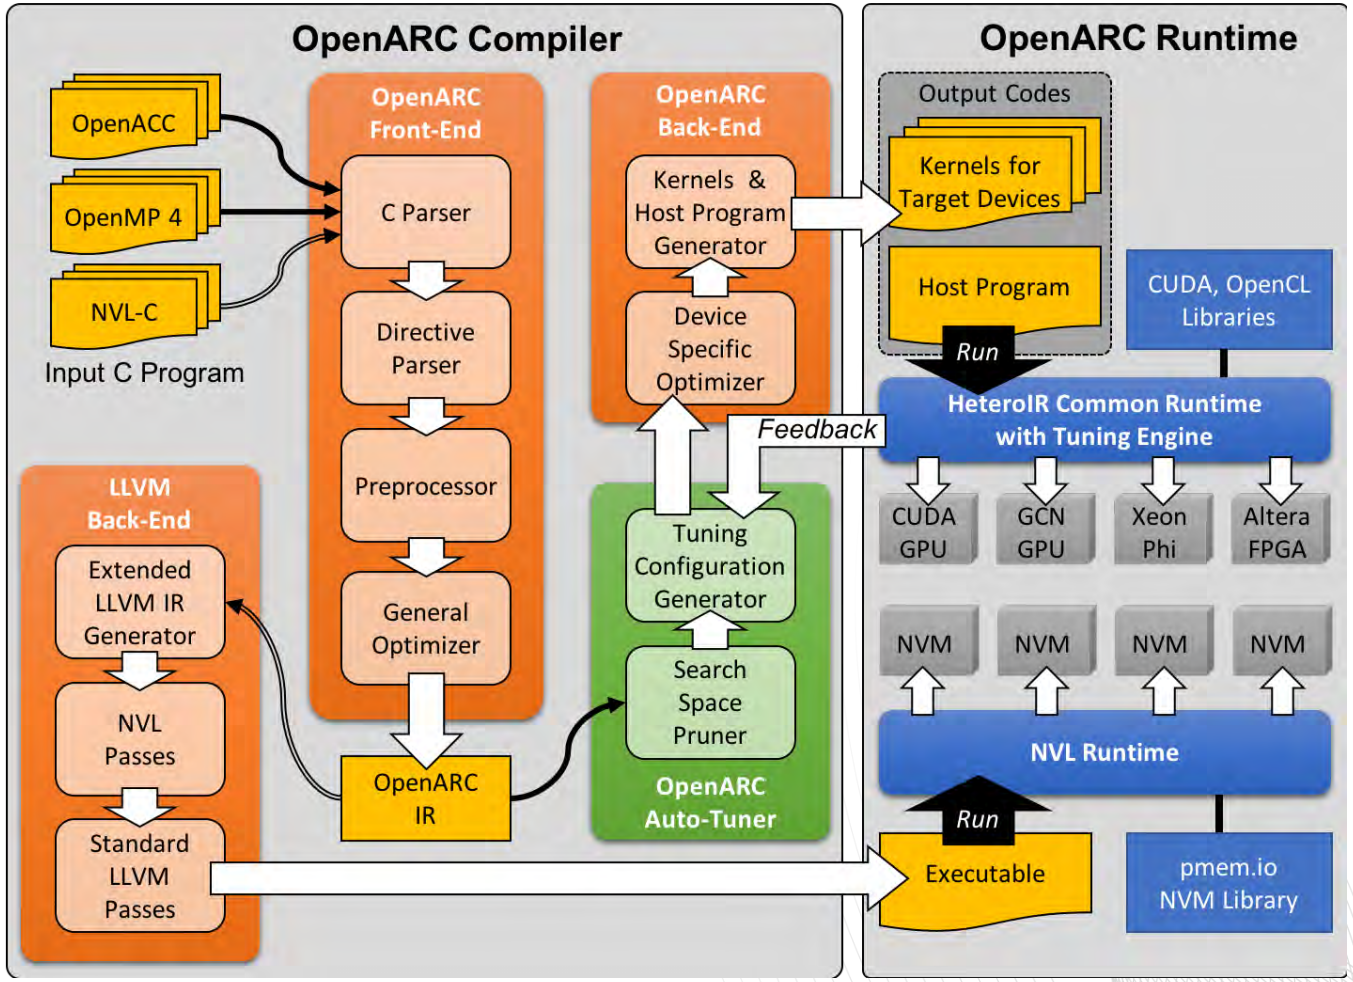


* OpenARC, Lee, HPDC '14.

# Heterogeneous Architecture Support by OpenARC

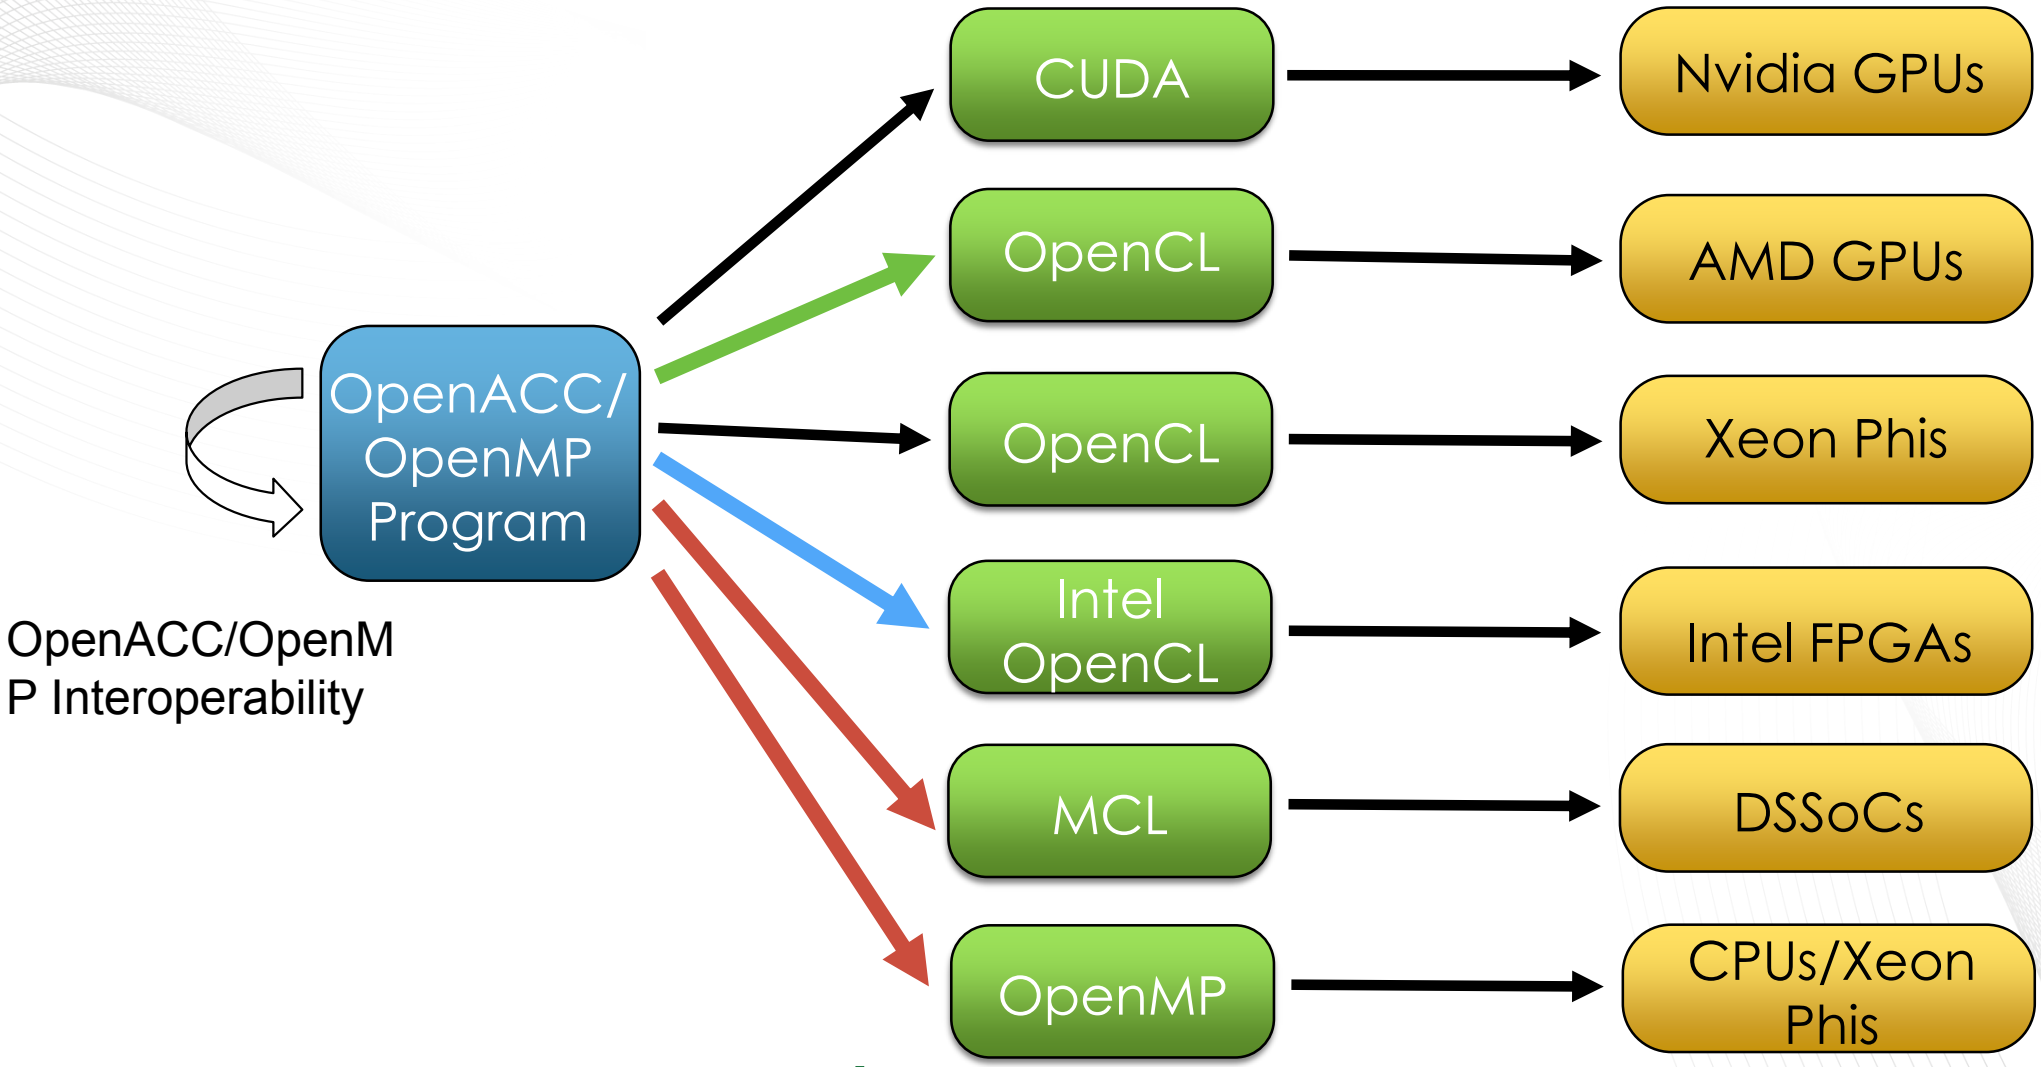


OpenACC/OpenM
P Interoperability

# Initial Performance of OpenACC GNU Radio Blocks on Xavier

- Ported the OpenACC version of the GNU Radio blocks to OpenMP3 (CPU target), CUDA (GPU target), and OpenCL (GPU target) and compared against the reference CPU version.
- Tested Platform: NVIDA Jetson Xavier (8-core ARM CPU, NVIDIA Volta GPU, two NVDLA Engines and VLIW Vision Processor)

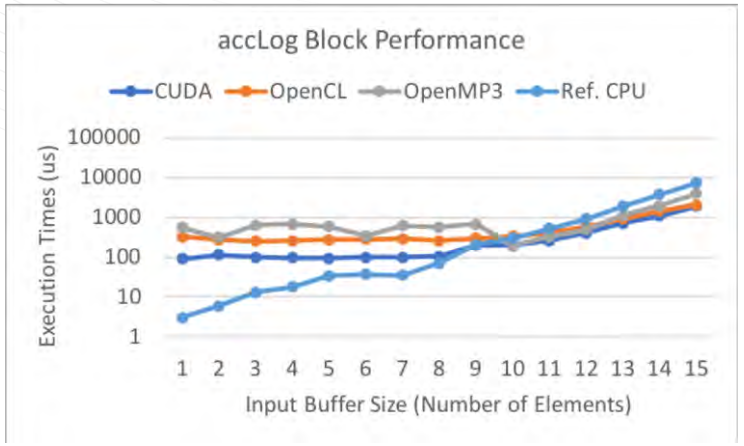


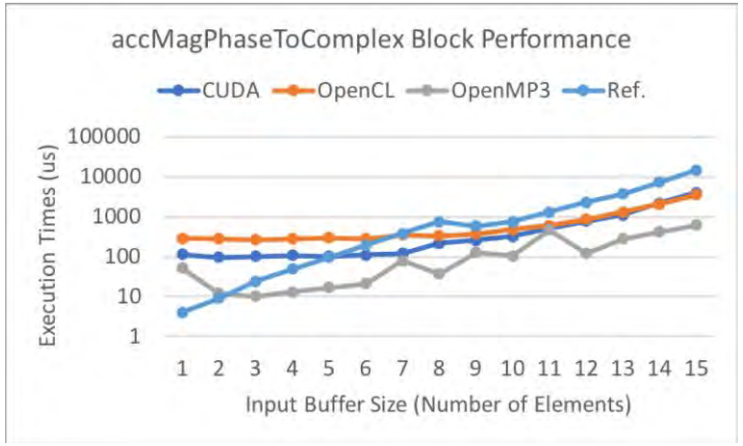


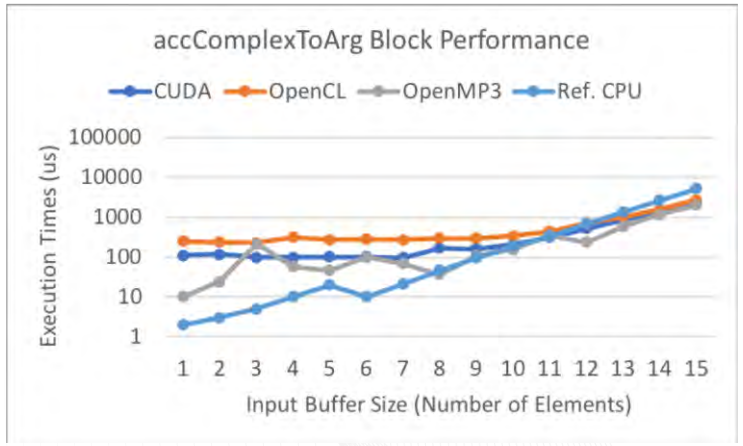


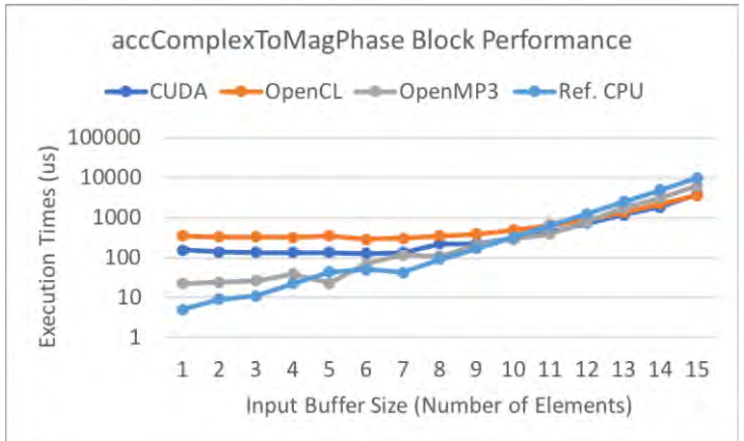


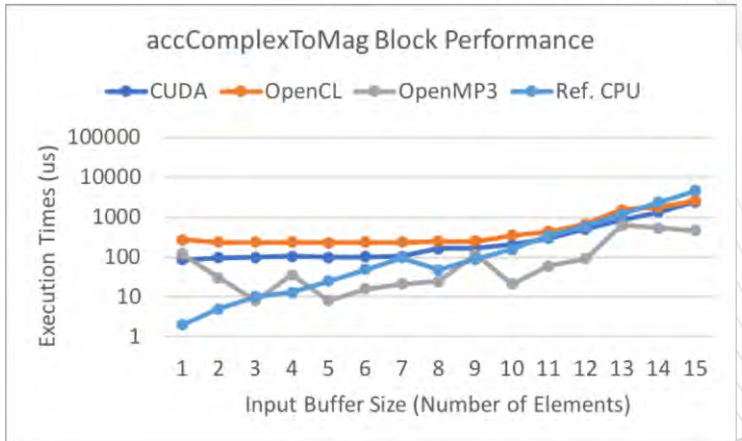


Omitted MCL since JIT compilation not yet factored

# Program Characteristics Prediction of GNU Radio Blocks using Aspen

- OpenARC automatically generates a structured Aspen performance model from the ported OpenACC code of the GNU Radio blocks.
- Aspen performance prediction tools digest the generated Aspen models and derive performance predictions for the target application.

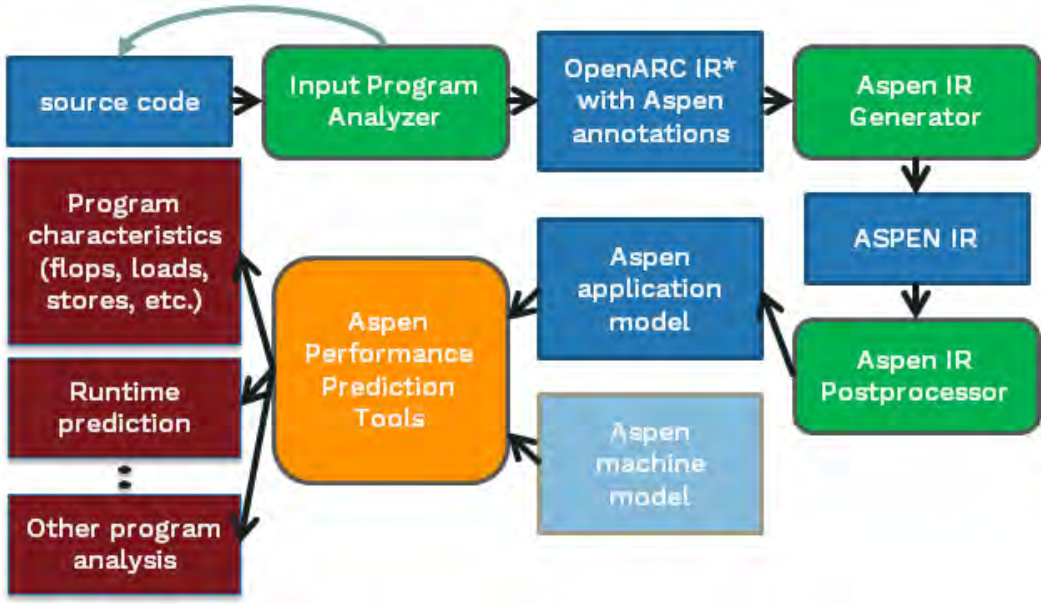


COMPASS: A Framework for Automated Performance Modeling and Prediction

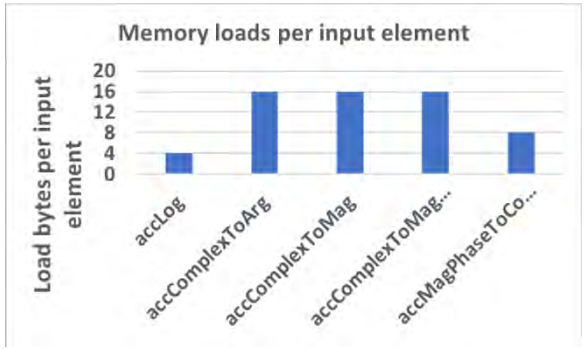


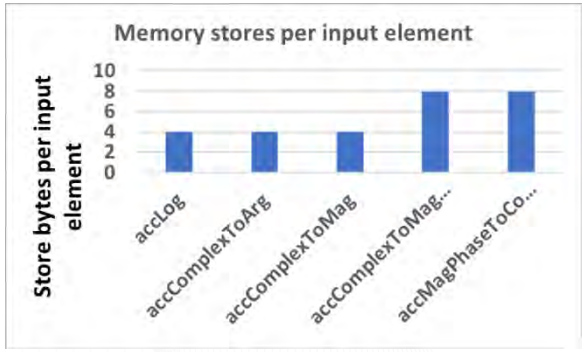


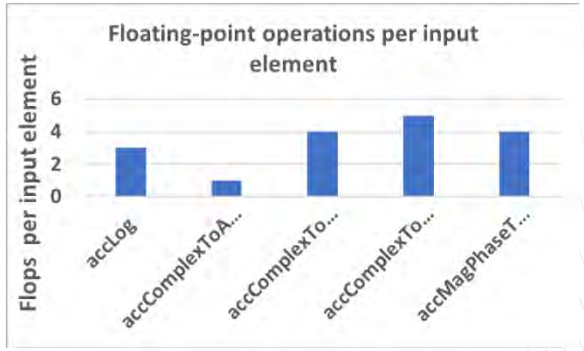


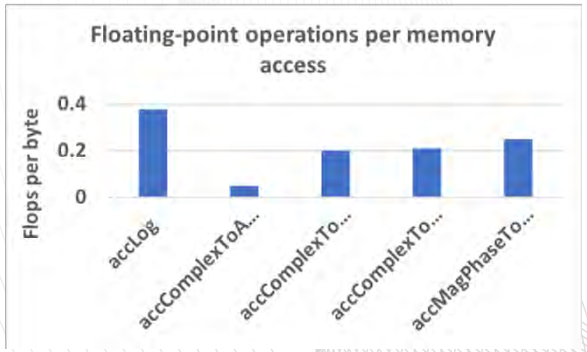

# Programming Systems

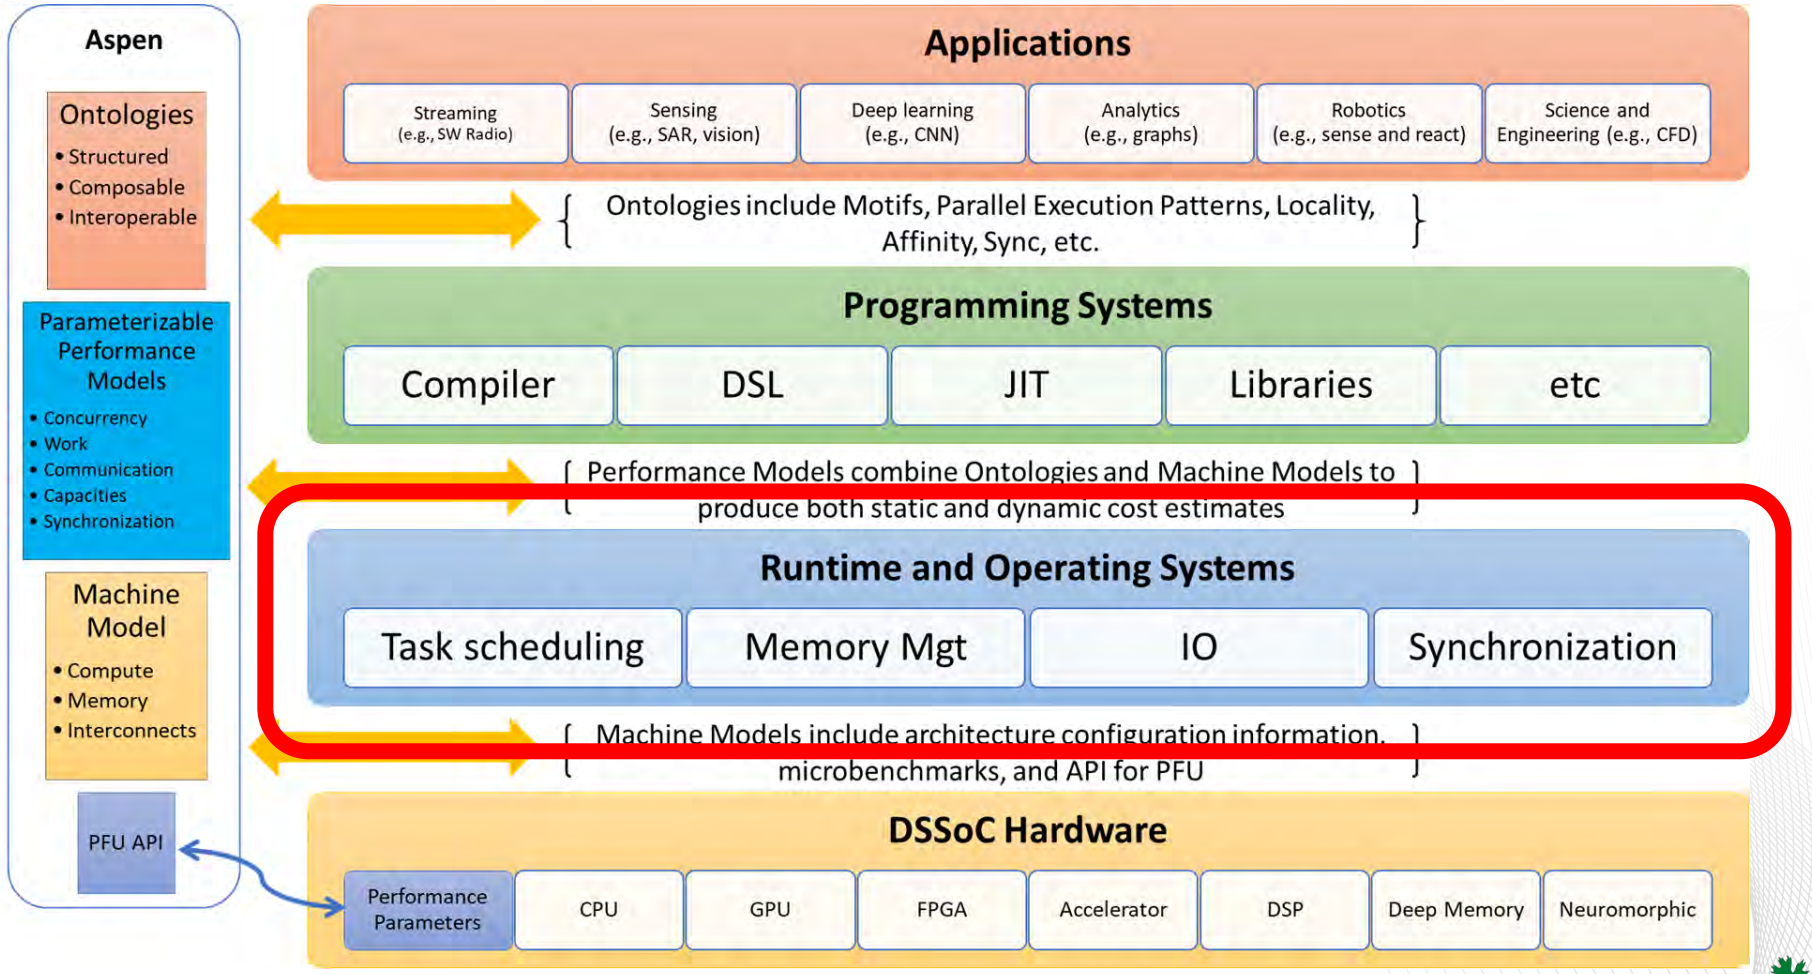

# Cosmic Runtime and Scheduler

- Framework for programming extremely heterogeneous systems
  - Programming model and programming model runtime
  - Maximize resource utilizations
  - Abstract low-level architecture details from programmers
  - Dynamically schedule work to available resources
- Key programming features:
  - Scheduler dispatches application tasks to available computing resources
  - **Asynchronous** execution of runnable tasks
  - Devices are managed by the scheduler and presented as “Processing Elements” to users
  - **Independent applications** submit tasks without having to synchronize with each other
  - Simplified APIs and programming model (e.g., compared to OpenCL)
- Flexibility:
  - Provides a scheduling framework in which new scheduling algorithm can be plugged in
  - **Multiple scheduling algorithms** co-exist
  - Users don’t need to port code when running on different systems
  - Executing tasks on different PEs doesn’t require user intervention or code modification
  - Resources allocated at the last moment

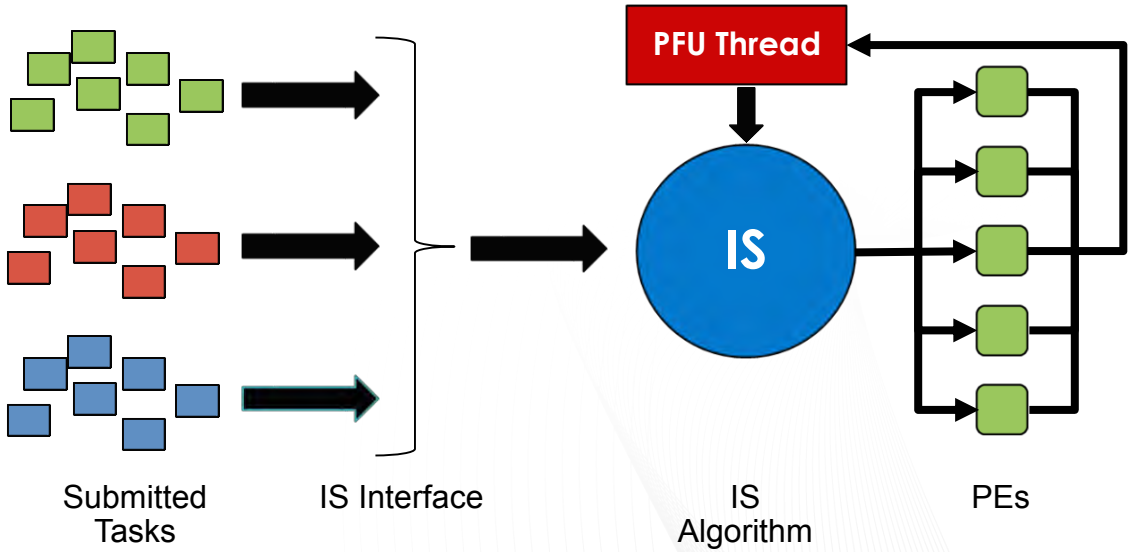

# Exploiting Parallelism in SDR

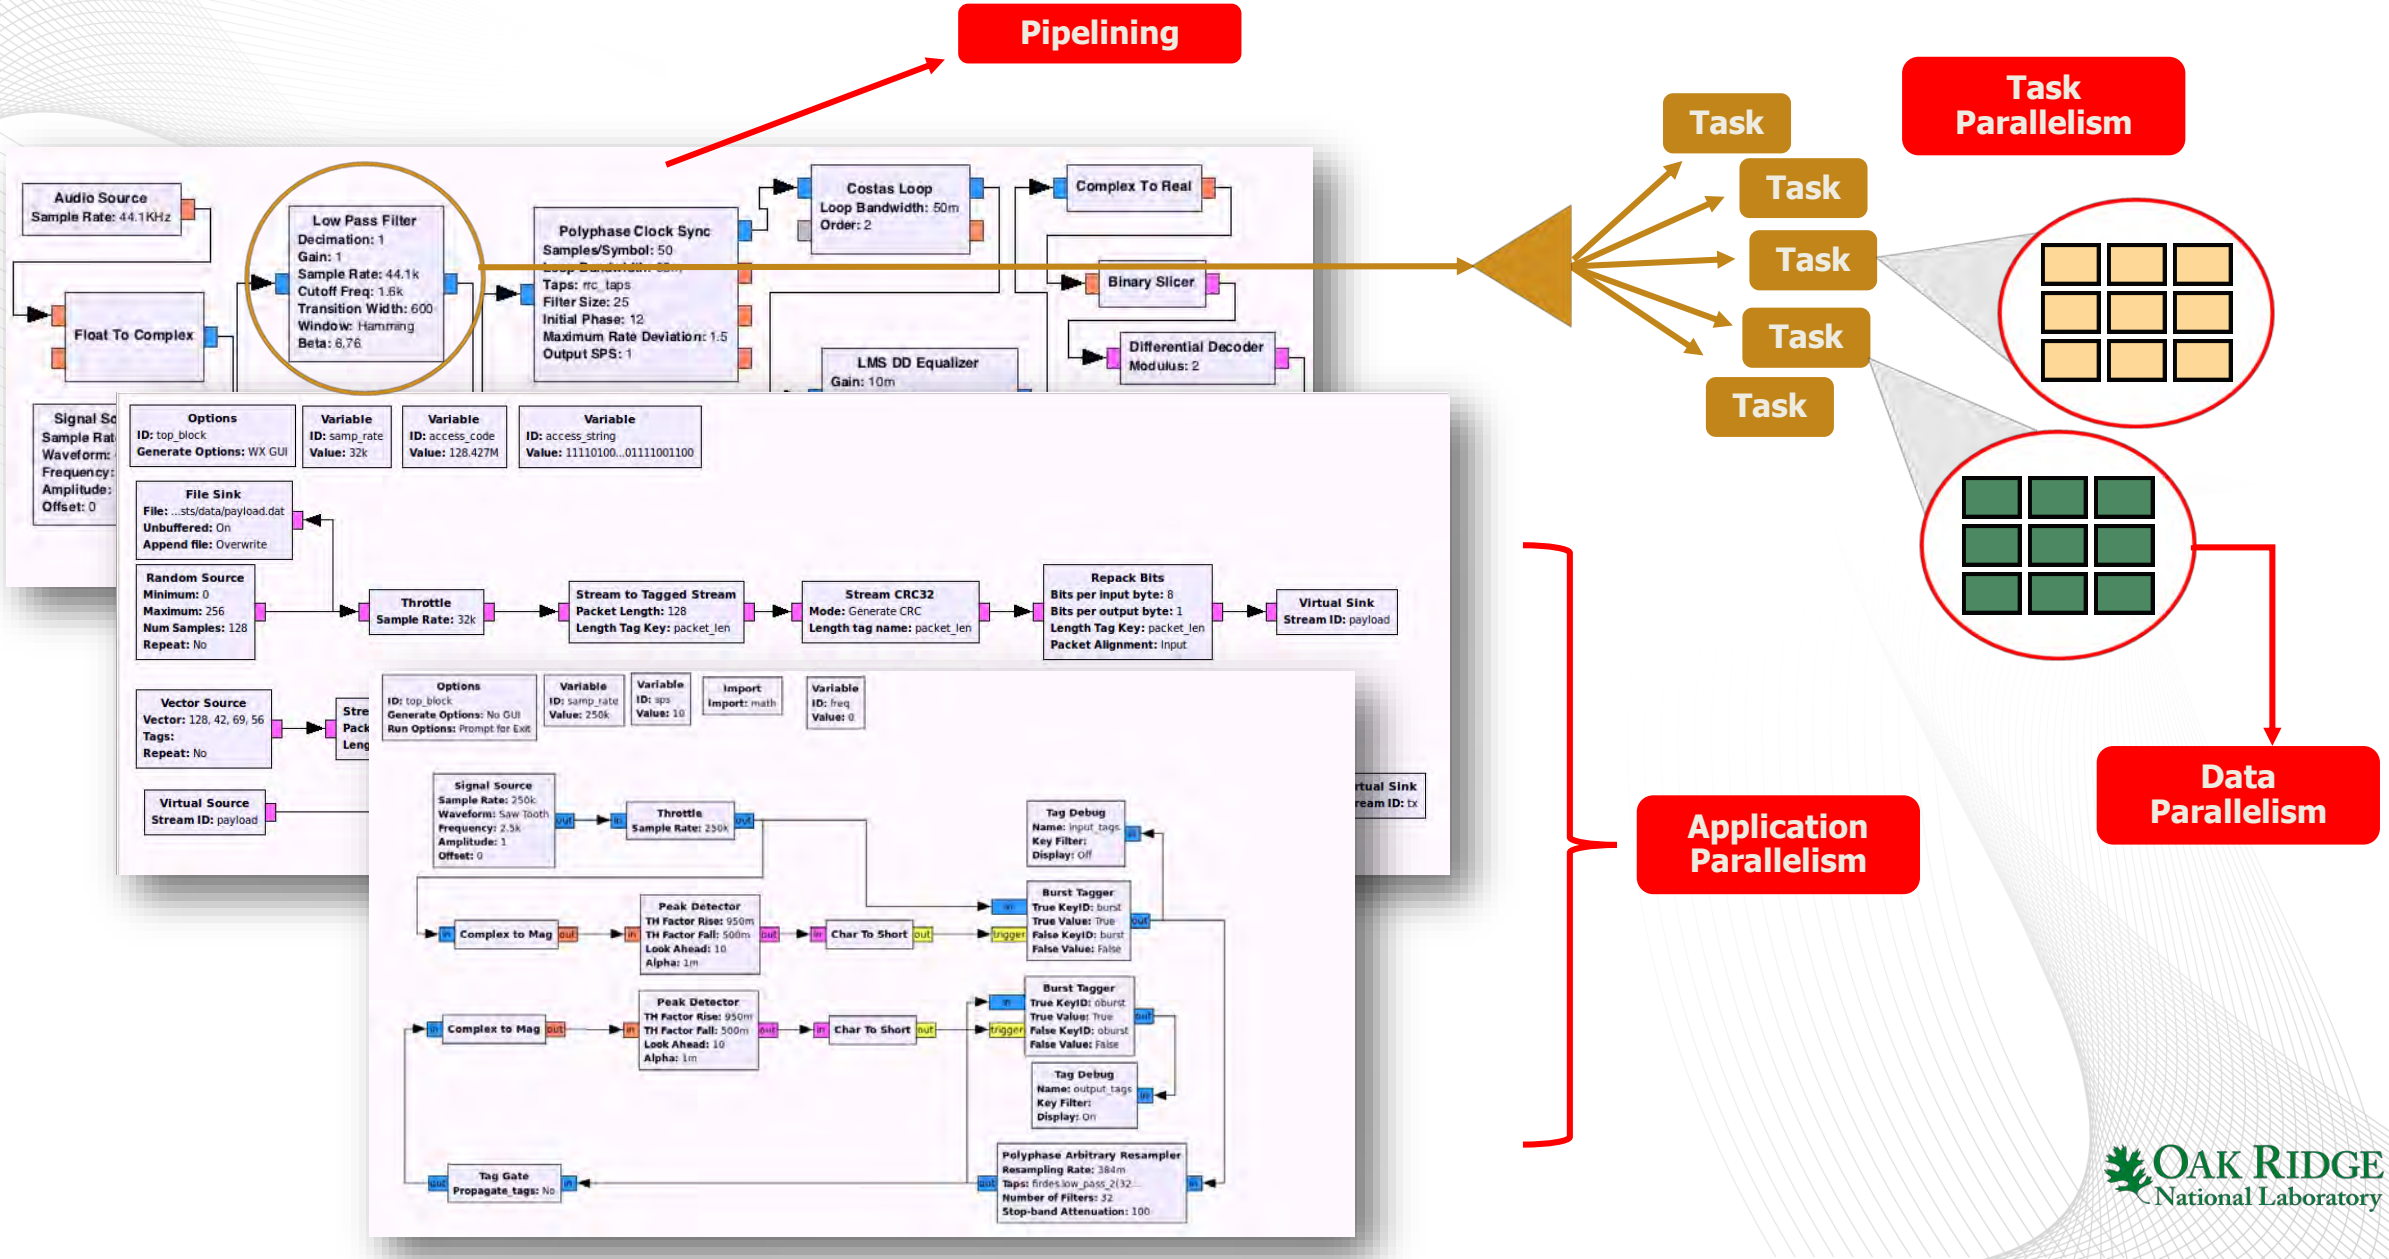

# Cosmic ISR / Prep

- ISR assumes that user application …
  - Written in high level, performance portable programming model (e.g., OpenACC or OpenCL if nec)
    - OpenARC – presented in last section
  - Has target code versions generated appropriate versions of target system code
    - JIT is possible for OpenCL targets (except FPGA)
- ISR contains specific RT modules for each device
- ISR sets up dependencies as specified by compiler, user
- ISR creates catalog of data to orchestrate data movement across disparate device memories

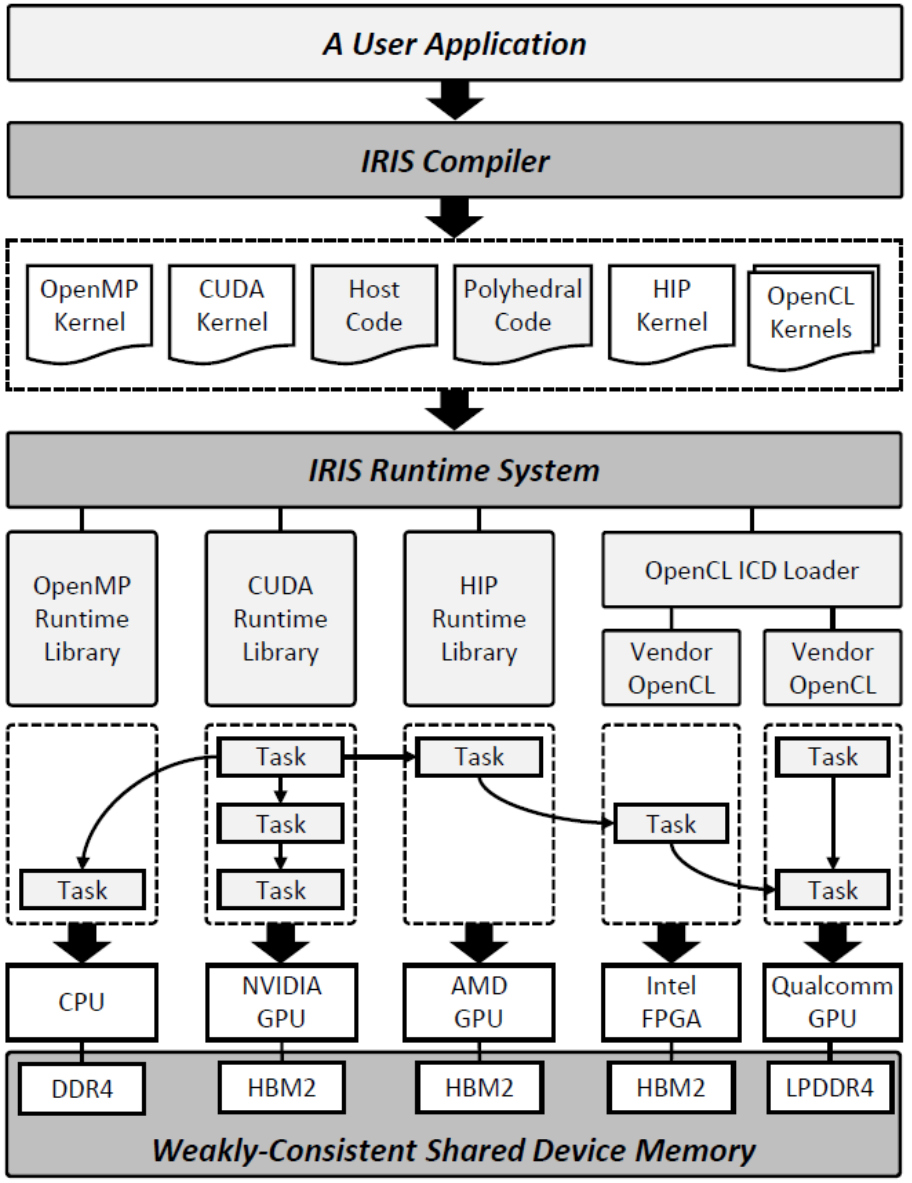

# Cosmic ISR / Execution

- During execution, ISR must
  - Discover available devices
  - Pick the most appropriate device for the task
  - Maintain dependencies
  - Orchestrate data movement
- Device selection uses any number of policies
  - Random, Round-robin, profiling, hints, ontology, performance models (Aspen)
  - ISR must also monitor existing device work to make tradeoffs
- Current support for
  - AMD GPU
  - NVIDIA GPU
  - CPU
  - Xeon Phi
  - Intel FPGA

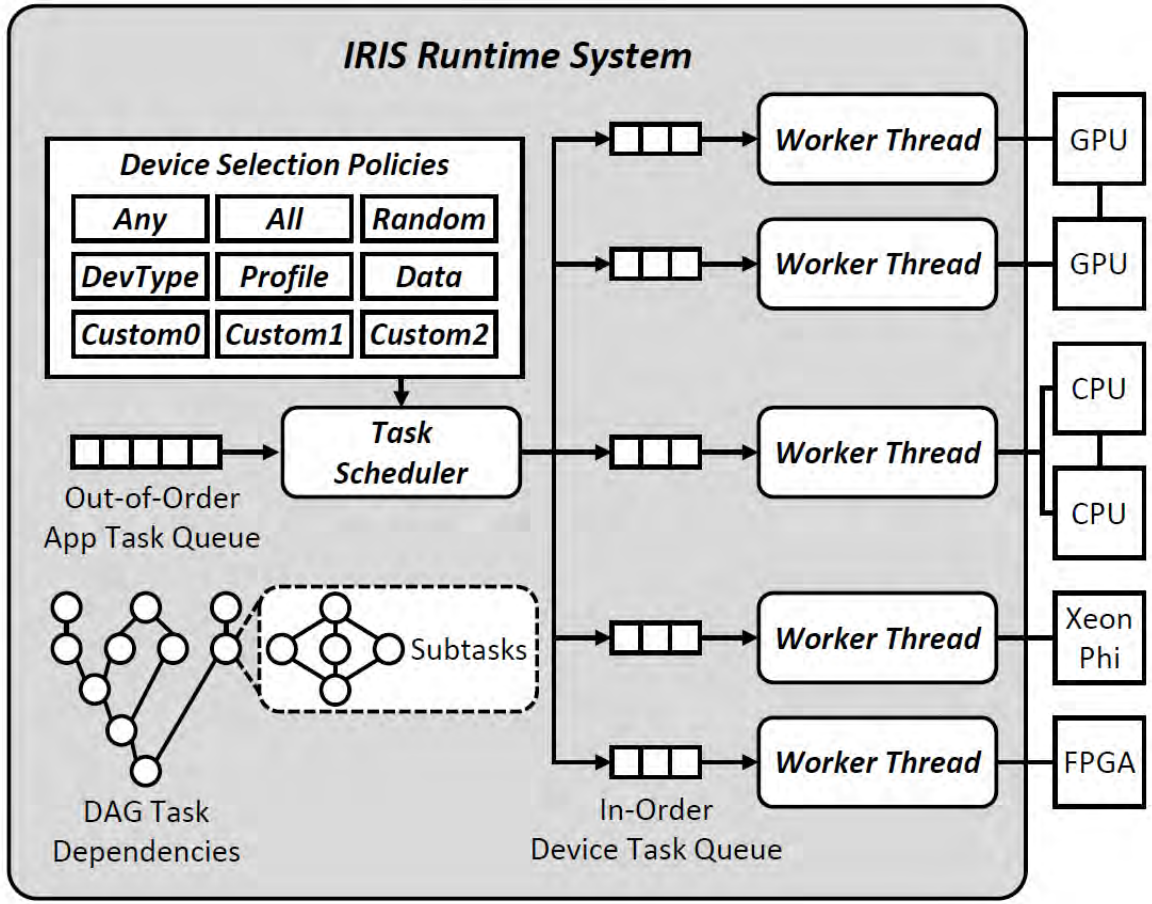

# Introspective Layer

- OpenCL provides a good compatibility layer but doesn't provide sufficient introspective feedback:
  - Performance counters
  - Interference counters
  - Power/Energy metrics
  - Temperature sensors

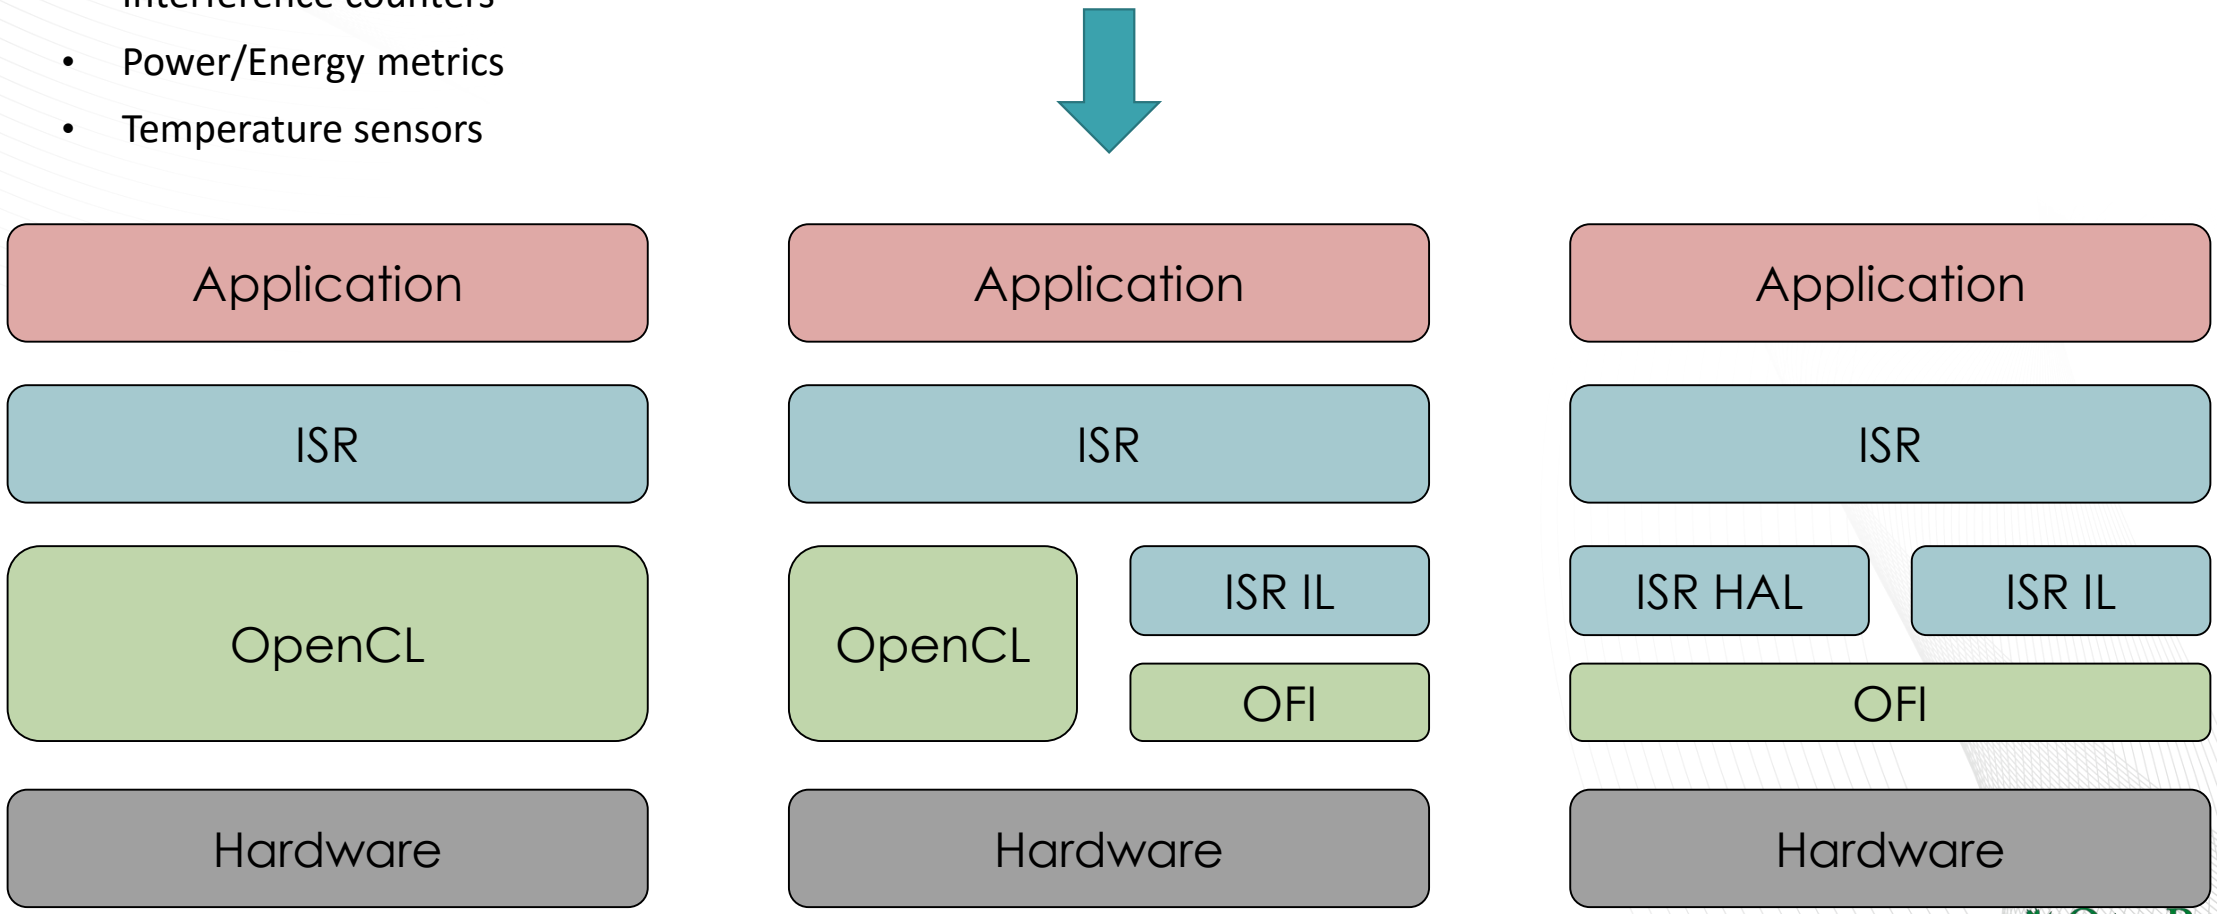

# Recap

- Motivation: Recent trends in computing paint an ambiguous future
  - Multiple architectural dimensions are being (dramatically) redesigned: Processors, node design, memory systems, I/O
  - Complexity is our main challenge
- Applications and software systems across many areas are all reaching a state of crisis
  - Need a focus on performance portability
- ORNL Cosmic project investigating design and programming challenges for these trends in SDR
  - Performance modeling and ontologies
  - Performance portable compilation to many different heterogeneous architectures/SoCs
  - Intelligent scheduling system to automate discovery, device selection, and data movement
  - Targeting wide variety of existing and future architectures (DSSoC and others)

- Visit us
  - We host interns and other visitors year round
    - Faculty, grad, undergrad, high school, industry
- Jobs in FTG
  - Postdoctoral Research Associate in Computer Science
  - Software Engineer
  - Computer Scientist
  - Visit https://jobs.ornl.gov
- Contact me vetter@ornl.gov

# Bonus Material

# ASCR Extreme Heterogeneity Workshop

## January 23-25, 2018 Virtual Meeting

https://orau.gov/exheterogeneity2018/
https://doi.org/10.2172/1473756

- **Goal: Identify Priority Research Directions for Computer Science needed to make future supercomputers usable, useful and secure for science applications in the 2025-2040 timeframe**
  - Note that quantum computing was defined as out of scope by ASCR.
- Primary focus on the software stack and programming models/environments/tools
- 150+ participants: DOE labs, academia, and industry
- White papers solicited (106 received!) to contribute to the FSD, identify potential participants, and help refine the agenda
- First ASCR workshop to use Basic Research Needs format (BES inspired)
  - Summit, Summit report, Factual Status Document, whitepapers, BRN/PRD result
- Organizing Committee
  - Jeffrey Vetter (ORNL), Lead Organizer and Program Committee Chair
  - Ron Brightwell (Sandia-NM), Pat McCormick (LANL), Rob Ross (ANL), John Shalf (LBNL)
  - Lucy Nowell, ASCR Program Manager
- Program Committee Members
  - Katy Antypas (LBNL, NERSC), David Donofrio (LBNL), Maya Gokhale (LLNL), Travis Humble (ORNL), Catherine Schuman (ORNL), Brian Van Essen (LLNL), Shinjae Yoo (BNL)

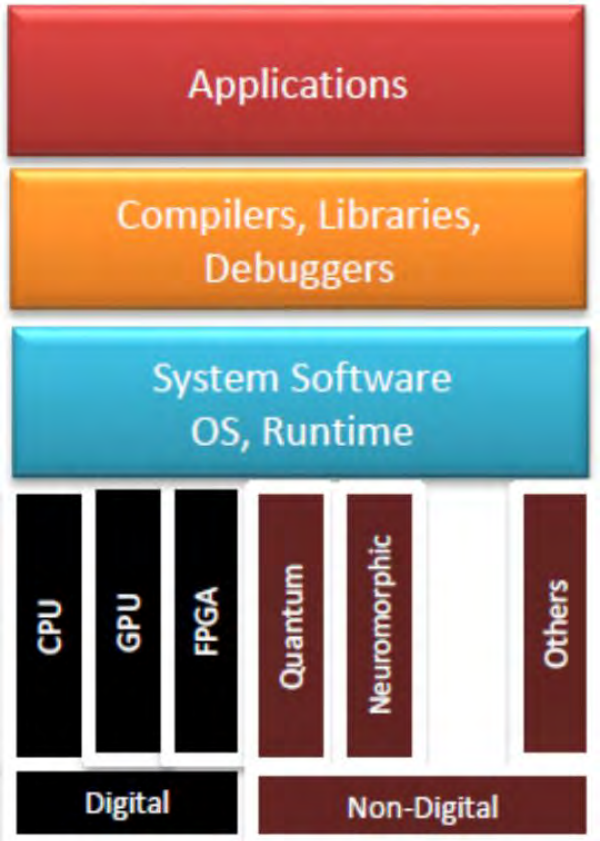


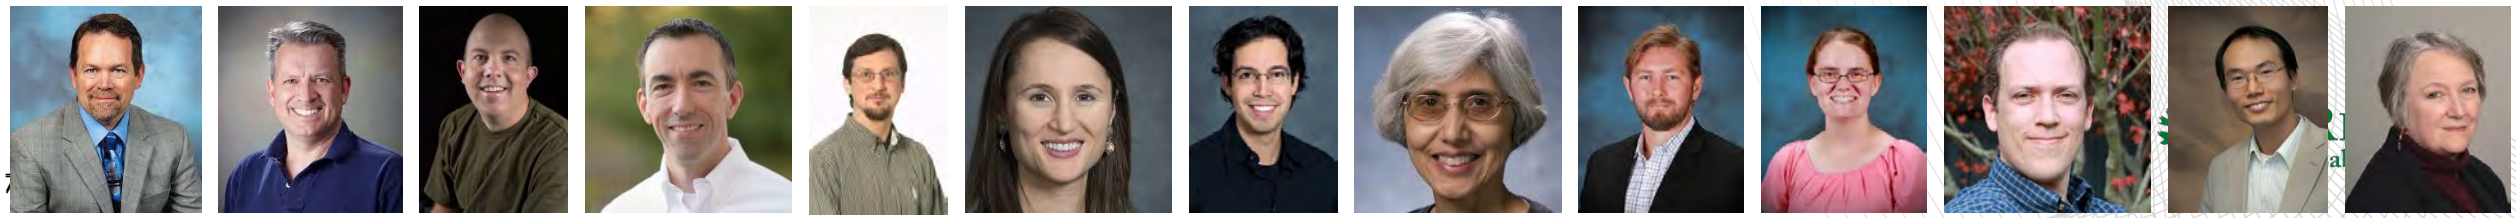

# Future Technologies Group (FTG)

*Jeffrey S. Vetter, Group Leader*

The Future Technologies Group performs research in core technologies for emerging generations of high-end computing architectures, including prototype computer architectures and experimental software systems. We investigate these technologies with the goal of improving the performance, energy efficiency, reliability, and productivity of these architectures for our sponsors and applications teams. See http://ft.ornl.gov.



https://www.thebroadcastbridge.com/content/entry/1094/altera-announces-arria-10-2666mbps-ddr4-memory-fpg

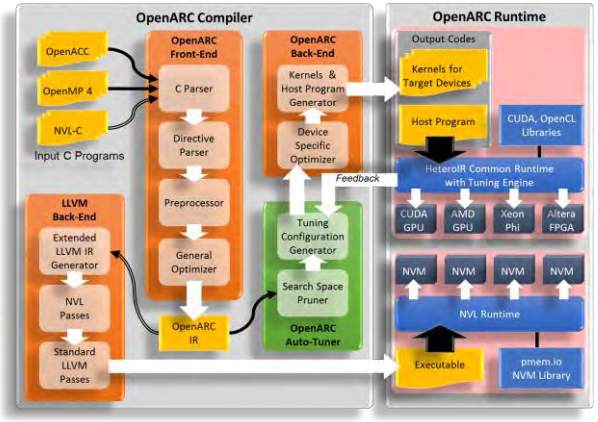


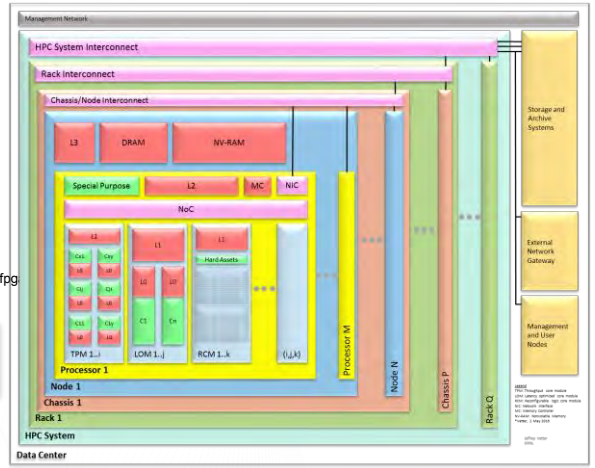


## Key Technical Areas

- Heterogeneous architectures
- Deep memory hierarchies including non-volatile memory
- Performance measurement, analysis, simulation, and modeling of emerging architectures.
- Programming systems to address emerging architectures
- Beyond Moore's Computing

## Software Artifacts

- Scalable Heterogeneous Computing Benchmarks (SHOC)
- mpiP
- DESTINY
- Aspen
- OpenARC
- Papyrus
- NVL-C
- Oxbow
- LLVM Clacc and Parallel IR
- DRAGON
- RISC-V Extensions

## Sponsors

- DOE ASCR, BER
- DOE Exascale Computing Project
- DOE SciDAC
- DARPA
- ORNL LDRD
- National Science Foundation
- Department of Defense
- NIH

## Impact

- Publications in SC, ICS, HPDC, TPDS, DATE, PLDI, IPDPS, Trans VLSI, etc.
- Two Gordon Bell awards
- NSF Keeneland
- DOE Titan
- IEEE TCHPC Early Career
- IEEE Fellows
- ~100 interns
- ~130 FTG seminars

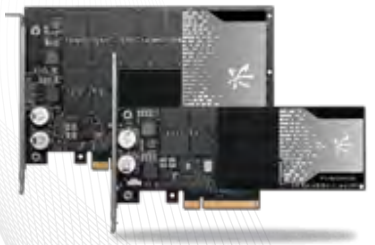

FusionIO

# Progression of Experimental Computing Technologies

**TRL 1-3 Basic Concepts**

- *Examples: carbon-nanotube computing, memristor-based neuromorphic computing, chip-level silicon photonics, universal quantum computing*

**TRL 4-6 Emerging**

- *Examples: FPGAs in HPC, TrueNorth, SpiNNaker, D-Wave, Emu, many SoC-based systems, TPU, Gen-Z NoCs, near-memory computing*

Evaluate, Select, and Improve Emerging Computing Technologies

**TRL 7-9 Operational**

- *Examples: Titan, Cori, Mira, Summit, BlueWaters, Keeneland, Stampede, Tsubame2.5*

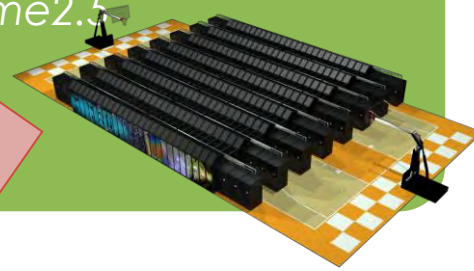

"Bench" System

Limited Access Testbed

Experimental Prototype

Production

CS & Math Research

| | | | | |
|---|---|---|---|---|
| **Programming** | Assembly language, or less | Few, if any, development tools | Language support and compilers. |
| **OS-R** | Manual | Specialized programming environments and OSs | Commodity OS & runtime systems |
| **Scale** | Small collections of devices | Single to hundreds of engineered processing elements | >10,000 processing elements |
| **Performance** | Analytical projections based on device empirical evaluation. | Analytical projections or simulation based on component or pilot system empirical evaluation. | Empirical evaluation of prototype and final systems. |
| **Apps** | Small encoded kernels | Architecture-aware algorithms; Mini-apps; Small applications | Numerical libraries; Full scale applications |
| **Example** | GPUs invented in 1999 | OpenGL in 2001; CUDA in 2007; OpenCL in 2008; OpenACC in 2010; DP in 2010; ECC in 2012 | GPUs are a fully supported compute technology in the HPC ecosystem. |

# ORNL ExCL Model

https://excl.ornl.gov

- Provide low-level access to emerging computer architectures to encourage experimentation and prototyping of new hardware and software solutions.
- Not just testbeds, but staff and software environments to support this mode of operation.

## ExCL Common Infrastructure

**Project and User management**
- Accounts
- Projects and Proposals
- Help

**Community**
- Workshops
- Online discussions forums and issues
- Consolidated
- News

**Shared Login and Gateway Nodes**
- Gateway nodes
- Data transfer nodes
- Consistent and secure access to private network compartments

**Authentication and Authorization**
- Secure operations
- Partition access to specific compartments
- System and account lifecycles
- Experience with management of export controlled and proprietary systems

**Shared Filesystems and Databases**
- Secure access to filesystems across pillars

**Monitoring and control systems**
- Manage access to shared resources
- Manage privileged access levels
- Lights out operation

**Source Code and Data sets**
- Source Code repos
- Performance databases for applications and architectures

**Web**
- Educational and reference materials
- Outreach
- Both Open and Controlled access

## ExCL Technology Pillars

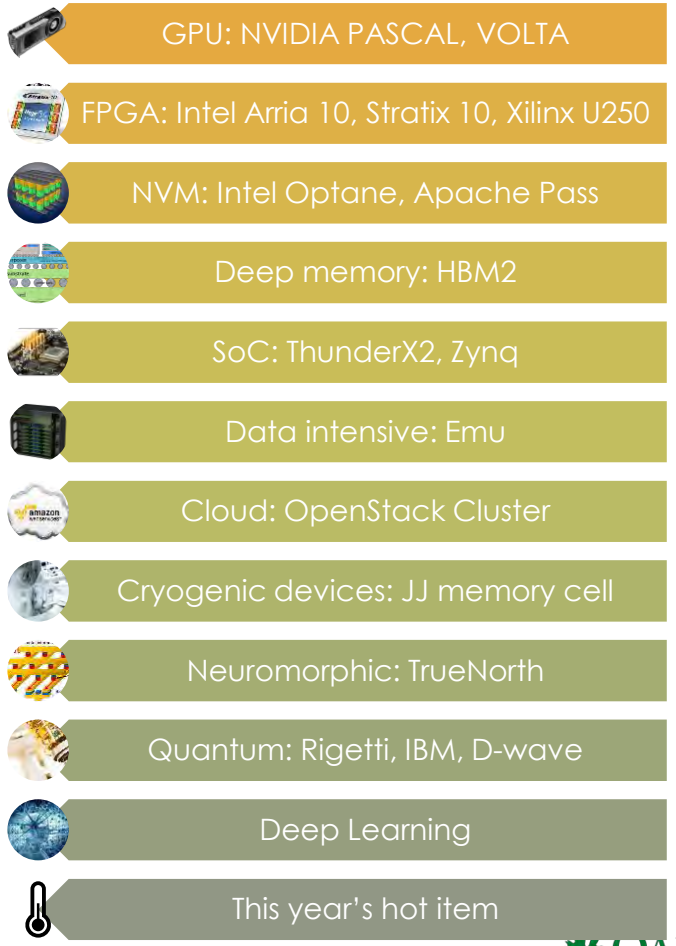


Pillars refreshed annually