Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

Bridging the Assessment Modeling Gap the PERFECT Way

KEVIN J BARKER, DARREN J KERBYSON, ADOLFY HOISIE, ANDRES MARQUEZ, JOSEPH MANZANO, ROBERTO GIOIOSA, ANTONINO TUMEO, NATHAN TALLENT, GOKCEN KESTOR, LEON SONG

Pacific Northwest National Laboratory

Modeling and Simulation Workshop, Seattle 2014

Test and Validation (TAV) in DARPA PERFECT

- DARPA's PERFECT (Power Efficiency Revolution for Embedded Computing Technologies)
  - Spearheading R&D into a multitude of diverse technologies
  - 75 Gflops/W for general-purpose embedding computing
  - Envisioning 7nm technologies in 2018 – 2020 timeframe
- Program split across 3 phases, beginning at the end of 2012
- 17 initial *Performer* teams
  - Device technology, Architecture, Systems, Software, and Optimization
- Test and Validation (TAV) team
  - Quantitatively assess PERFECT technologies individually with respect to the overall system & overall performance and power goals
  - Use a combination of benchmarking, modeling and simulation

PERFECT Assessment: Challenges and Strategy

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- PERFECT Challenge
  - Goal: Embedded system delivering "75 GFLOPS/W"
  - Performers contribute only part of a system (architecture to algorithms)
  - TAV must assess Performer's contribution w.r.t. entire system

- Three pillars to assessment strategy
  - Baseline Architectures: quantifying today's state of the art
  - PERFECT Suite: defining a workload
  - Proxy Architecture: modeling framework

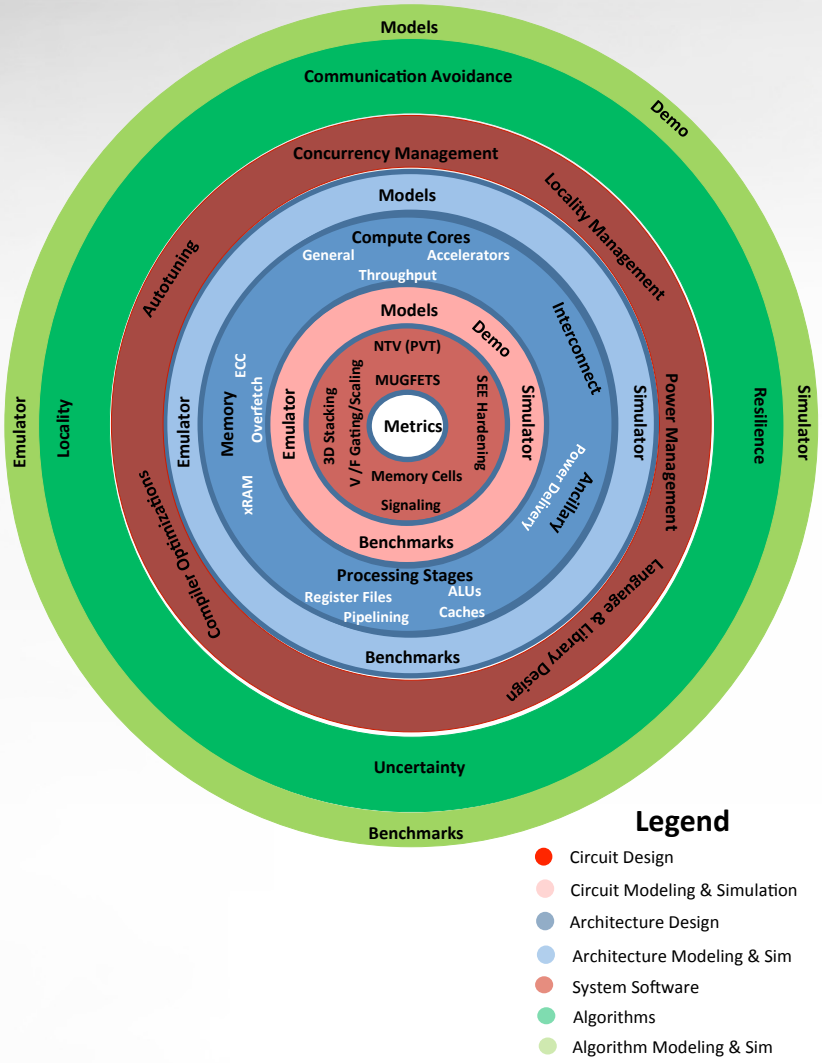

Modeling and Simulation Challenges

- Integrating Performance and Power Prediction
  - Overall PERFECT goals stipulate *efficiency*
  - PERFECT TAV effort has defined methodologies for Performance and Power prediction thrusts
- Defining a suitable interface with Performer teams
  - What information is required to parameterize the models?
  - What is the appropriate level of architectural abstraction?
- Although not the goal of PERFECT, how can these methodologies be extended to large-scale systems?
  - Potential for combining with existing scalability modeling methodologies with PERFECT TAV tools

TAV Overall Approach: 3 Pillars

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

Modeling and Simulation focus

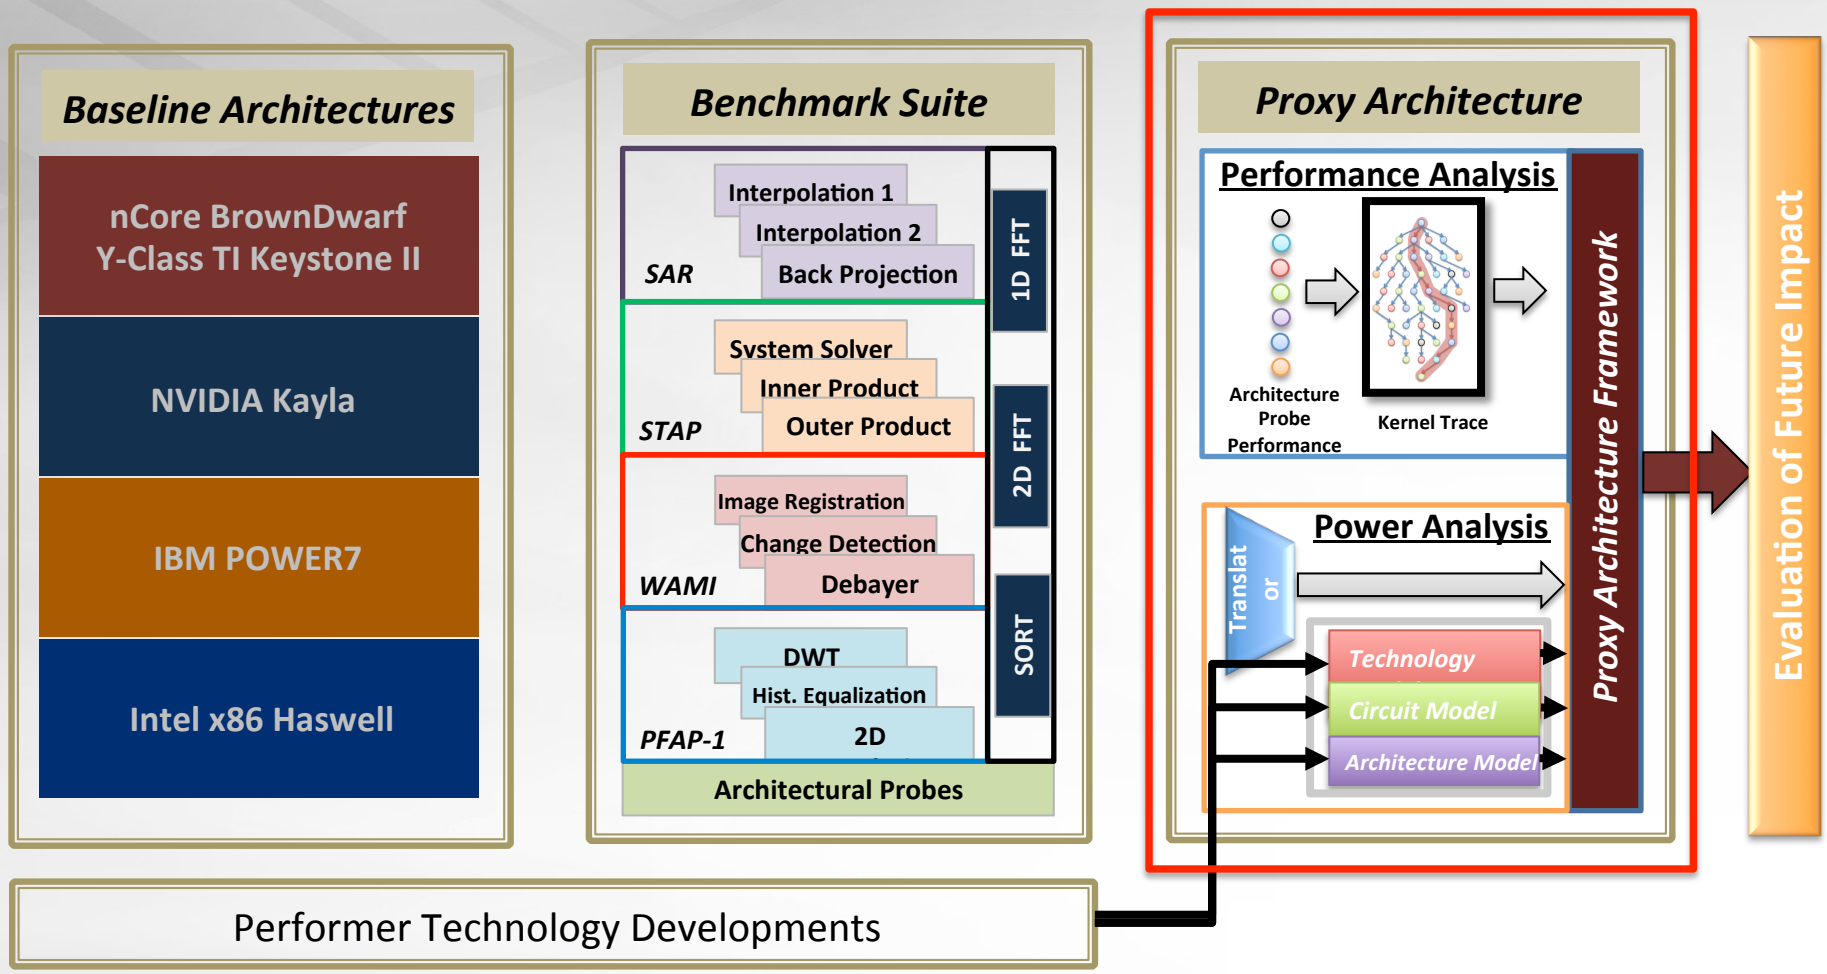

PERFECT TAV Pillar 1: Baseline Architectures

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- Architectures that reflect state-of-the-art systems in Phase 1
- Real world data points for power/ performance model validation
  - Calibration of Performer's modeling and simulation environments
  - Validation of TAV-developed models on existing platforms

Platform	# Cores (Threads)	Peak Perf (Gflops)	Clock (GHz)	Peak Power (Watts)	Gflops per Watt	Memory (GB)
nCore BD-Y TI Keystone II	16+96 (ARM + DSP)	614.4 (SP)	1.2 + 1.4	36-56	17.1 – 11 (SP)	56
Nvidia Kayla	4+2(384) (ARM+SMX)	300 (SP)	1.2 / 1.05	27	11.1	2/1
IBM Power7	8(32)	265 (DP)	4.2	240	1.1	16
Intel X86 Haswell	4(8)	295 (SP)	2.3	45	6.5	16

- Power Instrumentation:
  - Level 1: Watts-up power meters
  - Level 2: Internal architecture-supplied counters (e.g., RAPL, Ameester, etc.)
  - Level 3: High-fidelity DAQ
- Baseline Architectures in place in EHPC lab at PNNL and are being used


nCore


Kayla


Power7


Haswell

PERFECT TAV Pillar 2: PERFECT Suite

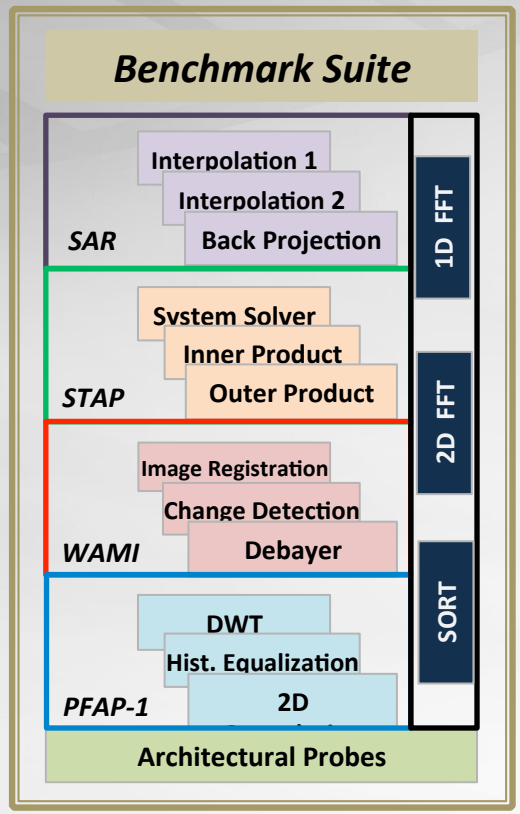

- Kernels and applications that represent application domains
- Benchmark-specific models will be validated on Baseline Architectures and used to predict performance and power consumption of Performer architectures
- Selection criteria
  - PERFECT's domain of interest
  - Alignment of app/kernels to Performer's projects
  - Reasonable input data set sizes selected

1. Synthetic Aperture Radar (SAR)
2. Wide Area Motion Imaging (WAMI)
3. Space Time Adaptive Processing (STAP)
4. PERFECT APPLICATION 1 (PFAP-1)
5. 3 "Core" Kernels (Sort, 1D and 2D FFTs)

- All serial and CUDA reference kernels available
  - http://hpc.pnnl.gov/projects/PERFECT

PERFECT TAV Pillar 3: Proxy Architecture Modeling Framework

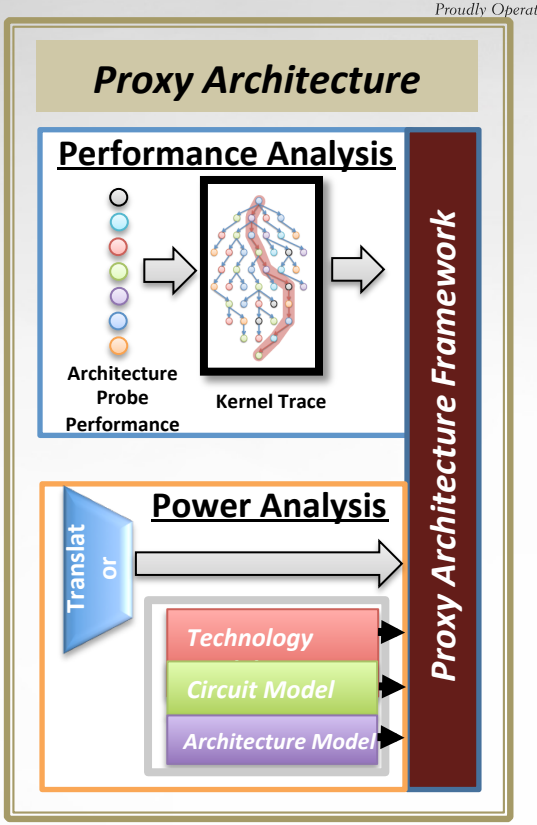

- An integrated approach to assist in the assessment of component technologies in the context of a complete architecture
- The *Proxy Architecture* is a mechanism that accommodates solutions from diverse technology research (architectures to algorithms)

Two thrusts:

- Performance Analysis
  - Benchmark-specific models parameterized in terms of architecture performance capabilities to derive predicted performance ranges
  - Solution providers define architectural operations, latencies, and throughputs
- Power Analysis
  - Utilizes McPAT, an open-source Power, Area, and Timing modeling framework
  - Input from Performer teams to define technology, circuit, and architecture parameters

Current Status of PNNL TAV Activities

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- Baseline Architectures installed in the EHPC lab at PNNL and are being remotely accessed by Performers
- Benchmark Suite in use by Performers and available to the community
  - Kernels available now; full scale applications available shortly
- Calibration of Performer's simulation environments in progress
  - Calibration against Baseline Architectures
  - Provides confidence in each simulation environment
- Proxy Architecture
  - Power thrust: P-McPAT Maintenance and Infrastructure
    - TAV P-McPAT Architectural Validation
  - Performance thrust: Internal validation of modeling methodology utilizing kernels from the TAV Benchmark Suite and Baseline Architectures
    - Initial decomposition of selected kernel into low-level operations
    - Architectural micro-kernel framework to measure latencies and throughputs

PERFECT TAV Summary

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- Key challenge being tackled by TAV in DARPA PERFECT is:
  - To quantitatively assess PERFECT technologies individually and with respect to the programs overall performance and power goals
- Approach is to use a combination of benchmarking, modeling, and simulation in 3 pillars:
  - Baseline Architectures
  - Benchmark Suite
  - Proxy Architecture
- Modeling and simulation work is in the Proxy Architecture pillar, and is the primary focus of current research
- Approach is not restricted to DARPA PERFECT, but is generally applicable for assessing the potential of future technologies
  - Unified approach capturing performance and power impacts
  - Current effort on integrating resilience modeling as well

Gaps and Opportunities

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- PERFECT TAV strategy defines a consistent evaluation methodology
  - Targeting diverse embedded computing technologies and architectures
  - However, we are not limited to embedded systems; strategy can be applied to systems across scales
- What are the gaps?
  - Architectural specification and benchmarking (e.g., metrics)
  - Third component of PPR – Resilience
- What are the opportunities?
  - Opportunity for large-scale modeling tool kit from “first principles”
  - Need well-defined interfaces between modeling layers
    - “Bag of tools” approach will allow different capabilities to be “plugged in”
    - Modeling tools selected based on desired levels of abstraction
    - Encourage interaction between researchers, designers, and vendors