

# SLEEC: Semantics-rich Libraries for Effective Exascale Computation

Milind Kulkarni, Arun Prakash, Vijay Pai and Sam Midkiff (Purdue University) and Michael Parks (Sandia National Labs)

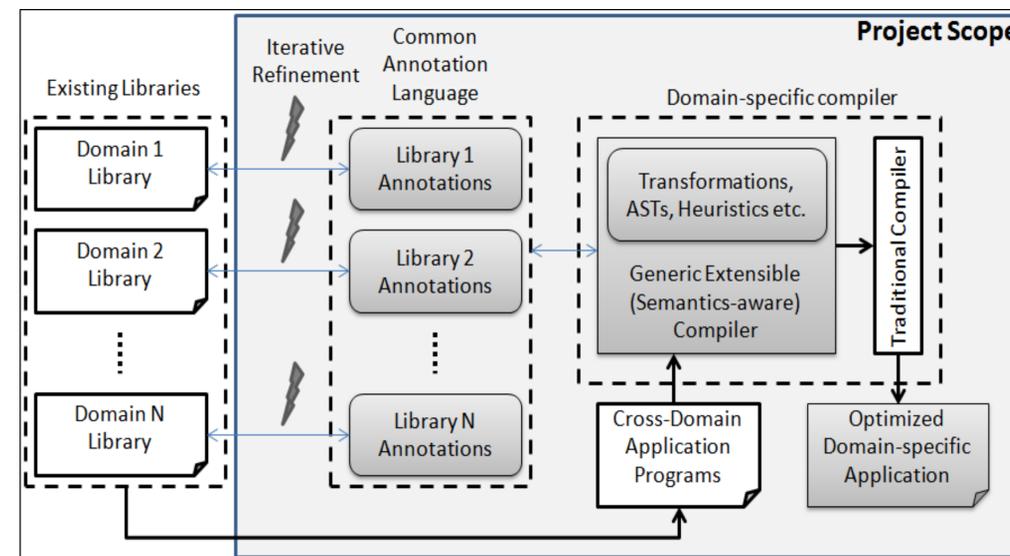
## Motivation

- ◆ Complex applications are composed of multiple, interacting library calls
- ◆ Traditional compilers are not able to optimize well across "black-box" components
- ◆ Nevertheless, domain libraries have rich semantic properties that can be exploited to optimize their use *if a compiler understood their semantics*
- ◆ Domain-specific compilers (DSCs) can help address this problem, but few shared resources between different DSCs, and no way to optimize across domains

## High level principles

- ◆ *Domain libraries* should convey semantic information (library writers are the domain experts, so are best-equipped to provide this information)
- ◆ Compiler should be *domain agnostic*, and able to perform transformations and optimizations independent of domain(s) of a program
- ◆ Compiler should be able to optimize for multiple objectives (not just performance!)

## Overview



## Approach

- ◆ Compiler provides high-level IR (e.g., expression trees or dataflow) to represent programs
- ◆ Annotations on domain libraries capture:
  - ◆ How to *abstract* library calls to IR
  - ◆ How to *transform* compositions of methods (e.g., by exploiting commutativity, associativity, etc.)
  - ◆ How to estimate cost(s) of methods
- ◆ Compilation process:
  - ◆ *Abstract* program into IR
  - ◆ *Transform* program to create multiple variants
  - ◆ *Search* space to find low-cost variant
  - ◆ *Lower* program back to original representation

## Research challenges

- ◆ What should the annotation language look? How can domain experts use it?
- ◆ How should the compiler generate search space of program variants and find lowest cost implementation?
- ◆ How can we deal with programs that cross domains?
- ◆ How can we deal with incomplete information?

## Deliverables

- ◆ Annotation language to capture semantic information
- ◆ Generic, extensible compiler that incorporates annotated libraries to create domain-specific compilers
- ◆ "Showcase" annotated domain libraries (e.g., Trilinos components)
- ◆ Proof-of-concept optimized applications across multiple domains

**PURDUE**  
UNIVERSITY

 Sandia  
National  
Laboratories

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.