

Traleika Glacier (X-Stack) – Intel Team

Goals and Objectives

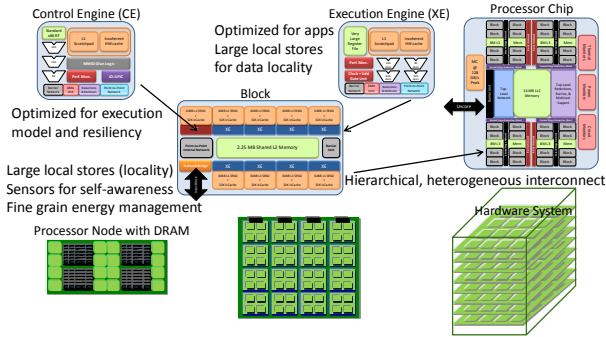
Goal:

Research and mature software technologies addressing major Exascale challenges and get ready to intercept by 2018-2020

Objectives:

Energy efficiency	SW components interoperate, harmonize, exploit HW features, and optimize the system for energy efficiency
Data locality	PGM system & system SW optimize to reduce data movement
Scalability	SW components scalable, portable to O(10 ⁹)—extreme parallelism
Programmability	New (Codelet) & legacy (MPI), with gentle slope for productivity
Execution model	Objective function based, dynamic, global system optimization
Self-awareness	Dynamically respond to changing conditions and demands
Resiliency	Asymptotically provide reliability of N-modular redundancy using HW/SW co-design; HW detection, SW correction

Straw-man System Architecture and Evaluation



Architecture embraces data-locality, and tapered BW

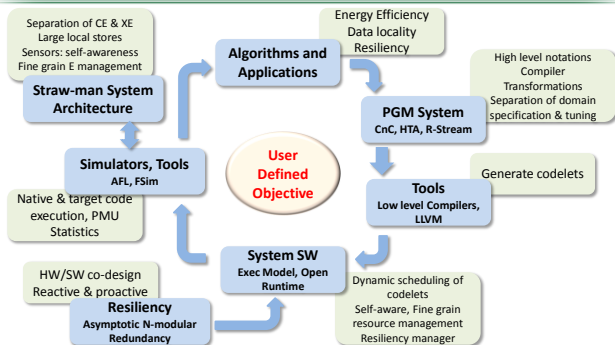
6

Runtime Research Agenda

- Locality aware scheduling—heuristics for locality/E-efficiency**
 - Extensions to standard Habanero-C runtime
- Adaptive boosting and idling of hardware**
 - Avoid energy expensive unsuccessful steals that perform no work
 - Turbo mode for a core executing serial code
 - Fine grain resource (including energy) management
- Dynamic data-block movement**
 - Co-locate codelets and data
 - Move codelets to data
- Introspection and dynamic optimization**
 - Performance counters, sensors provide real time information
 - Optimization of the system for user defined objective
 - (Go beyond energy proportional computing)

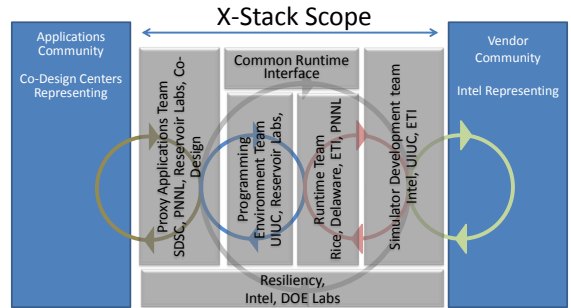
11

X-Stack Components Put Together



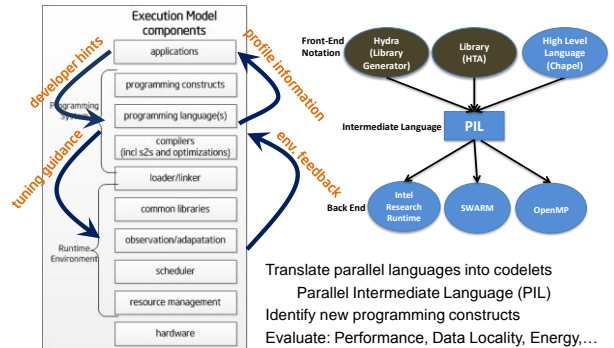
16

Scope of Traleika Glacier Project



4

Programming System Components



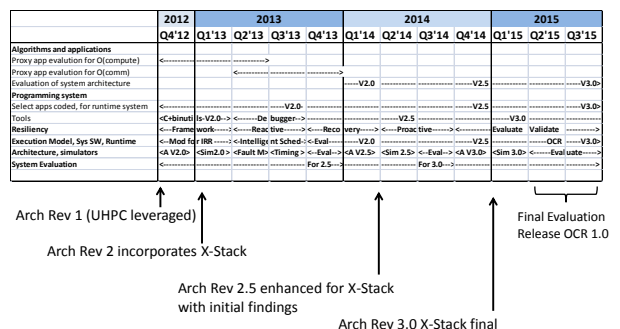
Translate parallel languages into codelets
Parallel Intermediate Language (PIL)
Identify new programming constructs
Evaluate: Performance, Data Locality, Energy...

Simulators and Tools, Leverage UHPC, and Go Beyond

Simulator	Pros	Cons	
AFL Application faking Library IRR API + Extensions	<ul style="list-style-type: none"> 64-bit Linux IRR implementation Near-native application code execution on host processor Rapid application development Epoch statistics as well as total statistics 	<ul style="list-style-type: none"> Does not model real architecture Does not model advanced ISA features Does not reflect expected timing of simulated system 	
FSim Functional Simulator	<ul style="list-style-type: none"> Every hardware unit modeled, exact memory and network hierarchies including messages Complete statistics and trace file >10 MIPS per core speed Massively parallel and distributed, limited only by machine pool 	<ul style="list-style-type: none"> Does not have a timing model Lower speed, highly detailed 	
Tool	Purpose	Advantage	Weakness
Power Estimator	<ul style="list-style-type: none"> Uses statistics/counters to make energy and area estimates for application behavior 	<ul style="list-style-type: none"> Scales from 45nm to 7nm projections Automatic analysis of outputs from FSim runs 	<ul style="list-style-type: none"> Only models dynamic power, uses circuit models for leakage Calibrated to existing commercial devices
Memory Analyzer	<ul style="list-style-type: none"> Detailed models for cache and/or scratchpad hierarchies, various levels & types of coherence Compares configurations 	<ul style="list-style-type: none"> Calls into Power Estimator Enables limited AFL trace power estimation 	<ul style="list-style-type: none"> Does not model instruction fetch/execution Limited to AFL-compatible memory traces at this time

12

3 Year Roadmap



18