



U.S. DEPARTMENT OF
ENERGY


Office of
Science



Parallelism




Data Movement




Programmability



Resiliency



Energy Efficiency




Deep Memory
Hierarchies

X-Stack and OS/R PI Meeting

Sonia R. Sachs

December 7-8, 2015



Parallelism

Data Movement

Programmability

Resiliency

Energy Efficiency

Deep Memory Hierarchies

Meeting Goals

- X-Stack and OS/R programs overview
- Introduction to the Exascale Computing Project
- Present the X-Stack and the OS/R Programs



X-Stack Program

- Community input collected in workshops held in 2011:
 - ASCR Exascale Programming Challenges Workshop:
 - <http://science.energy.gov/ascr/research/computer-science/programming-challenges-workshop/>
 - ASCR Exascale Tools Workshop:
 - <http://science.energy.gov/ascr/research/computer-science/exascale-tools-workshop/>
- FOA topics: Programming models, languages, compilers, runtime systems, and tools that address identified Exascale challenges
- FOA released November 22, 2011 with Deadline of February 6, 2012.
- 67 full proposals received and reviewed: 24 Lab proposals, 36 University proposals, 7 industry proposals
- Projects were funded for 3 years, starting on September 2012, until September 2015.
- We renewed many of the projects for one year until September, 2016, in order to enable the teams to mature their technologies and research prototypes.



U.S. DEPARTMENT OF
ENERGY

Office of
Science



<https://xstackwiki.modelado.org/>

X-Stack Portfolio (original)



DEGAS (Kathy Yelick)

Hierarchical and resilient programming models, compilers and runtime support.



Traleika (Shekhar Borkar)

Exascale programming system, execution model and runtime, applications, and architecture explorations, with open and shared simulation infrastructure.



D-TEC (Dan Quinlan and Saman Amarasinghe)

Complete software stack solution, from DSLs to compilers to optimized runtime systems.



XPRESS (Ron Brightwell)

Software architecture and interfaces that exploit the ParalleX execution model, prototyping several of its key components.



DynAX (Rishi Khan)

Novel programming models, dynamic adaptive execution models and runtime systems.



X-Tune (Mary Hall)

Unified autotuning framework that integrates programmer-directed and compiler-directed autotuning.

The University of Chicago



GVR (Andrew Chien)

Global view data model for architecture support for resilience.



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)

Semantics-aware, extensible optimizing compiler that treats compilation as an optimization problem.



X-Stack Portfolio changes

- Started September 2013, until September 2016



PIPER (Martin Schultz)

Tools for debugging and analysis of performance, power, and energy.



Vancouver 2 (Jeffrey Vetter)

Heterogeneous computing:
abstractions, runtime systems, tools
and low-level libraries

- Started in June 2014, until September 2016



Los Alamos National Laboratory

Delivering science and technology to protect our nation and promote world stability

ARES (Pat McCormick and Jeffrey Vetter)

High-level intermediate
representations (HLIR) and
optimizations to select architectures



X-Stack Portfolio (fourth year)



DEGAS (Kathy Yelick)

Hierarchical and resilient programming models, compilers and runtime support.



Traleika (Shekhar Borkar)

Exascale programming system, execution model and runtime, applications, and architecture explorations, with open and shared simulation infrastructure.



D-TEC (Dan Quinlan and Saman Amarasinghe)

Complete software stack solution, from DSLs to optimized runtime systems code.



XPRESS (Ron Brightwell)

Software architecture and interfaces that exploit the ParalleX execution model, prototyping several of its key components.



X-Tune (Mary Hall)

Unified autotuning framework that integrates programmer-directed and compiler-directed autotuning.



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)

Semantics-aware, extensible optimizing compiler that treats compilation as an optimization problem.



OS/R Program

- Challenges and strategies related to OS/R for Exascale platforms were identified in the ASCR Exascale Operating System and Runtime Software workshop October 4-5, 2012:

<https://collab.mcs.anl.gov/display/exasr/Workshop%20Overview>

- ASCR/NNSA Technical Council on OS/R met March 2012 – October 2012. Meeting results and OS/R workshop insights were captured in a report:
<http://science.energy.gov/~media/ascr/pdf/research/cs/Exascale%20Workshop/ExaOSR-Report-Final.pdf>
- Lab Solicitation topics: Runtime and operating systems that address identified Exascale challenges
- Lab solicitation released January 2, 2013 with revised deadline of February 25, 2013.
- Two full proposals received and reviewed
- Two projects were funded for 3 years, starting September 2013, until September 2016.



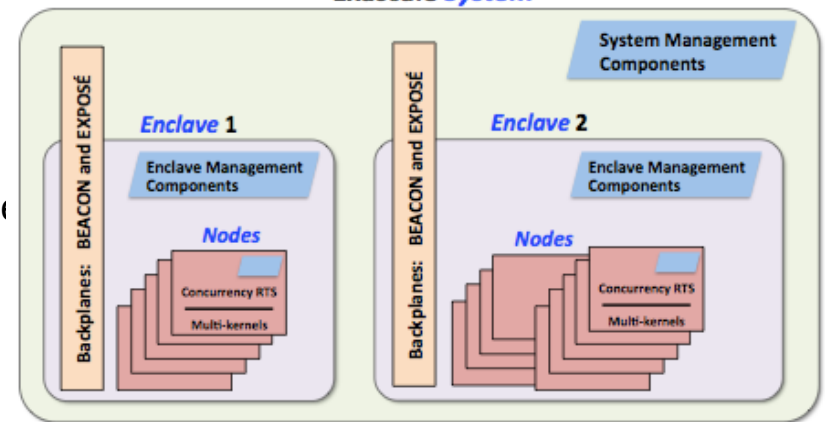
OS/R Portfolio (original)



ARGO (Pete Beckman)


Node OS/R providing extremely lightweight multithreading, heterogeneous hardware resource management, and dynamic node management services. Complete node, enclave, and global OS/R prototype.

Exascale System



HOBBS (Ron Brightwell)

Extensions of previous architectures and designs, providing complete node, enclave, and global OS/R prototype. Full system virtualization supports almost any RTS. Supports composition of applications developed for different programming systems.





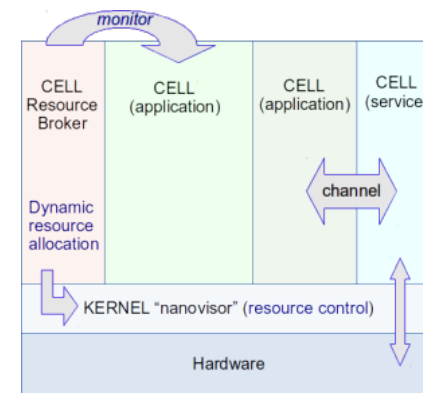
OS/R Portfolio change

- Started September 2013, until September 2016



X-ARCC (Steven Hofmeyr)

New approaches to Operating System design for exascale using *Adaptive Resource-Centric Computing (ARCC)*, which combines dynamic resource with Quality-of-Service (QoS). Based on the UCB Tessellation project.





U.S. DEPARTMENT OF
ENERGY

Office of
Science

Backup



X-Stack FOA: what we asked for

- “Sought are complete solutions that will address multiple components of the system software stack and that will have the following characteristics:
 - Scalability: enable applications to strongly scale to Exascale levels of parallelism;
 - Programmability: clearly reduce the burden we are placing on high performance programmers;
 - Performance Portability: eliminate or significantly minimize requirements for porting to future platforms;
 - Resilience: properly manage fault detection and recovery at all components of the software stack; and
 - Energy Efficiency: maximally exploit dynamic energy saving opportunities, leveraging the tradeoffs between energy efficiency, resilience, and performance.”
- “It is important to demonstrate the viability of such solutions in a broad high performance programming context by showing how the proposed solution:
 - interoperates with existing programming environments based on the MPI+X model, so that a smooth migration path is enabled, and/or
 - enables the automatic transformation of applications (possibly with users in the loop) from the “old” programming environment to the he “new” one, such that the transformations are semantics and performance preserving.
- Complete FOA:
 - http://science.energy.gov/~media/grants/pdf/foas/2012/SC_FOA_0000619.pdf



Projects Coordinating with X-Stack Program

Co-Design Centers

CoDEX

DMD

Blackcomb 2

CAL



<https://xstackwiki.modelado.org/>

Exascale MPI

ARGO

HOBBS

X-ARCC

Fast Forward (I and II)

Design Forward



OS/R FOA: what we asked for

- “...seeking collaborative proposals that address these key challenges [listed below], leading to platform-neutral prototypes of a complete exascale operating system and runtime software (OS/R).”
- Identified challenges:
 - Energy efficient hardware and software design constrained to a power envelope of 20 Megawatts(MW);
 - Managing unprecedented parallelism, especially at the node level of exascale systems;
 - Managing deep memory hierarchies and multi-stage storage systems under severe energy constraints; and
 - Resilience in the presence of predicted high rates of hard and soft faults.
- “The proposed research is expected to leverage the OS/R reference architecture [described in section 4 of <http://science.energy.gov/~media/ascr/pdf/research/cs/Exascale%20Workshop/ExaOSR-Report-Final.pdf>]. The use of alternative architectures is allowed, but proposals must contain extensive justifications for the proposed deviations. “
- Complete solicitation:
http://science.energy.gov/~media/grants/pdf/lab-announcements/2013/LAB_13-02.pdf