U.S. DEPARTMENT OF ENERGY | Office of Science

X-Stack Kickoff Meeting

Sonia R. Sachs

09/18/2012

U.S. DEPARTMENT OF ENERGY | Office of Science

Outline

- Meeting Goals
  - Acknowledgements
- X-Stack Background
- X-Stack Portfolio
  - The X-Stack puzzle
  - Strawman vision of X-Stack

Meeting Goals

- Review the X-Stack portfolio development process
- X-Stack Puzzle
  - Understand the "pieces of the puzzle"
  - Blending, reshaping, extending boundaries, sharpening/smoothing edges
- Start process of completing the X-Stack vision:
  - Detailed architecture and interfaces
  - Mapping of tasks from projects to components in the vision

U.S. DEPARTMENT OF ENERGY | Office of Science

Acknowledgements

- Intel for hosting our meeting
  - Shekhar and Wilf
- Intel support staff for meeting logistics
  - Jim and Barbara
- Committee for Sessions
  - Kathy, Shekhar, Dan, Saman, and Ron
- Tina Macaluso: lead scribe
- X-Stack PIs for all meeting materials
- Guests of X-Stack portfolio

U.S. DEPARTMENT OF ENERGY | Office of Science

X-Stack Background: Programming Challenges Workshop

Understand and Prioritize Challenges for Programming Exascale systems:

- Bulk-synchronous to asynchronous computing.
- Expressing and managing up to a billion separate threads
- Expressing and managing hierarchical locality and data movement
- Dealing with heterogeneity across the system
- Dealing with demands for adaptive, dynamic scheduling of work and resources
- Dealing with energy and resilience constraints

Parallelism

Data Movement

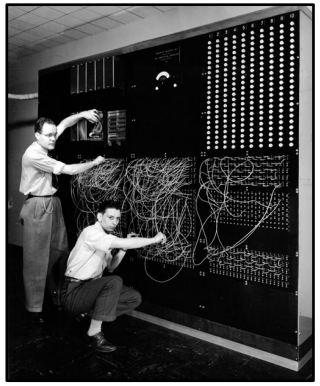

Programmability

Resiliency

Bill Harrod: multiple presentations

X-Stack Background: Programming Challenges Workshop

Understand and prioritize new approaches for programming Exascale systems in:

- Compilers
- Programming Models
- Programming Languages and environments
- Runtime systems

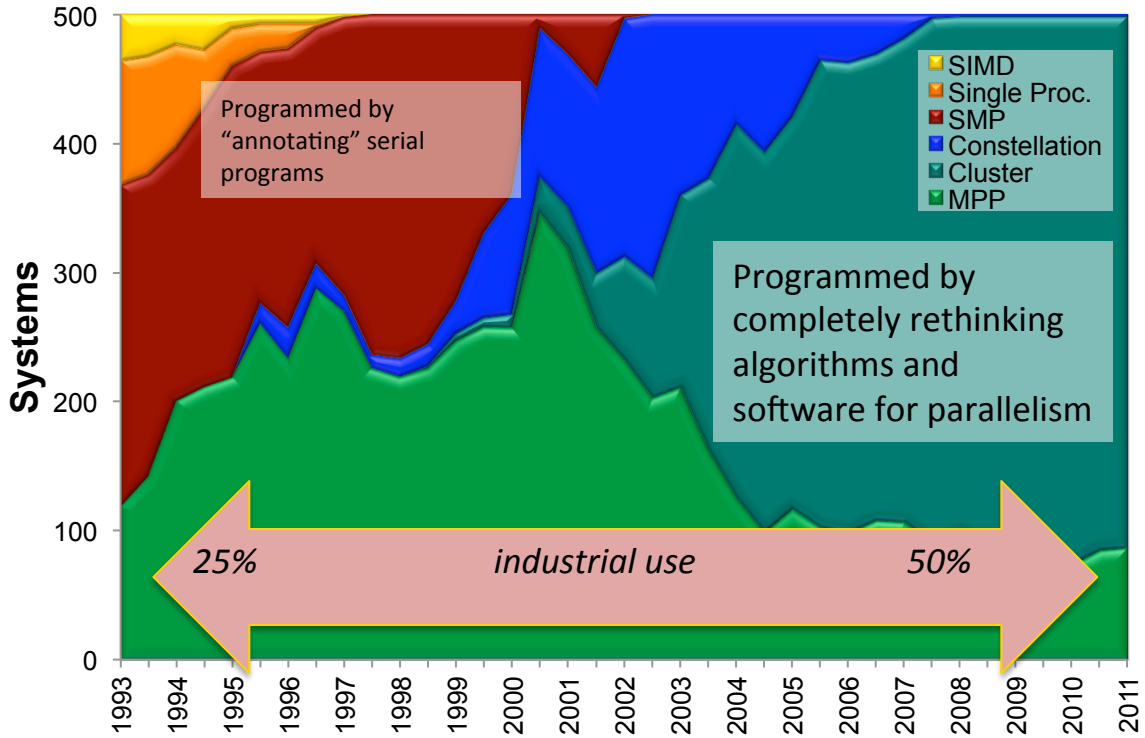

Kathy Yelick: "Software Stack and Co-Design"

U.S. DEPARTMENT OF ENERGY | Office of Science

X-Stack Background: Conclusions of Programming Challenges Workshop

- **Programming stack**
  - **abstract specification** of computation: DSLs and embedded DSLs for expressing the mathematics of systems under study
  - Selection of algorithm implementations*
  - Map computations and data to machine representation and select communications mechanisms*
  - Select computational library routines*
  - Optimization via compiler transformations and runtime mechanisms
  - **runtime optimized code**

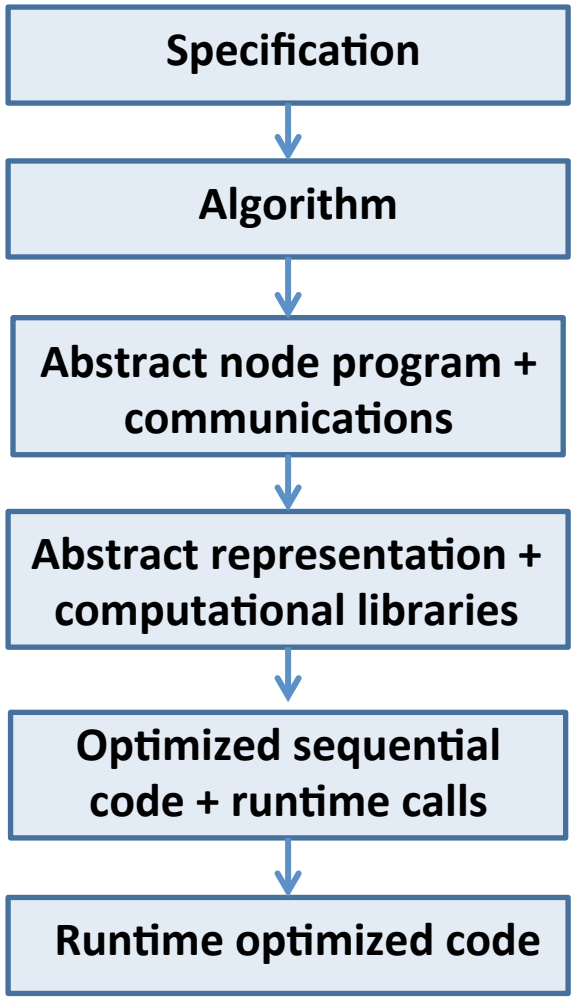

* **Significant automation**

X-Stack Background: Conclusions of Programming Challenges Workshop

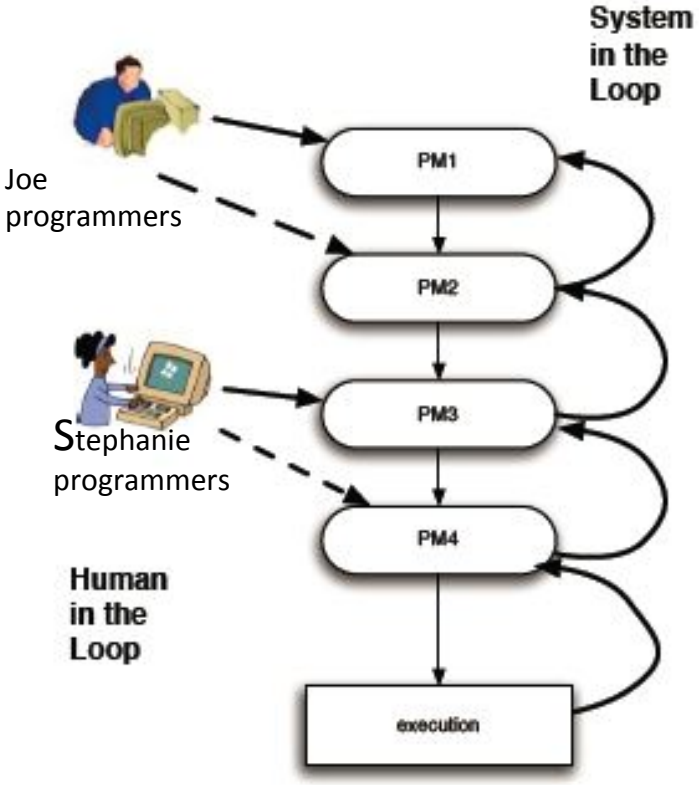

Multi-resolution Models

- Mappings from high-level representations to low level ones are **semantics and performance preserving**
  - Mostly automated
- Reverse mappings (system in the loop)
  - for debugging and tuning
- Multiple levels of different programmers (human in the loop)
  - Mostly high level programming (Joe programmers)
  - Occasionally low level interference (Stephanie programmers)

X-Stack FOA: Programming Challenges, Runtime Systems, and Tools

- Focus Areas
  - Programming models, languages, runtime systems, tools, and related technologies
  - New energy-efficient and resilient programming techniques that are portable across multiple future machine generations
- What was encouraged:
  - **Evolutionary and Revolutionary** solutions demonstrating support for:
    - billions of threads,
    - much more constrained memory systems,
    - heterogeneous cores,
    - deep memory hierarchies,
    - asynchronous data movement and irregular applications,
    - performance portability,
    - active energy and power management mechanisms
  - Evolutionary and Revolutionary solutions that **converge to a uniform treatment of parallelism intra- and inter-node.**

U.S. DEPARTMENT OF ENERGY | Office of Science

X-Stack FOA

- Proposals were expected to:
  - Articulate complete solutions integrating multiple components of the system software stack and addressing Exascale challenges:
    - Scalability
    - Programmability
    - Performance Portability
    - Resilience
    - Energy Efficiency
- Proposal were required to have:
  - Description of plans for **developing prototypes** of the proposed solution;
  - Description of the **proposed path to integration and/or interoperation** with existing programming environments;
  - **Evaluation plan** using compact applications, mini-applications

U.S. DEPARTMENT OF ENERGY | Office of Science

X-Stack Timeline

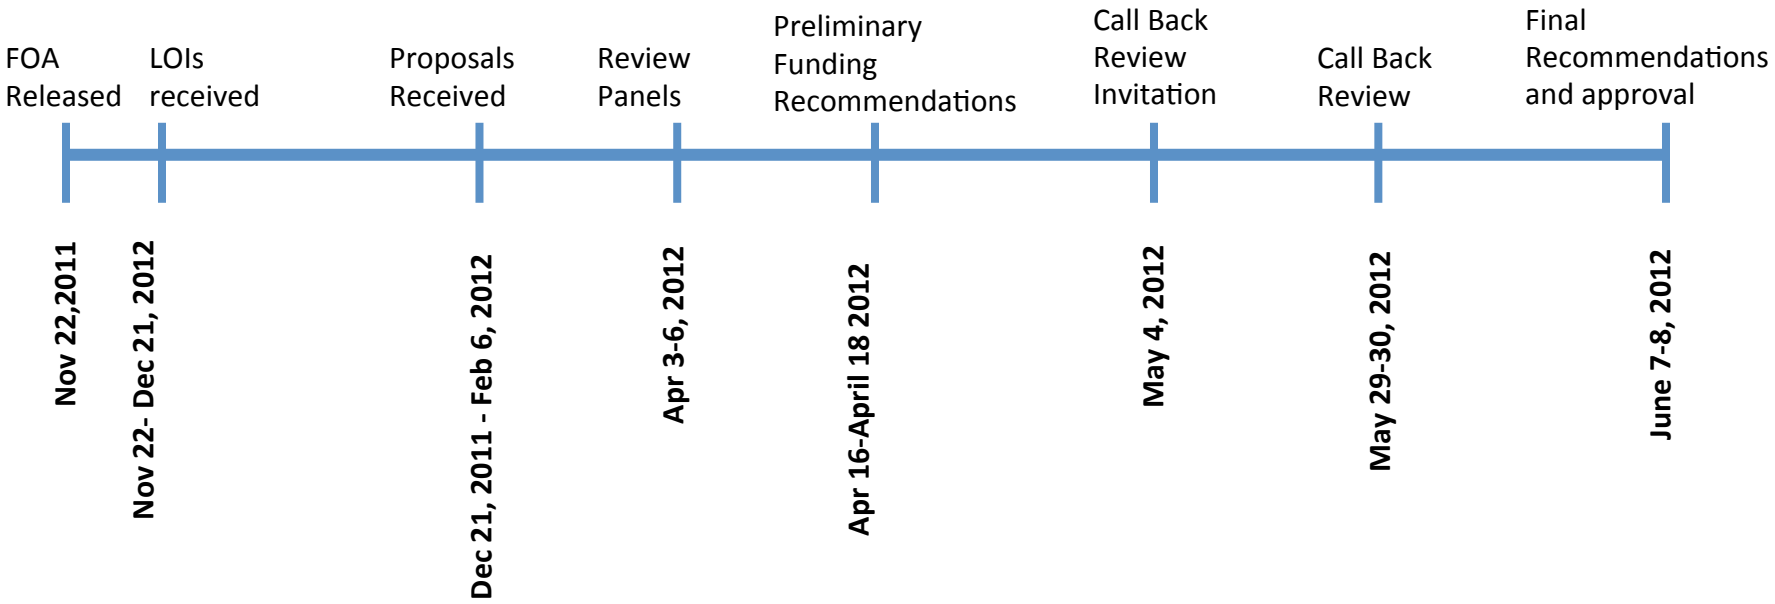

Available funding up to $15,000,000 per year for three years. Anticipated project funding between $500,000 - $4,000,000 per project.

ASCR received full proposals: 68

- 24 Lab-led proposals
- 36 University-led proposals
- 7 Industry-led proposals
- 1 Declined without review

U.S. DEPARTMENT OF ENERGY | Office of Science

The X-Stack Portfolio

DAX (ETI):

SLEEC: Purdue

D-TEC: LLNL and MIT

Traleika Glacier:Intel

GVR: U. Chicago

DEGAS: LBNL

Autotunig: U. Utah

CORVETTE: UC Berkeley

XPRESS: Sandia

Project Title	Lead PI(s)	Description
Traleika Glacier	Shekhar Borkar	Simulation infrastructure. Compiler optimization, execution models, and runtime environments.
DEGAS	Kathy Yelick	Hierarchical programming models, language design, compilers, communications layer, adaptive runtime, resilience.
D-TEC: DSL Technology for Exascale Computing	Dan Quinlan and Saman Amarasinghe	Complete solution for X-Stack: DSL, compilers, abstract machine model, refinement and transformation framework, adaptive runtime systems.
XPRESS	Ron Brightwell	Runtime system implementing Parallex, co-designed with an OS. Framework to translate MPI and OpenMP legacy codes.
DAX: Dynamically Adaptive X-Stack	Rishi Khan	Programming Models, Compilers and Runtime Systems for Dynamic Adaptive Event-Driven Execution Models.
Autotuning for Exascale	Mary Hall	A unified autotuning framework that seamlessly integrates programmer-directed and compiler-directed autotuning.
Global View Resilience	Andrew Chien	Cross layer error management architecture.
CORVETTE	Koushik Sen	Automated bug finding methods to eliminate non- determinism in program execution and to make concurrency bugs and floating point behavior reproducible.
SLEEC	Milind Kulkarni	Annotation language. Function semantics exposed in DSL libraries. Design and development of a semantics-aware, extensible optimizing compiler that treats compilation as an optimization problem.

U.S. DEPARTMENT OF ENERGY | Office of Science

X-Stack Software Vision

Aligned with the Exascale Research Initiative

- ECI Goals: Deploy exascale computers:
  - 500 to 1,000 more **performance** than today's HPC systems
  - Under **20MW** Power
  - Highly **programmable**
- ECI Strategy:
  - Conduct critical R&D efforts.
  - Develop exascale software stacks.
  - Fund computer technology vendors
  - Fund the design and development of exascale computer systems.
  - **Joint effort with NNSA.**
  - Collaboration with other government agencies and other countries.

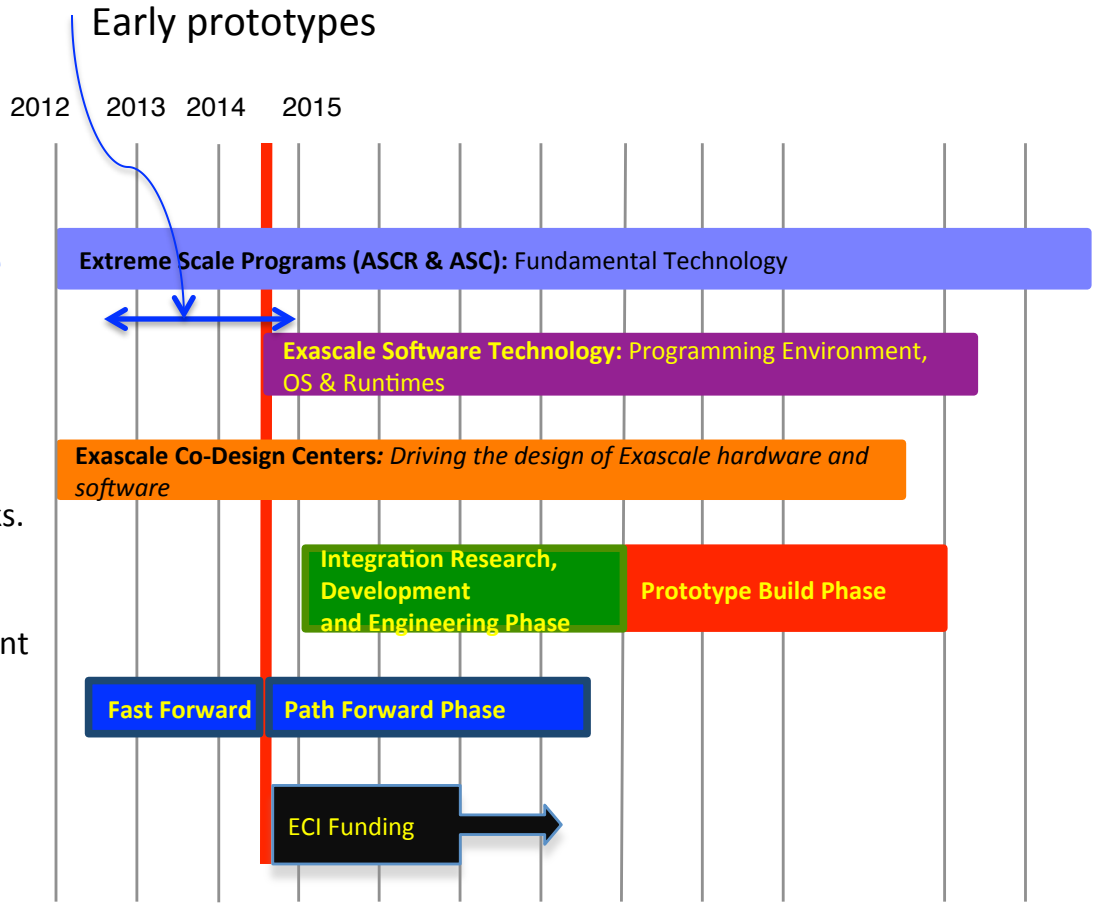

Strawman X-Stack Vision

Energy Efficiency, Resilience, Programmability, Scalability, Performance Portability, Interoperability

Evaluation Plans

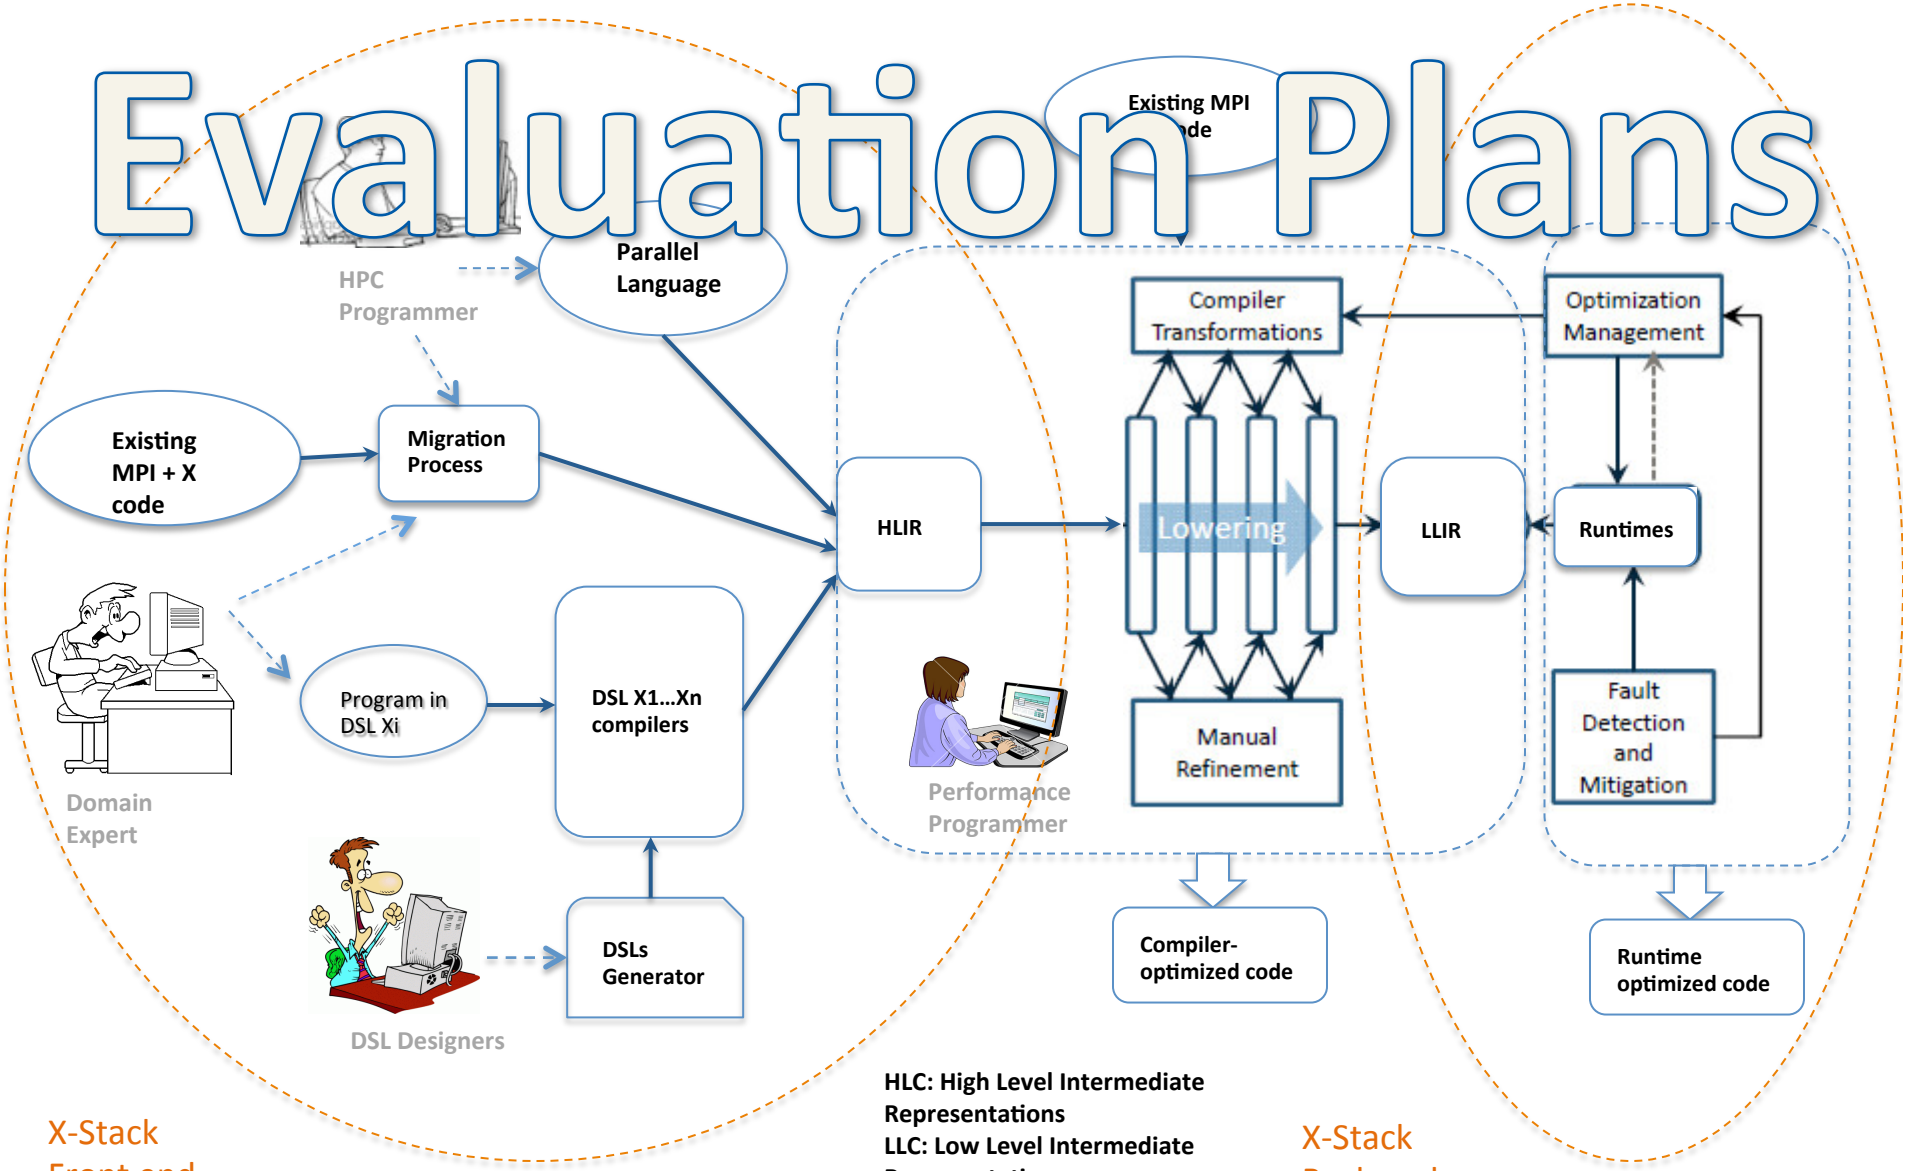

U.S. DEPARTMENT OF ENERGY | Office of Science

Agenda – First Day

7:00am – 8:00am	**Full Breakfast Registration**
8:00am – 8:30am	Opening Remarks (Bill, Sonia)
8:30am – 9:15am	Traleika Glacier (Intel team)
9:15am – 10:00am	D-TEC (LLNL)
10:00am – 10:30am	Break & posters
10:30am – 11:15am	DEGAS (LBNL)
11:15pm – 12:00pm	XPRESS (Sandia)
12:00pm – 1:00pm	Lunch & posters
1:00pm – 1:30pm	Programming Models, Compilers and Runtime Systems for Dynamic Adaptive Event-Driven Execution Models (ETI)

1:30pm – 2:00pm	**Exploiting Global View for Resilience (GVR) – Andrew Chien**
2:00pm – 2:30pm	Program Correctness, Verification and Testing for Exascale (Corvette)
2:30pm – 3:30pm	Break & posters
3:30am – 4:00pm	Autotuning for Exascale: Self-Tuning Software to Manage Heterogeneity in Algorithms, Processors and Memory Systems (Utah team)
4:00pm – 4:30am	SLEEC: Semantics-rich Libraries for Effective Exascale Computation (Purdue team)
4:30pm – 6:00pm	Discussions & posters

Agenda – Second Day

Energy Efficiency, Resilience, Programmability, Scalability, Performance Portability, Interoperability

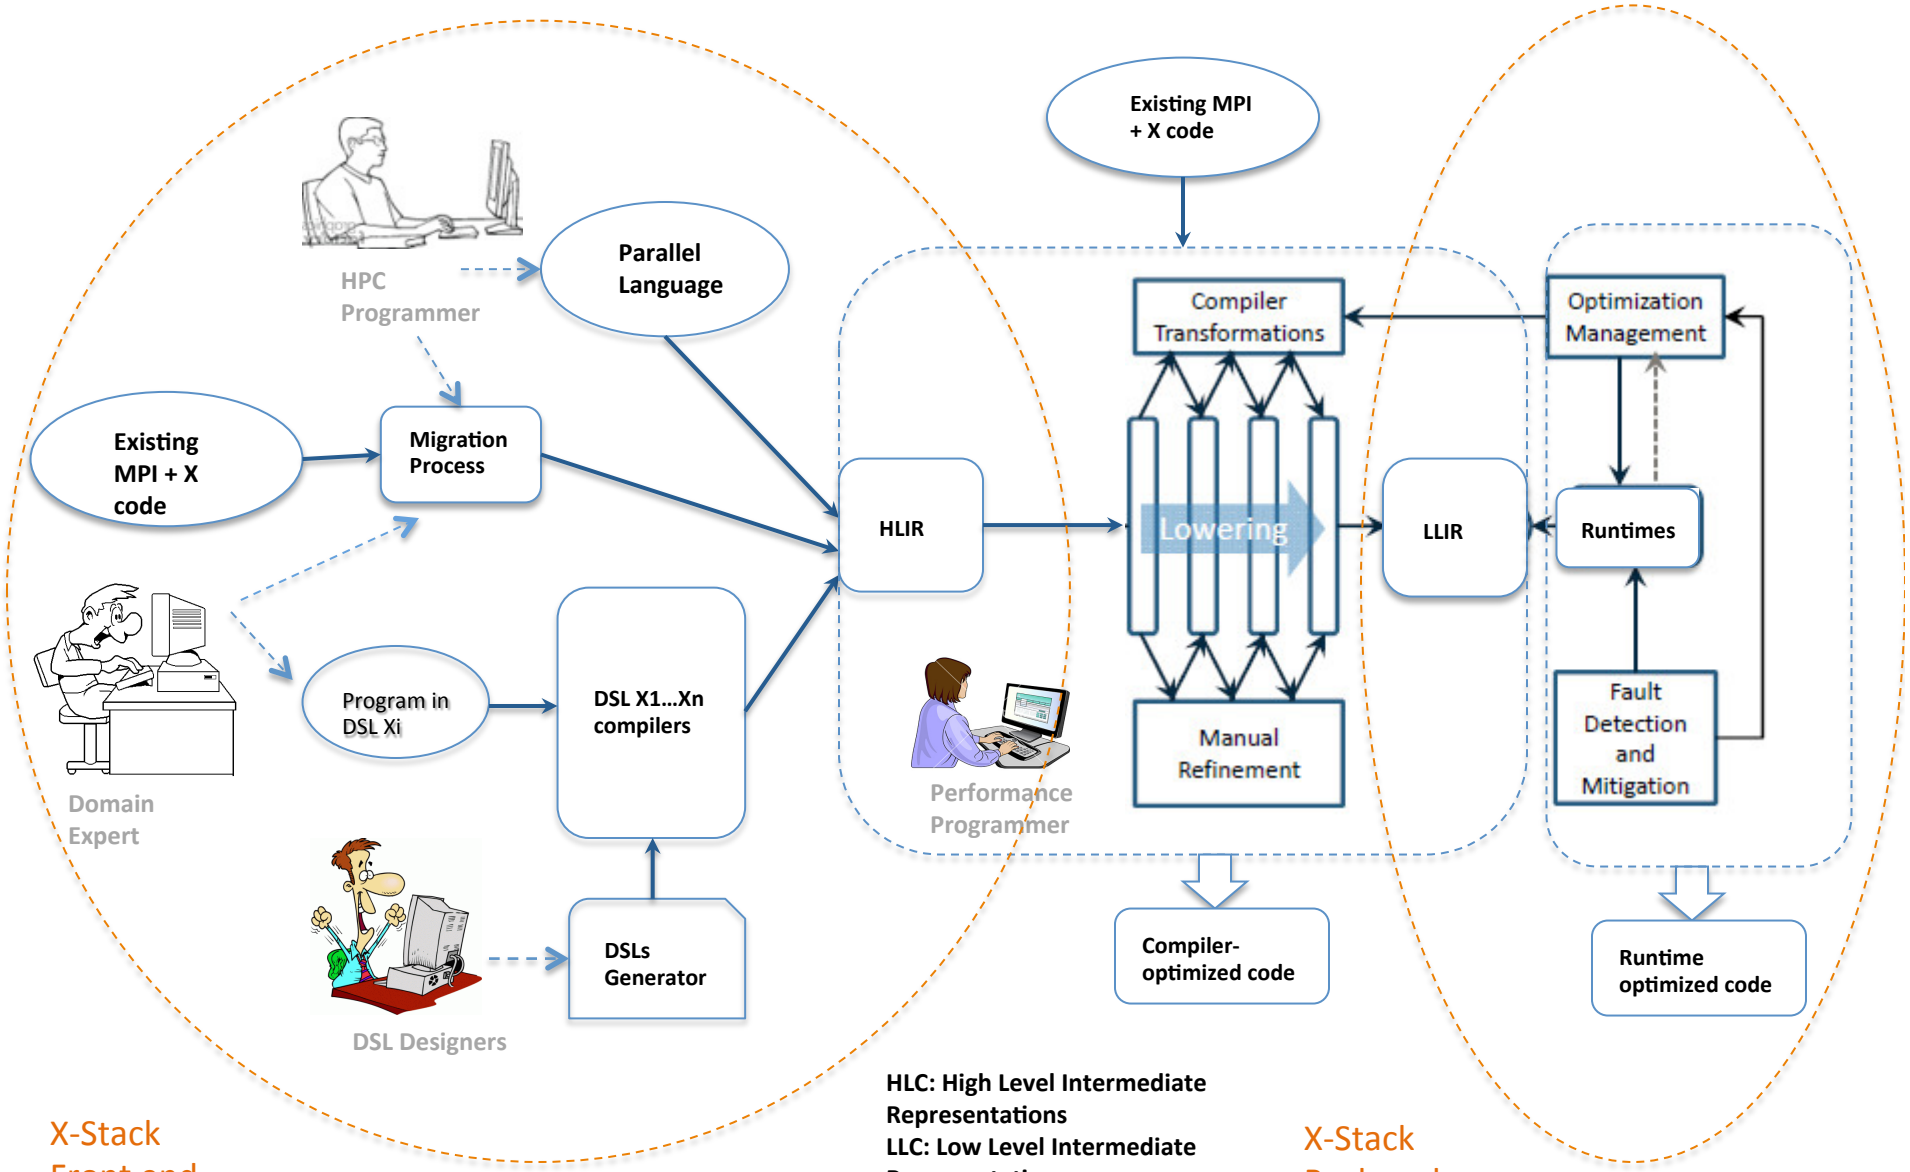

U.S. DEPARTMENT OF ENERGY | Office of Science

Agenda – Second Day

7:00am – 8:00am	**Full Breakfast Registration**
8:00am – 8:10am	X-Stack postdocs pool (Sonia)
8:10am – 9:10am	Co-design centers: explaining the proxy apps (15 minutes for each center)
9:10am – 10:45am	**Parallel session I**: X-Stack front end **Parallel session II**: X-Stack back end
10:45am – 11:00am	Break
11:00am – 12:00 noon	**Parallel Session I**: continue **Parallel Session II**: continue
12:00am – 12:30am	Reporting on parallel sessions.
12:30pm – 1:30pm	Lunch & posters

1:30pm – 2:30pm	**Panel:** High Level Representations: Programming Models, DSLs, parallel languages, MPI+X, resilience
2:30pm – 3:30pm	**Panel:** Runtime systems (ARTS, OCR, HPX, X10-SEEC, SWARM), resilience
3:30pm – 4:00pm	Break & posters
4:00pm – 5:00pm	**Panel:** Low Level Representations: synthesis, refinements, transformations, resilience
5:00pm – 5:30pm	Discussions and path forward