



DSL Technology for Exascale Computing (D-TEC)

Lead PI and DOE lab: Daniel J. Quinlan Lawrence Livermore National Laboratory

Co-PIs and Institutions

Saman Amarasinghe, Armando Solar-Lezama, Adam Chlipala, Srinivas Devadas,
 Una-May O'Reilly, Nir Shavit, Youssef Marzouk @ MIT
 John Mellor-Crummey & Vivek Sarkar @ Rice University
 Vijay Saraswat & David Grove @ IBM Watson

P. Sadayappan & Atanas Rountev @ Ohio State University
 Ras Bodik @ University of California at Berkeley
 Craig Rasmussen @ University of Oregon
 Phil Colella @ Lawrence Berkeley National Laboratory
 Scott Baden @ University of California at San Diego


INTRODUCTION

- Domain Specific Languages define high level abstractions that make it efficient for the development of application code.
- High-level abstractions improve productivity
- Enables domain-specific performance and energy optimizations
- However, it is very difficult to efficiently leverage DSLs due to lengthy design and development of the corresponding software stack support, including languages, compilers, runtime and tools.

GOAL

- Making DSLs effective for exascale
- Support both embedded and general (syntax extended) DSLs
 - Address all layers of the exascale software stack: languages (DSLs), compilers, abstract machine, runtime, and tools
 - Address multiple exascale challenges: scalability, programmability, performance portability, resilience, energy efficiency, correctness, and heterogeneity
 - Include interoperability with MPI+X through translation to low level code
 - Provide a migration strategy for legacy code through source-to-source technology


The D-TEC approach addresses the full Exascale workflow



APPROACH


- Languages (C1):
 - Discovery of domain specific abstractions through collaborative efforts from both computer scientists and application experts
- Rosebud: define DSLs specification and plug-in
- Refinement/Transformation(C2):
 - Series of manual refinements steps (code rewrites) define the transformations
 - Equivalence checking between steps to verify correctness
 - Machine learning to drive optimizations
- Compiler (C3):
 - Leverage a source-to-source compiler infrastructure, ROSE, to create an DSL framework to support code rewriting, analysis and optimizations
- Parameterized Abstract Machine (C4): extraction of machine characteristics
- Runtime System (C5): leverages X10 and extends it with SEEC support
- Tools (C6): source-to-source migration tools and tools for mappings between DSL layers to support future tools

Management Organization Plan and Collaboration Paths with Advisory Board and Outside Community



Summary

DSLs are expected to be a transformational technology to bridge the gap between diverse algorithms/applications and complex exascale machines. We are working on a comprehensive DSL framework to facilitate the definition, implementation and optimizations of DSLs for addressing multiple challenges arisen from the future exascale computing. *More info at <http://www.dtec-xstack.org>.*



This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.