

# ExaSlang: A Domain-Specific Language for Highly Scalable Multigrid Solvers

Christian Schmitt<sup>‡</sup>, Sebastian Kuckuk<sup>†</sup>, Frank Hannig<sup>‡</sup>, Harald Köstler<sup>†</sup>, Jürgen Teich<sup>‡</sup>

<sup>‡</sup>Hardware/Software Co-Design, <sup>†</sup>System Simulation,  
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)  
WOLFHPC, New Orleans, LA, USA; November 17, 2014







## Motivation

Why not concentrate on the algorithmic description?

```
Function Smoother () : Unit {  
  communicate Solution  
  loop over Solution {  
    Solution = Solution + 0.8 * (1.0 / diag(Laplace)) *  
      (RHS - Laplace * Solution)  
  }  
}
```

## Motivation

Why not concentrate on the algorithmic description?

```
Function Smoother () : Unit {  
  communicate Solution  
  loop over Solution {  
    Solution = Solution + 0.8 * (1.0 / diag(Laplace)) *  
      (RHS - Laplace * Solution)  
  }  
}
```

- **Productivity**
  - Algorithm description at high-level
  - Hide low-level details from programmer
- **Portability**
  - Support different target platforms from the same description
  - Support different target languages from the same description
- **Performance**
  - Portable: high performance on different target platforms
  - Competitive: comparable performance to hand-written code

# ExaSlang




## ExaSlang


- **ExaStencils language**
- Abstract description for generation of massively parallel geometric multigrid solvers
- Multi-layered structure → set of Domain-Specific Languages (DSLs)
- Top-down approach: From abstract to concrete
- Very few mandatory specifications at one layer  
→ room for decisions at lower layers based on domain knowledge
- External Domain-Specific Language
  - Better reflection of extensive ExaStencils approach
  - Enables greater flexibility of different layers
  - Eases tailoring of DSL layers to users
  - Enables code generation for large variety of target platforms

## Basic Multigrid Ideas


Residual on fine grid



Smoother applied



Residual on coarse grid







## Basic Multigrid Ideas

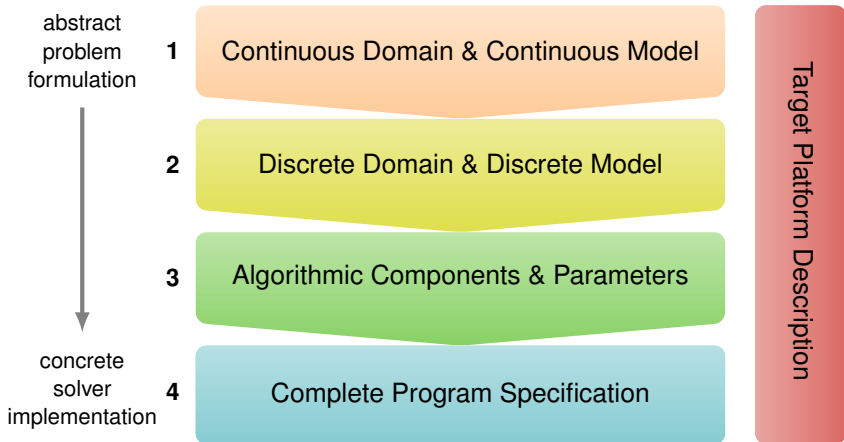
### Multigrid method

1. Pre-smoothing
2. Calculation of residual
3. Restriction
4. Recursive call(s) or solve (at coarsest level)
5. Prolongation
6. Correction
7. Post-smoothing



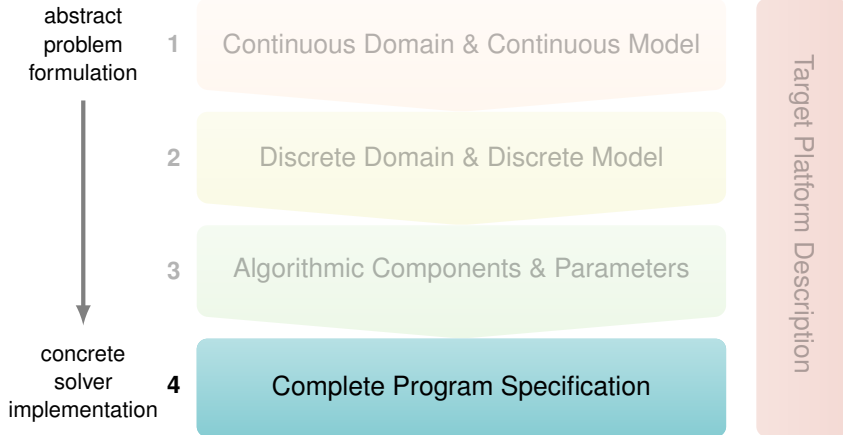
## ExaSlang: Multi-layered DSL Structure

Different layers of DSL tailored towards different users and knowledge.



## ExaSlang: Multi-layered DSL Structure

Different layers of DSL tailored towards different users and knowledge.



# ExaSlang 4: Complete Program Specification

## Properties

- Procedural
- Statically typed
- External DSL
- Syntax partly inspired by Scala

# ExaSlang 4: Complete Program Specification

## Properties

- Procedural
- Statically typed
- External DSL
- Syntax partly inspired by Scala

## Specification of

- Operations depending on the multigrid level
- Loops over computational domain
- Communication and data exchange
- Interface to 3rd-party code

# Data Types

## Simple and aggregate data types

- Real, Integer, String, Boolean
- Complex<Real>, Complex<Integer>

## Algorithmic data types

### Field

- Correspond to discretized (mathematical) variables
- Communication scheme via **Layout**
- Specify **Slot** number for multiple copies

### Stencil

- Correspond to discretized (mathematical) operators
- (Nearly) arbitrary expressions possible

## Computations

Loop over computational domain split into `loop over fragments`

- Fragments stem from distribution across different cluster nodes
- Corresponds to global operation
- Optionally: reduction operators

and `loop over <field>`

- Iteration over parts of fields possible
- Corresponds to local operation
- Optionally: Reduction operators

```

Function NormResidual @(coarsest and finest) () : Real {
  Variable res : Real = 0
  loop over fragments with reduction(+ : res) {
    loop over Residual @current with reduction(+ : res) {
      res += Residual @current * Residual @current
    }
  }
  return ( sqrt(res) )
}

```

## Level Specifications

### Multigrid is inherently hierarchical and recursive

→ We need

- Multigrid recursion exit condition
- Access to other levels' data & functions

→ Additionally, we want

- Relative addressing
- Aliases for certain levels
- Variable definitions per level

### Implementation

- Numerical values, e. g., @0 for bottom level
- Aliases, e. g., @all, @current, @coarser, @coarsest
- Simple expressions, e. g., @(coarsest + 1)
- Lists, e. g., @(1, 3, 5)
- Ranges, e. g., @(1 to 5)



## Level Specifications: Example

### Disjunct function definition

```

Function VCycle @((coarsest+1) to finest) () : Unit {
  repeat 3 times {
    Smoother @current ()
  }
  UpResidual @current ()
  Restriction @current ()
  SetSolution @coarser (0)
  VCycle @coarser ()
  Correction @current ()
  repeat 3 times {
    Smoother @current ()
  }
}

Function VCycle @coarsest () : Unit {
  /* ... solve directly ... */
}

```

## Level Specifications: Example

### Disjunct function definition

```

Function VCycle @((coarsest+1) to finest) () : Unit {
  repeat 3 times {
    Smoother @current ()
  }
  UpResidual @current ()
  Restriction @current ()
  SetSolution @coarser (0)
  VCycle @coarser ()
  Correction @current ()
  repeat 3 times {
    Smoother @current ()
  }
}

Function VCycle @coarsest () : Unit {
  /* ... solve directly ... */
}

```

## Level Specifications: Example

No disjunction needed due to overloading

```

Function VCycle @(coarsest to finest) () : Unit {
  repeat 3 times {
    Smoother @current ()
  }
  UpResidual @current ()
  Restriction @current ()
  SetSolution @coarser (0)
  VCycle @coarser ()
  Correction @current ()
  repeat 3 times {
    Smoother @current ()
  }
}

Function VCycle @coarsest () : Unit {
  /* ... solve directly ... */
}

```

## Level Specifications: Example

Level Specification can be simplified further

```

Function VCycle @all () : Unit {
  repeat 3 times {
    Smoother @current ()
  }
  UpResidual @current ()
  Restriction @current ()
  SetSolution @coarser (0)
  VCycle @coarser ()
  Correction @current ()
  repeat 3 times {
    Smoother @current ()
  }
}

Function VCycle @coarsest () : Unit {
  /* ... solve directly ... */
}

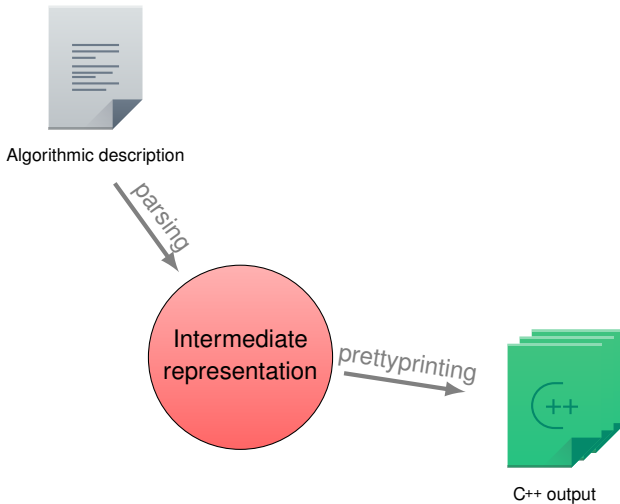
```

# ExaStencils Transformation Framework



# ExaStencils Framework

Abstract workflow:



## ExaStencils Framework

Using a simple 1-step concept, we can do some refinements, e. g.,

```
loop over Solution {
    // ....
}
```

is processed to

```
for (int z = start_z; z < stop_z; z += 1) {
    for (int y = start_y; y < stop_y; y += 1) {
        for (int x = start_x; x < stop_x; x += 1) {
            // ....
        }
    }
}
```

## ExaStencils Framework

Using a simple 1-step concept, we can do some refinements, e. g.,

```
loop over Solution {
    // ....
}
```

is processed to

```
for (int z = start_z; z < stop_z; z += 1) {
    for (int y = start_y; y < stop_y; y += 1) {
        for (int x = start_x; x < stop_x; x += 1) {
            // ....
        }
    }
}
```

But what about the calculations? What about more complex things?  
 Optional code modifications? Parallelization? Vectorization? Blocking?  
 Color splitting?

→ Very cumbersome with 1-step approach. Need something more flexible!



# ExaStencils Framework

## Current workflow

1. DSL input (Layer 4) is parsed
2. Parsed input is checked for errors and transformed into the IR
3. Many smaller, specialized transformations are applied
4. C++ output is prettyprinted

# ExaStencils Framework

## Current workflow

1. DSL input (Layer 4) is parsed
2. Parsed input is checked for errors and transformed into the IR
3. Many smaller, specialized transformations are applied
4. C++ output is prettyprinted

## Concepts

- Major program modifications take place only in IR
- IR can be printed to C++ code
- Small transformations can be enabled and arranged according to needs
- Central instance keeps track of generated program: [StateManager](#)
- Variant generation by duplicating program at different transformation stages

# ExaStencils Framework

## Transformations

- Transform program state into another one
- Are applied to program state in depth-first order
- May be applied to only a part of the program state
- Are grouped together in Strategies

# ExaStencils Framework

## Transformations

- Transform program state into another one
- Are applied to program state in depth-first order
- May be applied to only a part of the program state
- Are grouped together in Strategies

## Strategies

- Are applied in transactions
- Standard strategy that linearly executes all transformations is provided
- Custom strategies possible

# ExaStencils Framework

## Transactions

- Before execution, a snapshot of the program state is made
- May be committed or aborted

## Checkpoints

- A copy of program state during compilation
- Restoration of program states
- Acceleration of variant generation for design space exploration

## ExaStencils Framework

Example transformations:

```

var s = DefaultStrategy("example strategy")

// rename a certain stencil
s += Transformation("rename stencil", {
  case x : Stencil if(x.identifier == "foo")
    =>
    {
      if(x.entries.length != 7) error("invalid stencil size")
      x.identifier = "bar"; x
    }
})

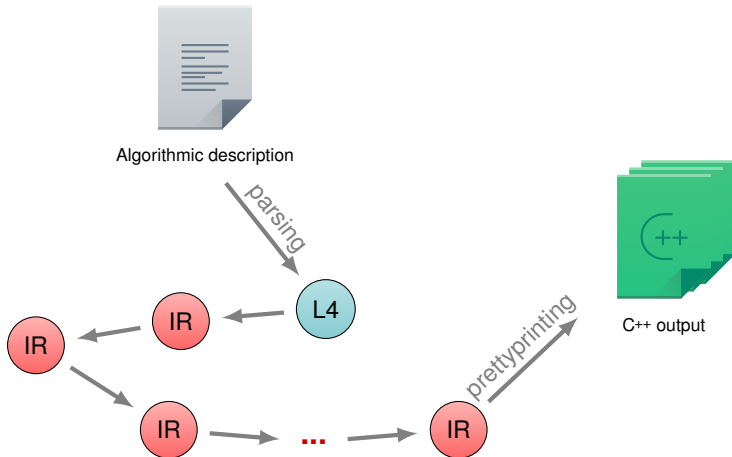
// evaluate additions
s += Transformation("eval adds", {
  case AdditionExpression(l : IntegerConstant, r : IntegerConstant)
    => IntegerConstant(l + r)
})

s.apply // execute transformations sequentially

```

# ExaStencils Framework

Implemented workflow:




# First Results






## Program Sizes during Transformation

- V(3,3) cycle, Jacobi smoother, CG coarse grid solver
- Hybrid MPI/OpenMP

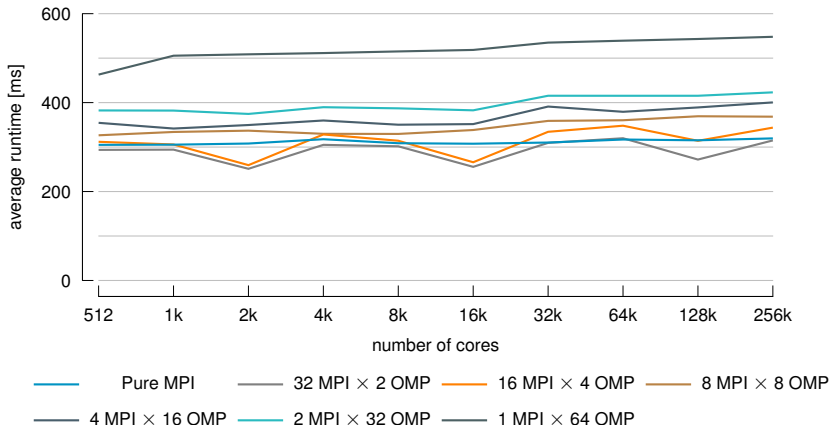


## Generated Lines of Code



## Weak-Scaling Results

Solution of Poisson's equation in 3D, V(3,3) cycle, Jacobi, CG  
 Average time per V-cycle on JUQUEEN



## Summary

### Presented

- Multi-layered DSL ExaSlang for multigrid-based numerical solvers
- Framework for specification of transformations and code generation
- Generation of highly scalable C++ code

### Conclusions

- Code generation is a viable approach to generation of multigrid codes
- Specialized DSLs allow for concise algorithmic descriptions (**Productivity**)
- Generation of solvers for different target platforms (**Portability**)
- More work on optimizations needed, but scalability already good (**Performance**)

Thanks for listening. Questions?

# ExaStencils

ExaStencils – Advanced Stencil Code Engineering

<http://www.exastencils.org>

ExaStencils is funded by the German Research Foundation (DFG)  
as part of the Priority Program 1648 (Software for Exascale Computing).