Traleika Glacier

Shekhar Borkar

Acknowledgment: TG Team

Intel Corp.

May 28, 2014

Outline

- Technology outlook and challenges
- Vision & Status
- SW Stack for both:
  - Evolutionary & Revolutionary approaches
- Open community runtime for research
- Applications and results
- Summary

Exascale Technology Challenges

NTV Logic & Memory for low energy

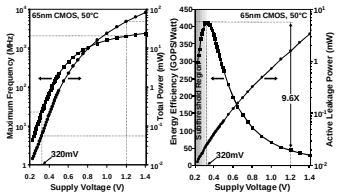

Break the Vmin barrier
FPU, logic & latches
Register Files,
SRAM

Extreme parallelism *O(billion)*

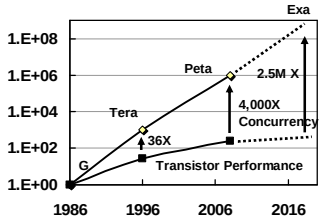

Programming model
Data locality
Legacy compatibility

Fine grain energy & power management

Voltage Regulator
Buck or Switched Cap
Power gating, frequency control

New introspective execution model

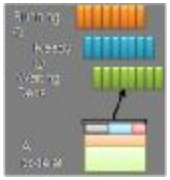

Tiny threads
Synchronization
Dynamic scheduling
Runtime system

Hierarchical, heterogeneous IC fabric

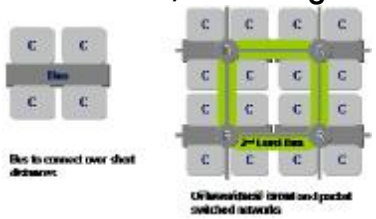

Busses
X Bars
Circuit & Packet
Switched

Self awareness

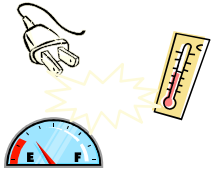

Observation based
Monitor, continuously adapt
Objective function based
runtime optimization

Data movement becomes expensive

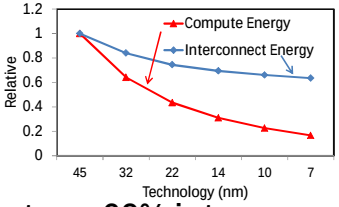

6X compute vs 60% interconnect energy

System level resiliency research

Straw-man HW System Architecture

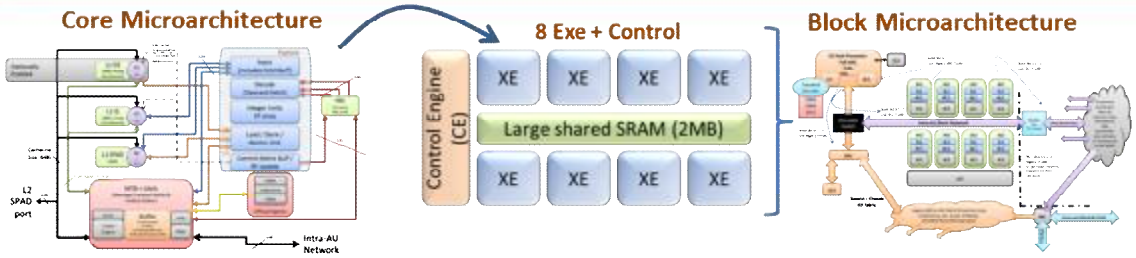

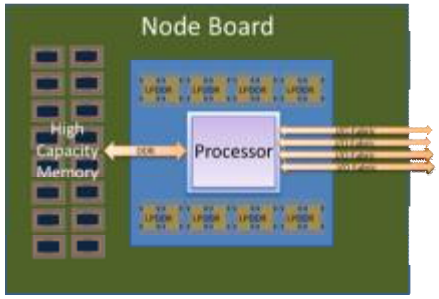

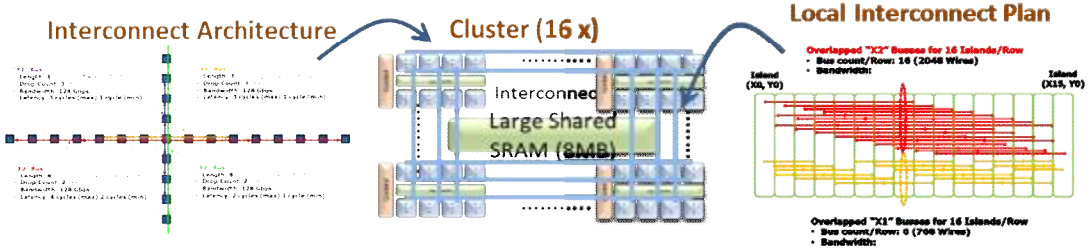

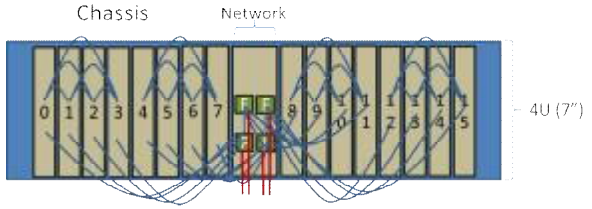

Exascale System (2022)

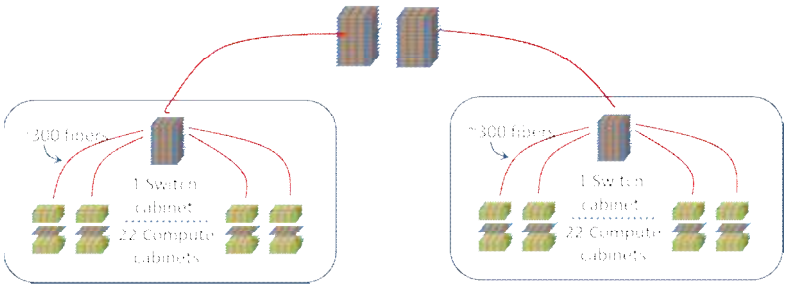

1 EF Peak, 20 MW

Switch Cabinet

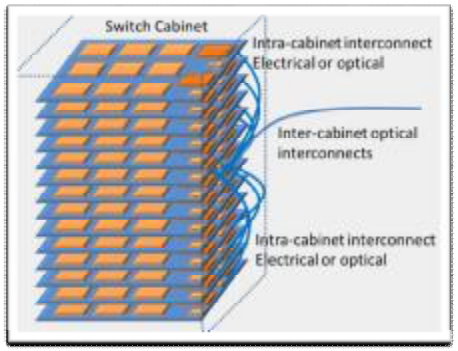

Cabinet

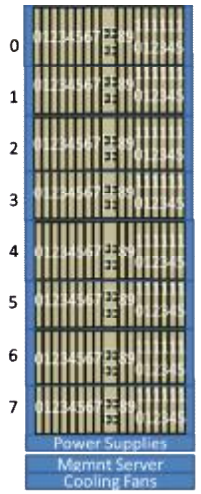

X-Stack (TG) Software Stack

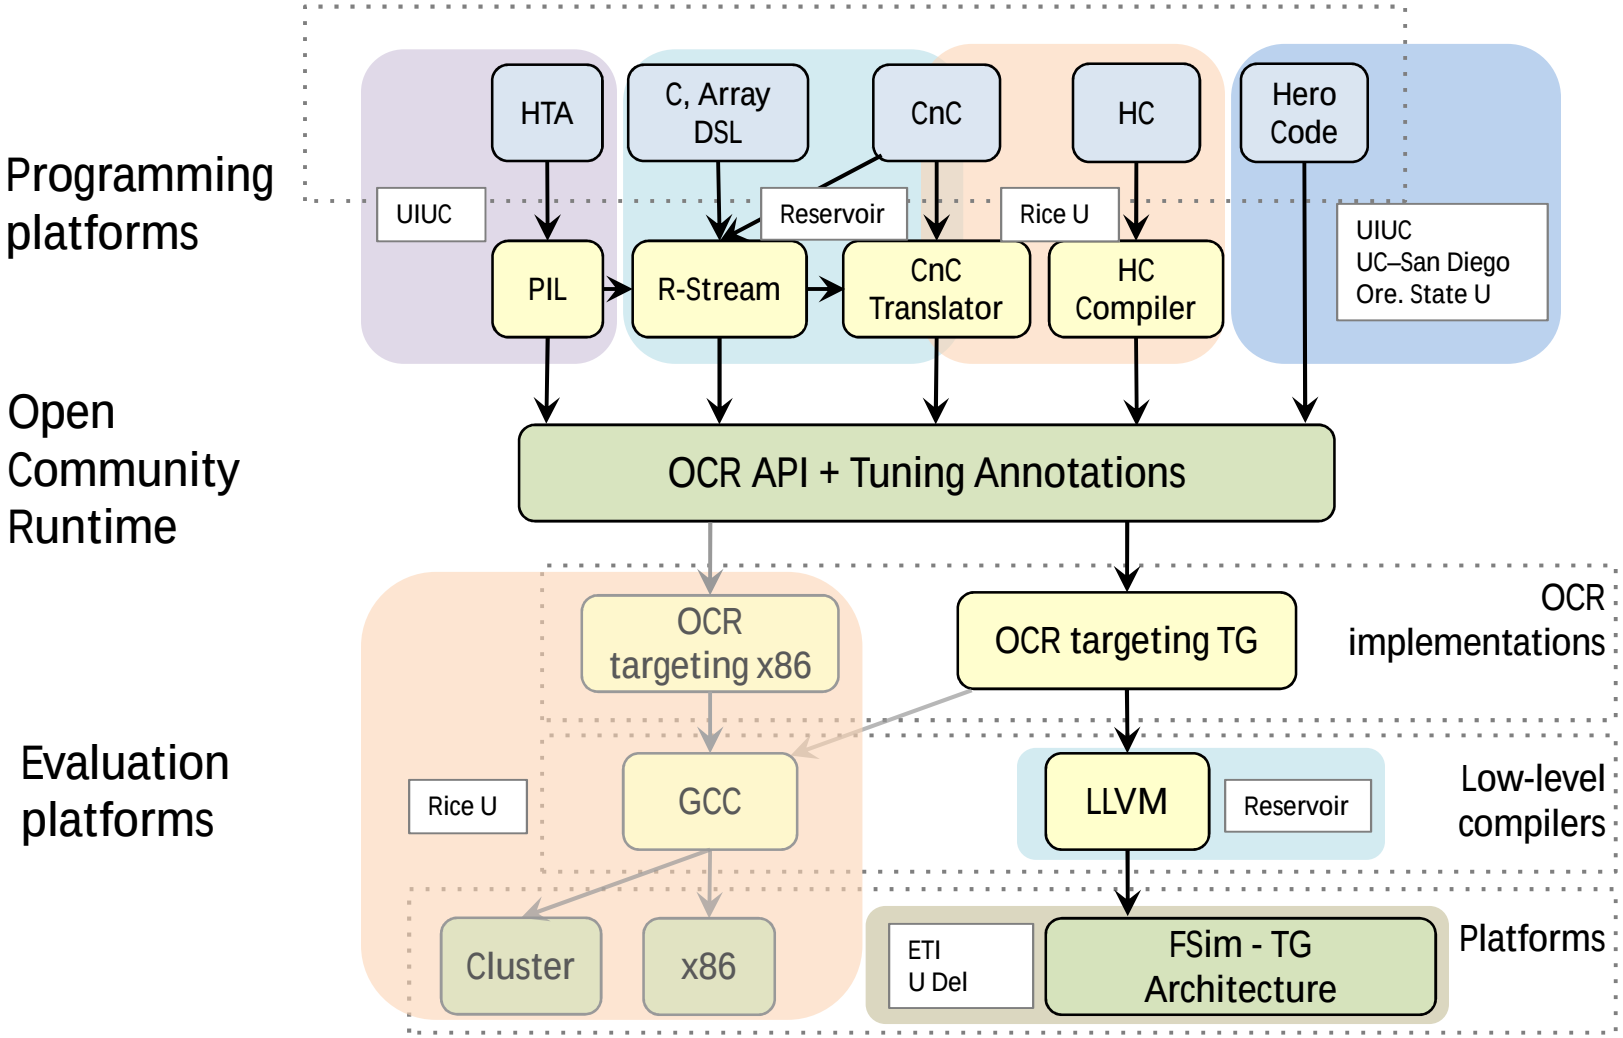

Legacy Support

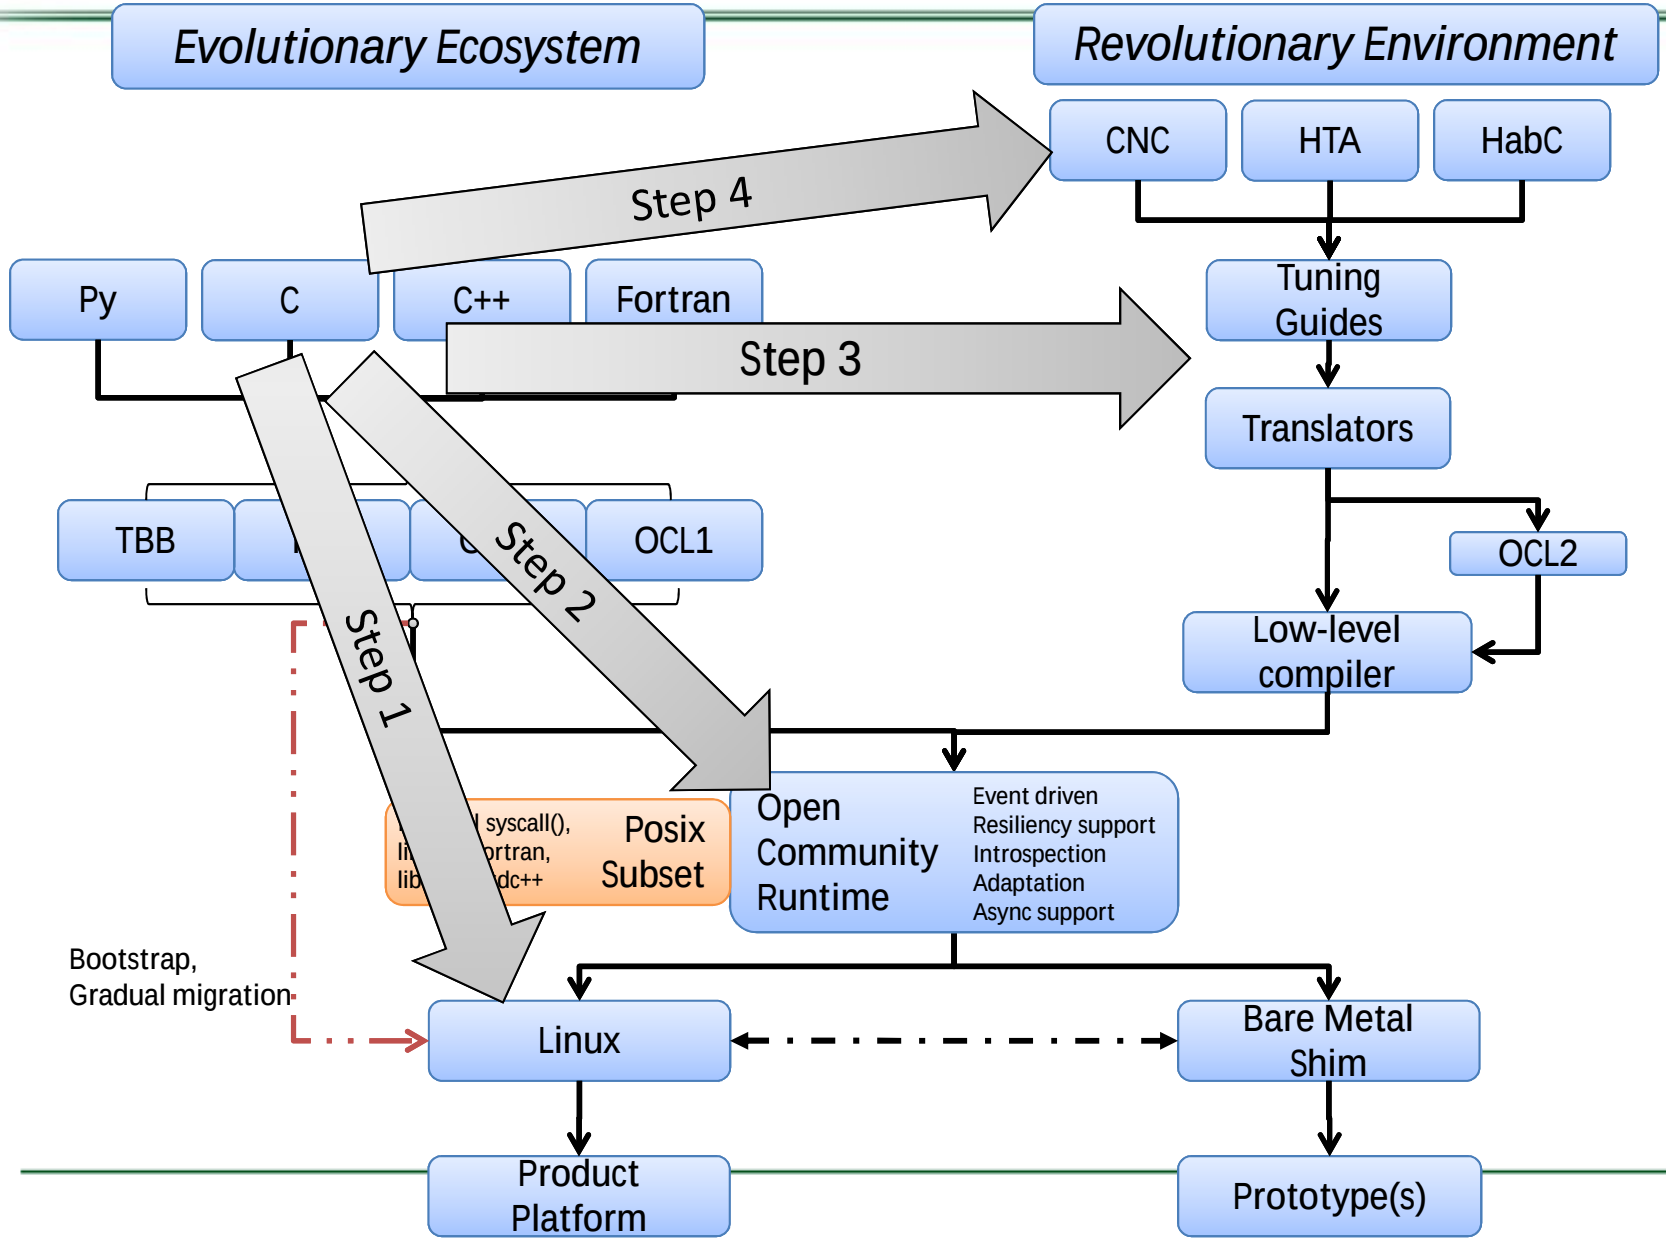

Software Components Put Together for Evaluation

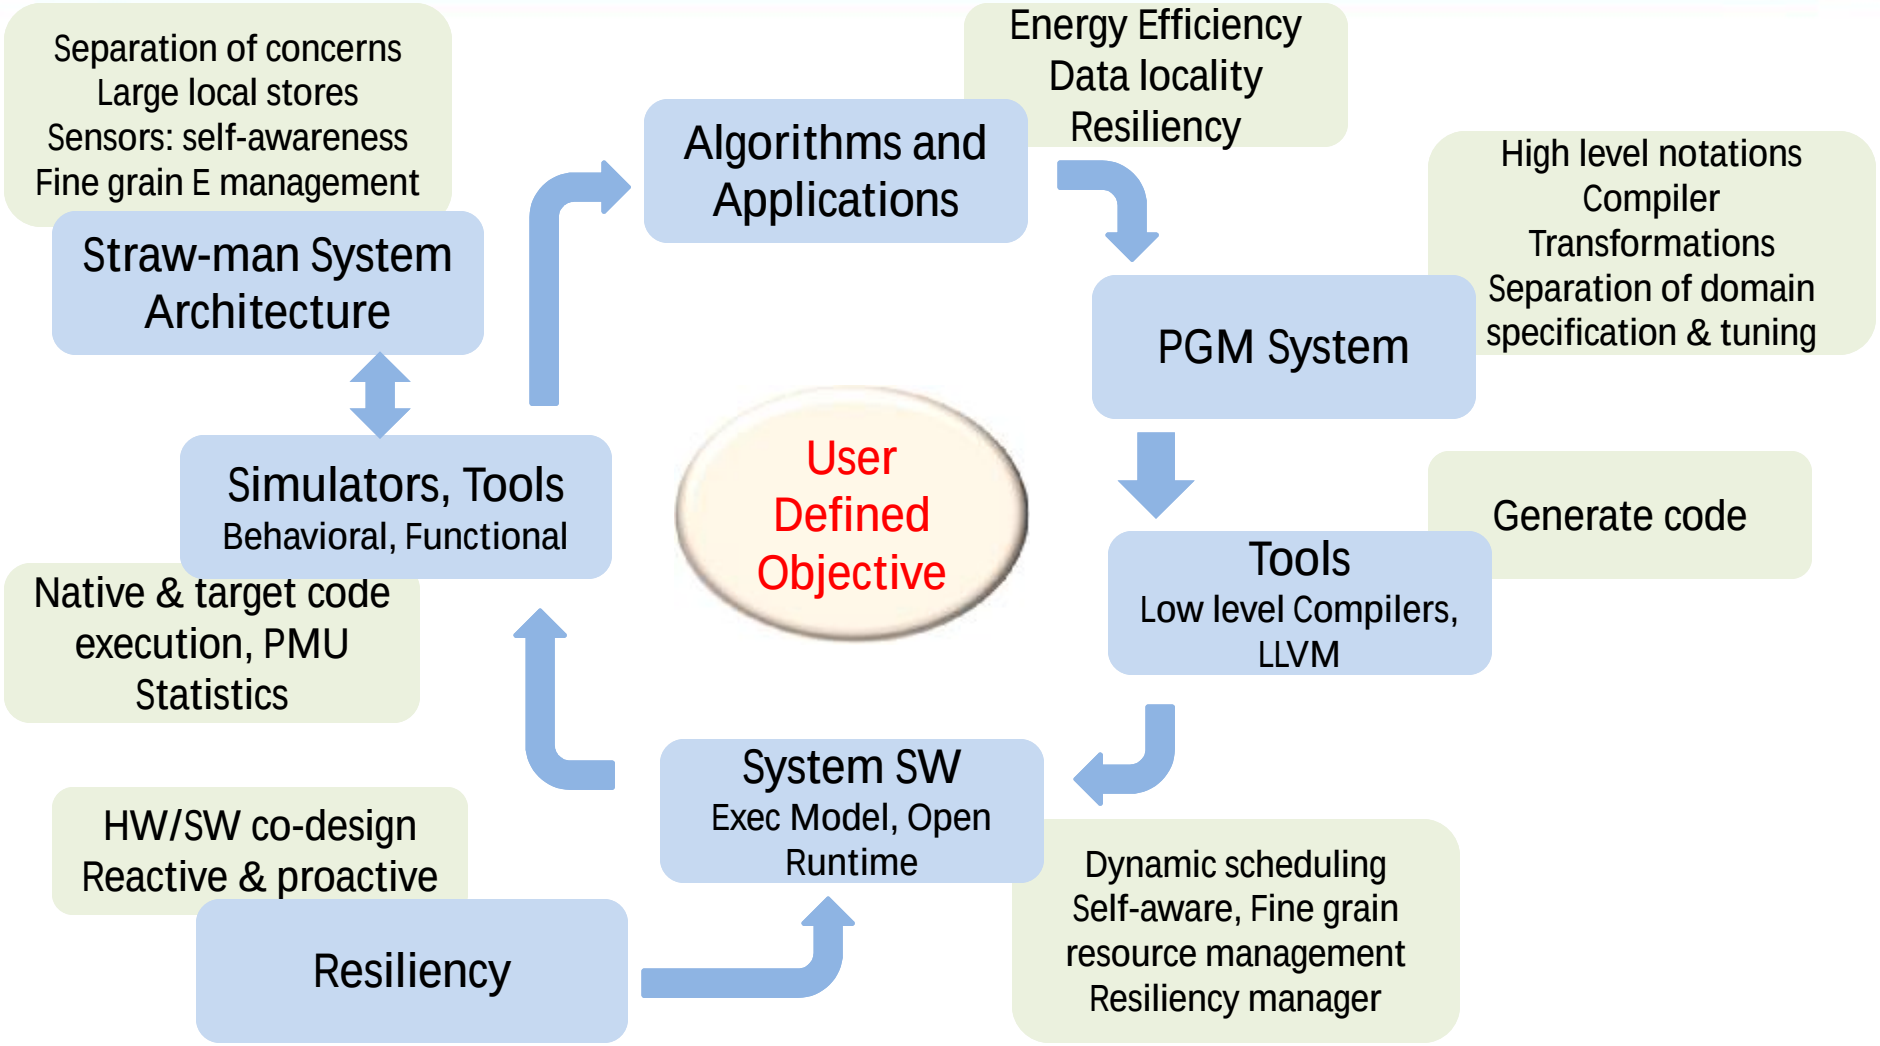

7

Dataflow-inspired Programming Model

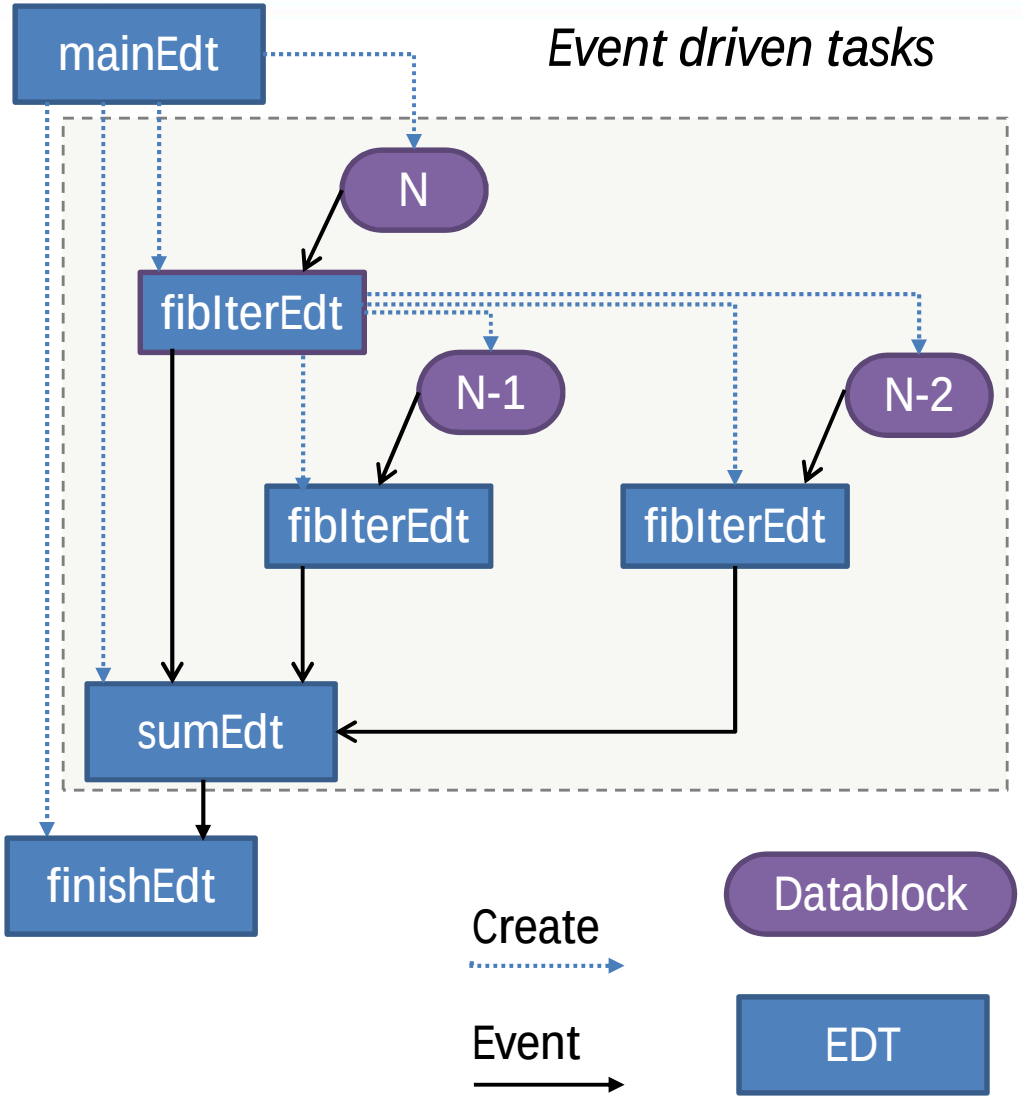

Runtime maps the constructed data-flow graph to architecture

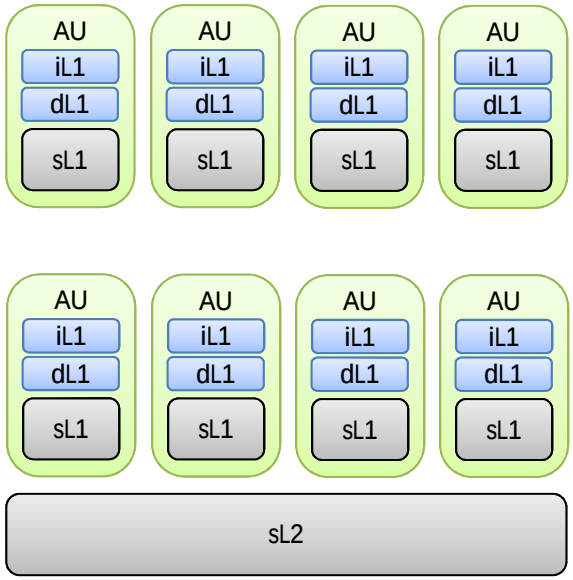

Open Community Runtime—Research Platform

- A research platform to evaluate revolutionary concepts
  - Event driven programming model
  - Introspection based resource management
  - Self-awareness
  - Resiliency
- Provides framework for future research
- Provides a reference implementation
- Provides runtime statistics for evaluation
  - Energy consumption
  - Data movement and computation

OCR modules

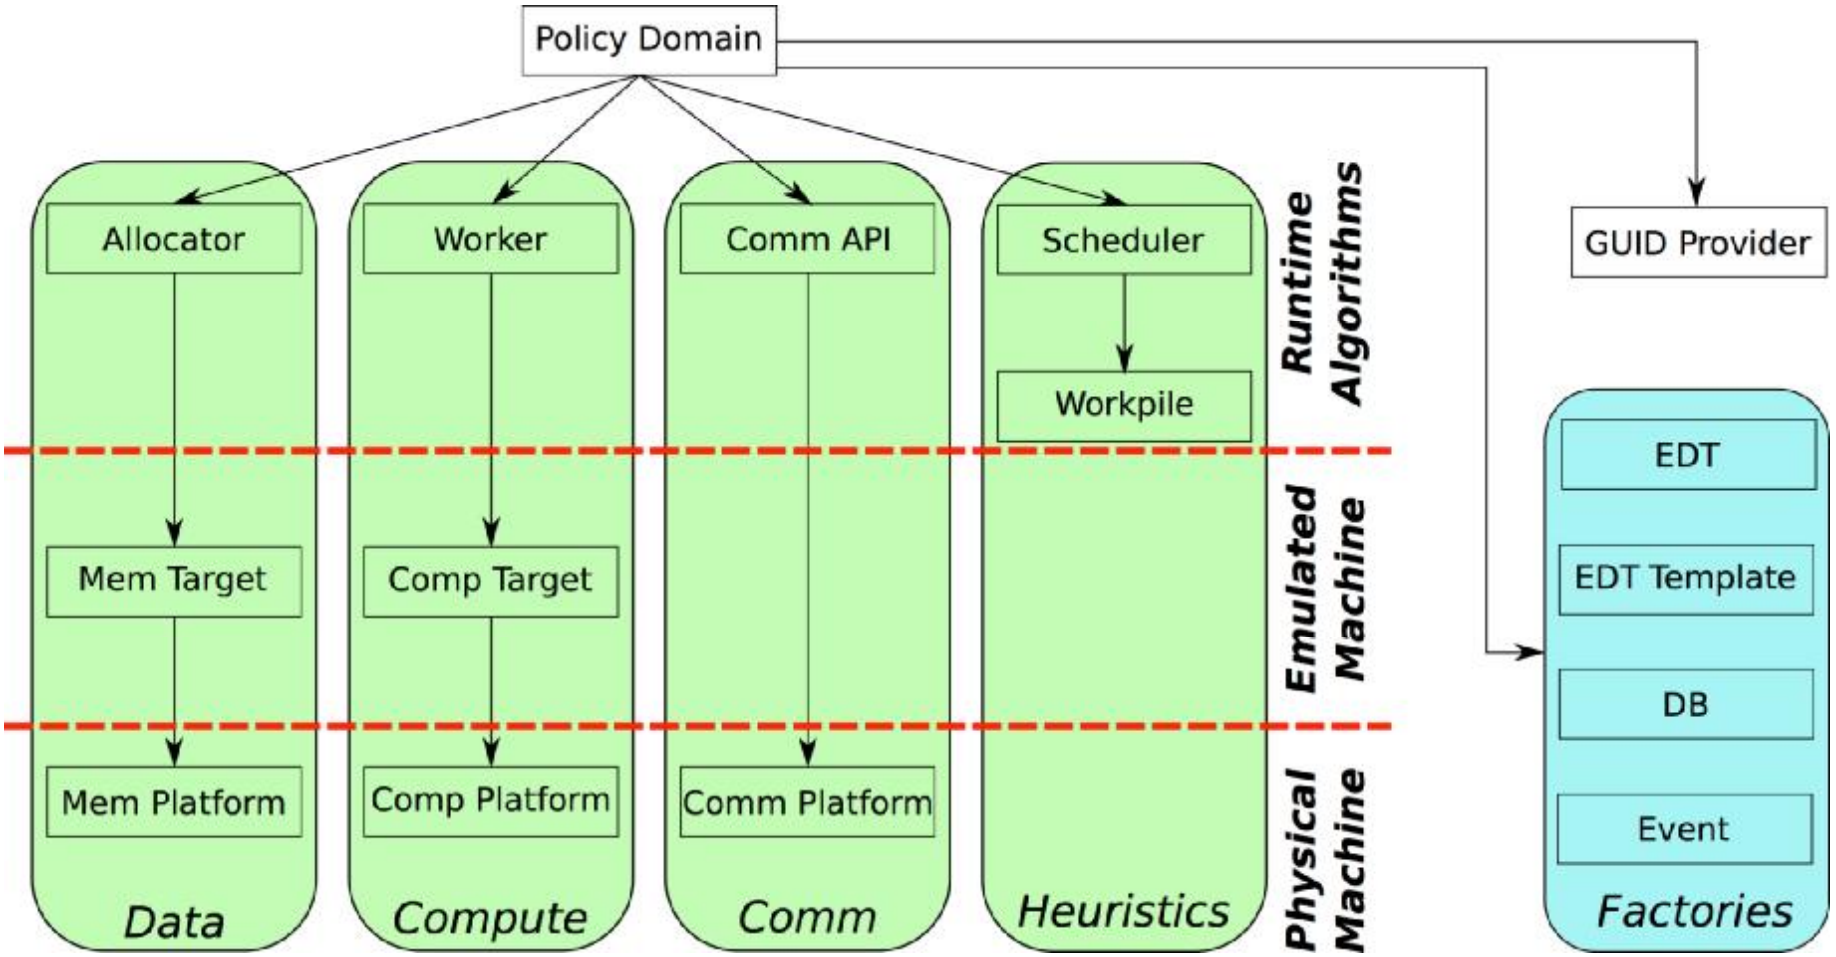
Policy Domain
GUID Provider
Allocator
Worker
Comm API
Scheduler
Workpile
Mem Target
Comp Target
Mem Platform
Comp Platform
Comm Platform
EDT
EDT Template
DB
Event
Data
Compute
Comm
Heuristics
Factories
Runtime Algorithms
Emulated Machine
Physical Machine

EDT Profiling using OCR

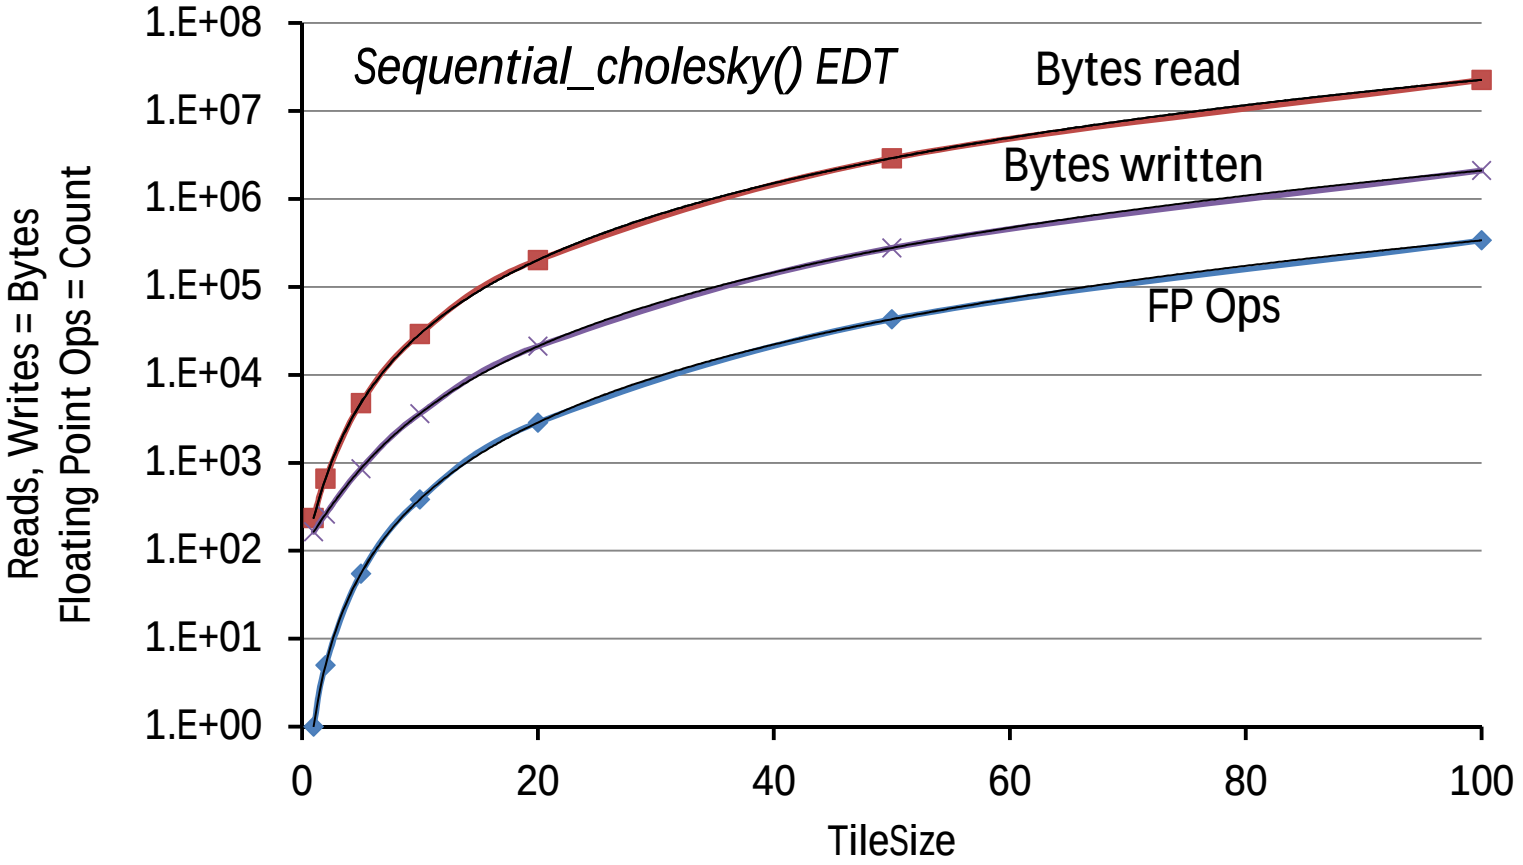

- Runtime statistics to evaluate benefits of the new approach

Introspection-based Projection

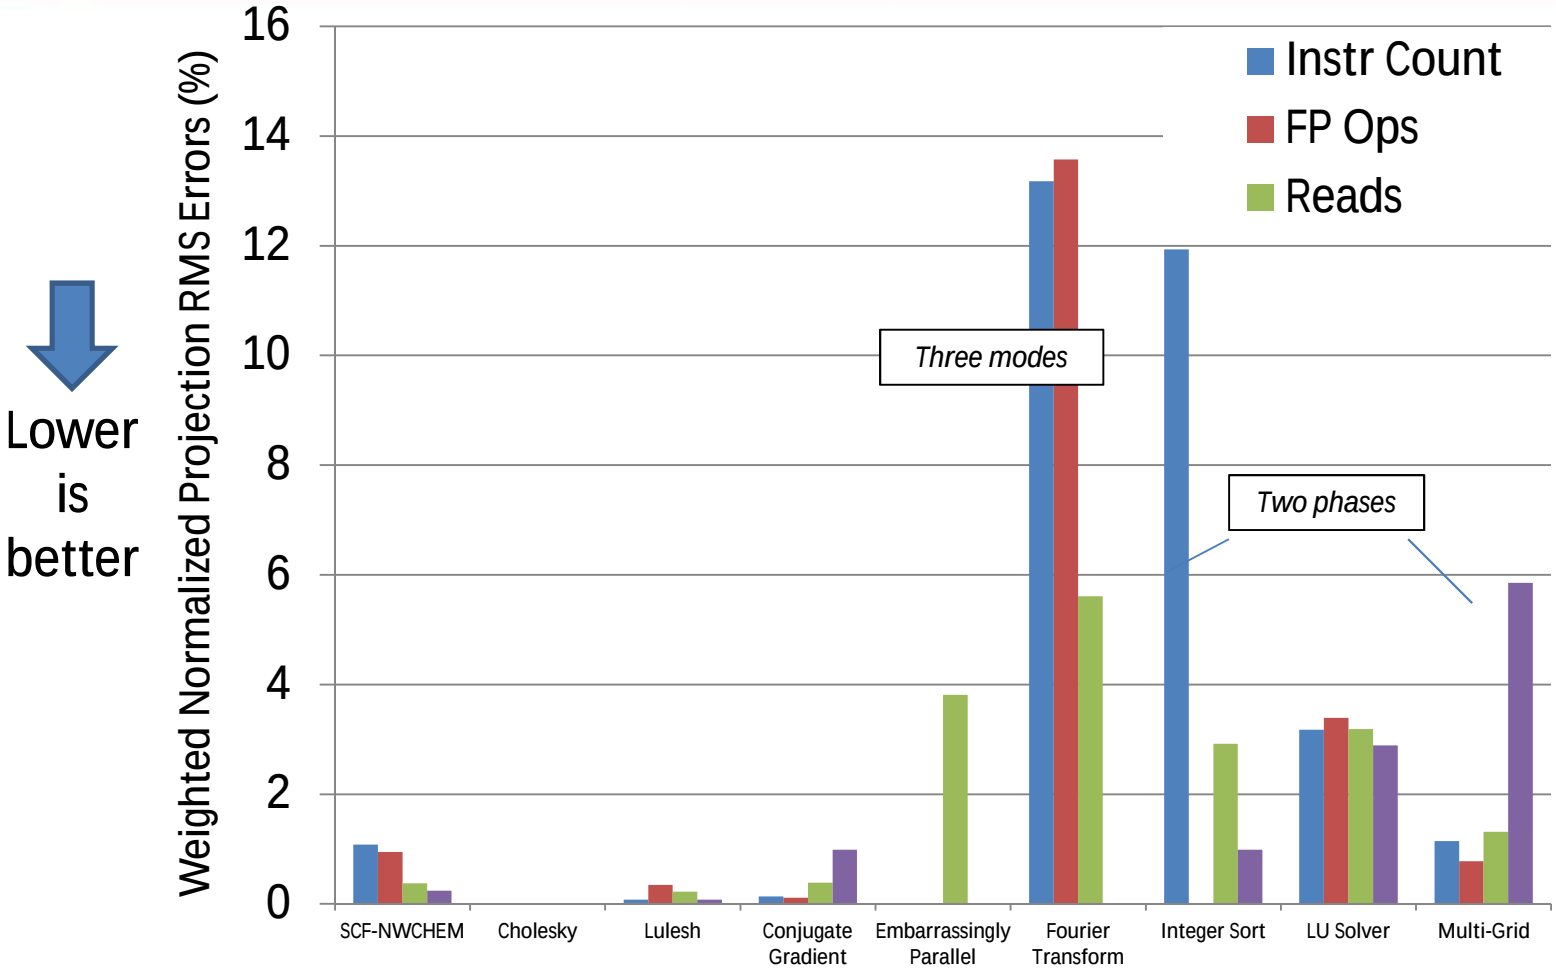

- Introspection based resource management looks promising

Acknowledgment: Chih-Chieh Yang, Adam Smith (NAS), Roger Golliver (LULESH), Jamie Arteaga (SCF)

OCR on 32-Thread SNB (Reservoir)

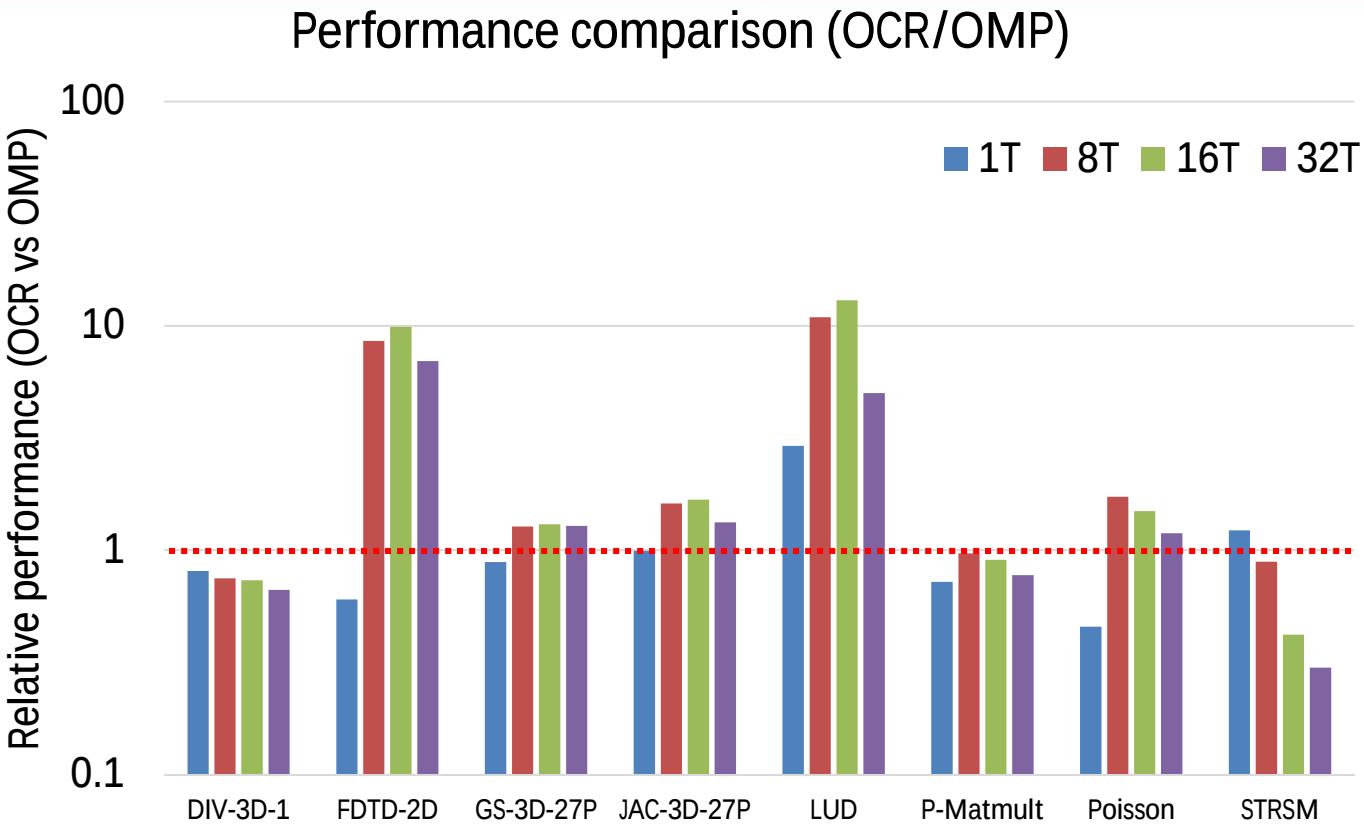

- Incomplete OCR implementation (no data blocks)
- Simple work-stealing scheduler
- Still… OCR is comparable or better than OMP in several benchmarks

(Shows a few benchmarks, for details please contact Reservoir Labs)

Thread Scaling with OCR (OR State University)

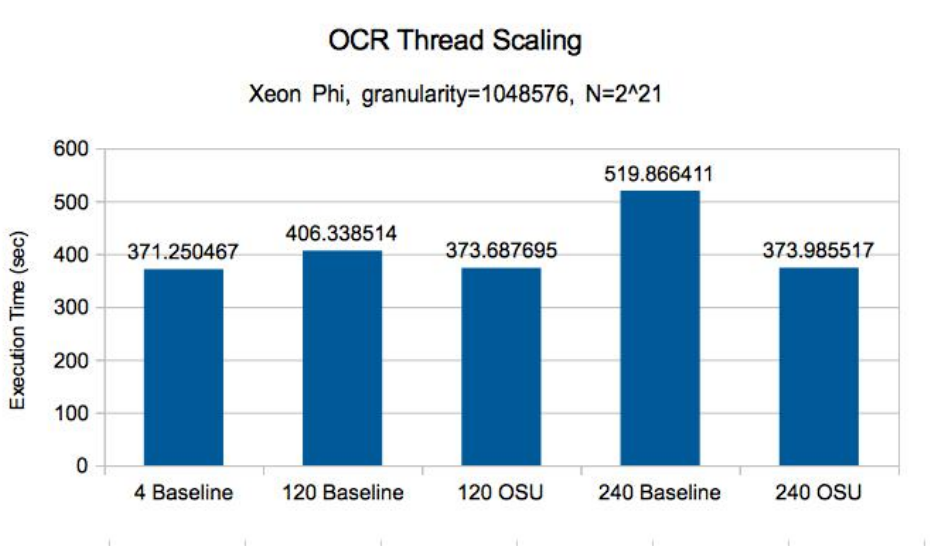

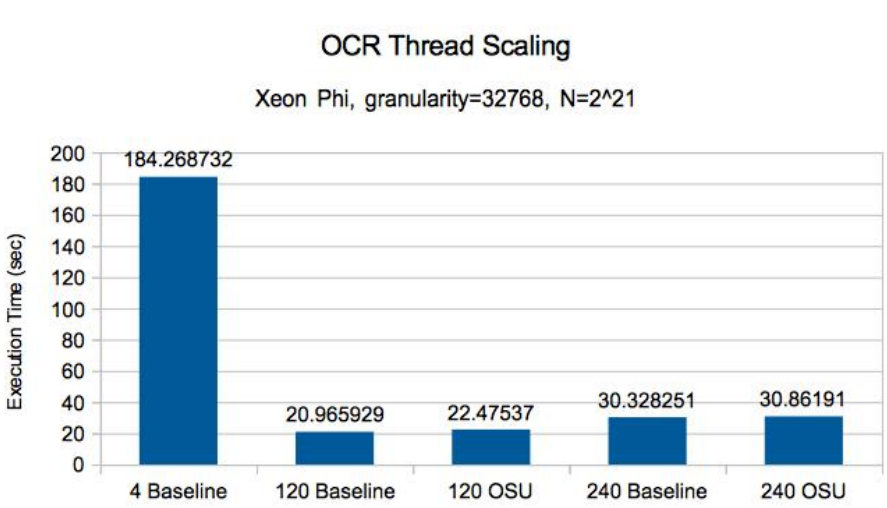

- Large data-blocks limit performance gains
- OSU scheduler (with back-off) increases performance due to fewer work-stealing attempts (less work present)

- Small data-blocks provide better speedup
- Runtime overhead seems to manifest itself

Applications Evaluated on TG Stack with OCR

- Hand-written OCR:
  - Cholesky (ETI)
  - CoMD (UCSD)
  - FFT (Oregon State)
  - HPCG (Oregon State and, separately, UCSD)
  - Lulesh 1.x (Roger Golliver)
  - SAR (Roger Golliver)
  - Stream (Oregon State)
- Written in CnC (Nick Vrvilo):
  - Smith-Waterman
  - Cholesky
- Written in HTA (UIUC):
  - Subset of NAS benchmarks (CG, FFT, Integer sort, LU decomposition, Multigrid solver) (UIUC)
- Several converted from RStream (Reservoir)

Summary

- TG SW Stack established
- Supports evolutionary & revolutionary approaches
- Open community runtime makes progress
- Results from applications look promising