Metaprogramming for CSE Applications meets performance engineering

PD Dr. Harald Köstler ISC, Frankfurt 2017



Contents



Motivation

ExaStencils Code Generation Framework

- Applications
- Performance Engineering
- Next: Metaprogramming



Definition: Code generation in a strict sense is used to describe the process of producing code that can be executed on a certain platform from an abstract representation

Why? Because programming is time-consuming.

Definition: Metaprogramming is a very general term and can be defined as writing software that can itself produce new software.

The final goal is to develop software that can itself solve problems without human interaction.

Approaches: Genetic Programming, Deep Learning

CSE optimization problem





Metaprogramming group @ LSS

- Code generation
 - Parallel Data Structures
 - Performance engineering
 - GPU, accelerators
- Metaprogramming
 - Lisp, AnyDSL
 - Scala
 - Python, Ilvm
 - C++
 - Genetic Programming, Deep Learning
- Applications
 - Data analysis, medical image processing
 - Solvers for PDEs, DG discretizations
 - Tsunami simulation, material sciences, CFD





ExaStenci

SKAMPY







DAAD



The ExaStencils Framework





Project ExaStencils



- Sebastian Kuckuk
- Georg Altmann
- Harald Köstler
- Ulrich Rüde





- Alexander Grebhahn
- Sven Apel



- Christian Schmitt
- Frank Hannig
- Jürgen Teich





- Hannah Rittich
- Lisa Claus
- Matthias Bolten



BERGISCHE UNIVERSITÄT WUPPERTAL





- Stefan Kronawitter
- Armin Größlinger
- Christian Lengauer





- Scope:
 - (Geometric) **multigrid solvers** for
 - Finite difference and finite volume discretizations of elliptic PDEs
 - On regular (staggered) quad and hex meshes
- Multi-layered, external DSL as input, whole program as output
- Transformation-based code generation framework written in Scala
- Support for OpenMP, MPI and/or CUDA
- Many optimizations can be applied automatically (loop transformations, vectorization, sophisticated CSE, etc.)
- Applications from **CFD**, image processing and quantum chemistry





We propose a <u>multilayered</u>, <u>external</u> DSL



Schmitt, Kuckuk, Hannig, Köstler, Teich. ExaSlang: A Domain-Specific Language for Highly Scalable **Multigrid Solvers**.







```
/** Layer 2 **/
// domain - information is carried over from L1
Domain global
// fields from L1
Field Solution@all with Real on Node of global = 0.0
Field RHS@all with Real on Node of global = 0.0
// boundary conditions from L1
Field Solution@finest on boundary = vf boundaryCoord x^{**2} - 0.5 * ...
Field Solution@(all but finest) on boundary = 0.0
// (discretized) operators from L1
Operator Laplace from Stencil {
  [ 0, 0] => 2.0 / ( vf_gridWidth_x ** 2 ) + 2.0 / ( vf_gridWidth_y ** 2 )
  [-1, 0] = -1.0 / (vf gridWidth x ** 2)
  [1, 0] = -1.0 / (vf gridWidth x ** 2)
  [0, -1] = -1.0 / (vf gridWidth v ** 2)
  [ 0, 1] => -1.0 / ( vf gridWidth v ** 2 )
}
```



```
/** Layer 3 **/
/* fields and operators from L2 are progressed automatically */
// new fields
Field Residual from Solution
// override inherited boundary conditions where applicable
override bc for Residual@finest with 0.0
```

// generate default inter-grid operators for node-based discretizations
Operator RestrictionStencil from default restriction on Node with 'linear'
Operator CorrectionStencil from default prolongation on Node with 'linear'

```
// create smoother function
Function Smoother@all {
    Val omega : Real = 0.8
    repeat 3 times {
        Solution =
            Solution + omega * diag_inv ( Laplace ) * ( RHS - Laplace * Solution )
    }
}
```



```
/** Layer 3 **/
// create v-cycle function
Function VCycle@((coarsest + 1) to finest) {
  Smoother ()
  Residual = RHS - ( Laplace * Solution )
  RHS@coarser = RestrictionStencil * Residual
  Solution@coarser = 0.0
 VCycle@coarser ( )
  Solution += CorrectionStencil * Solution@coarser
  Smoother ()
}
Function VCycle@coarsest {
  /* CGS */
}
```



```
/** Layer 4 **/
// create main function
Function Application ( ) : Unit {
    /* init code */
    Var resStart : Real = NormResidual@finest ( )
    Var curRes : Real = resStart
    repeat until res_0 < 1.0E-5 * resStart_0 {</pre>
         VCycle@finest ( )
         curRes = NormResidual@finest ( )
     }
    /* de-init code */
}
```

}



```
loop over p {
  solve locally {
   u@[0, 0, 0] => A_u@[0, 0, 0] * u@[0, 0, 0] == rhs u@[0, 0, 0]
      + vf cellWidth y * vf cellWidth z * ( p@[-1, 0, 0] - p@[0, 0, 0] )
   u@[1, 0, 0] => A_u@[1, 0, 0] * u@[1, 0, 0] == rhs u@[1, 0, 0]
      + vf cellWidth y * vf cellWidth z * ( p@[0, 0, 0] - p@[1, 0, 0] )
   v@[0, 0, 0] => A v@[0, 0, 0] * v@[0, 0, 0] == rhs v@[0, 0, 0]
     + vf cellWidth x * vf cellWidth z * (p@[0, -1, 0] - p@[0, 0, 0])
   v@[0, 1, 0] => ...
   w@[0, 0, 0] => A_w@[0, 0, 0] * w@[0, 0, 0] == rhs_w@[0, 0, 0]
     + vf cellWidth x * vf cellWidth y * ( p@[0, 0, -1] - p@[0, 0, 0] )
   w@[0, 0, 1] => ...
   p@[0, 0, 0] => rhs p ==
        integrateOverEastFace ( u * rho ) - integrateOverWestFace ( u * rho )
      + integrateOverNorthFace ( v * rho ) - integrateOverSouthFace ( v * rho )
      + integrateOverTopFace ( w * rho ) - integrateOverBottomFace ( w * rho )
 }
```



Applications: It is all about Grids





- Work with Haase (U. Graz) and Vasco (U. Santiago de Chile)
- Simulation of non-Newtonian/non-isothermal fluids
- 3D FV discretization of staggered grids
- Recent extension towards FAS-FMG
- (Automatic) parallelization using OMP/MPI/CUDA



Vasco, Moraga, Haase. Parallel finite volume method simulation of three-dimensional fluid flow and convective heat transfer for viscoplastic non-Newtonian fluids

Kuckuk, Haase, Vasco, Köstler. Towards Generating Efficient Flow Solvers with the ExaStencils Approach

Scope: Towards Ocean Simulation



- Overall goal: simulation of ocean behavior
- Discretization using discontinuous Galerkin (DG) methods
- Performance and scalability





Hybrid Grids: Unstructured Grids are slow(er)!



Generation of Block-Structured Grids



1. Rough (unstructured) tessellation of the computational domain



Generation of Block-Structured Grids



- 1. Rough (unstructured) tessellation of the computational domain
- 2. Uniform subdivision





- 1. Rough (unstructured) tessellation of the computational domain
- 2. Uniform subdivision
- 3. Adaptation of vertex positions





- 1. Rough (unstructured) tessellation of the computational domain
- 2. Uniform subdivision
- 3. Adaptation of vertex positions
- 4. Triangulation





- What is already possible?
 - Finite volume discretizations for simple model problems on regular, non-uniform patches (cell-centered, 2D)
 - Using the complete ExaStencils pipeline
- What is missing?
 - Suitable abstractions and language extensions for DG
 - Extension to prisms
 - (Specialized) communication routines for triangle and prism data







Towards Performance Engineering









Algorithmic Performance Engineering



FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG TECHNISCHE FAKULTÄT





- For each choice you require a valid performance model that estimates the runtime for your settings!
- Therefore there is a need to automize the creation of performance models

Performance Model: Runtime Prediction





Gmeiner, Köstler et al, Parallel multigrid on hierarchical hybrid grids: a performance study on current high performance computing clusters, Concurrency and Comp.: Practice and Experience, 2012.

Performance Modeling



- We aim for an automatic derivation of performance models
- Prototype implementation (optimistic):



• Ongoing work: extension to ECM

Runtimes 3D Variable Coefficients

8

6

4

 $\mathbf{2}$

0

1

 $\mathbf{2}$

3

4

 $\mathbf{5}$

Multi Grid Level n

6

Multiples of Estimate



1

 $\mathbf{2}$

3

4

5

Multi Grid Level n

6

8

9

CPU dual socket Xeon E5620, 2.4 GHz clock, 4 cores / 8 hw-threads Microarchitecture Westmere-EP/Gulftown (2010) L3 Cache 12 MB smart cache Main Memory 24 GB Interconnect 2x QPI 5.86 GT/s

7

8

9

measured

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG TECHNISCHE FAKULTÄT



Towards Metaprogramming



What is the problem?

Goal: Solve partial differential equation

$$\Delta u = f \quad in \ \Omega$$
$$u = 0 \quad on \ \partial \Omega$$

After discretization one requires an efficient iterative solver for sparse systems

$$Au_h = f_h$$

Multigrid solver has complexity O(N)















- Linear system of interest as Ax = b, where A is the system matrix, x is the unknown vector and b is the right hand side vector.
- Denote the iteration matrix with M and it is represented as a tree
- The approximated solution in the n+1 iteration is obtained using the old solution and the iteration matrix using the relation $x_{n+1} = Mx_n$.

• A tree in any stage of the framework is a valid expression based on the defined matrix operations i.e. the nodes of the tree are selected such that the generated tree is consistent in terms of dimensions and the generated expressions are secured to be computable.



- Terminal Set: It is composed of arbitrary symbolic matrices.
 - **A**, **b**, diagonal of **A**, inverse of the diagonal of **A**, **A** minus its diagonal, lower triangular part of **A**, inverse of the lower triangular part of **A** and upper triangular (1st diagonal) part of **A**.

• Operation Set:

•The operation set consists of addition, subtraction and multiplication.

• Fitness Evaluation:

• The symbolic matrix corresponding to the genetic expression of an individual is simplified using the symbolic Math Toolbox of MATLAB and then evaluated to get the numerical version of the iteration matrix.

- spectral radius r is used to rank individuals, fitness function f for r <
- 1 is defined $f(r) = 1/(1 r^2)$ and for r > 1 as $f(r) = a_0 + a_1 r$



1D Example





$$M_{2} = \left[-\frac{a_{22}a_{31}a_{22}^{2} - a_{21}a_{23}a_{32}a_{22} + a_{21}a_{32}}{a_{11}a_{22}^{2}a_{33}}, \frac{a_{22} - a_{22}a_{22}a_{22}}{a_{22}^{2}a_{32}}, \frac{a_{22}}{a_{22}a_{32}}, \frac{a_{23}}{a_{22}a_{32}}, \frac{a_{23}}{a_{2}$$

2D Poisson Example



$$\boldsymbol{A} = \begin{bmatrix} 4 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 4 & -1 & 0 & -1 & 0 & 0 & 0 & -1 \\ 0 & -1 & 4 & 0 & 0 & -1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 4 & -1 & 0 & -1 & 0 & 0 \\ 0 & -1 & 0 & -1 & 4 & -1 & 0 & -1 & 0 \\ 0 & 0 & -1 & 0 & -1 & 4 & 0 & 0 & -1 \\ 0 & 0 & 0 & -1 & 0 & 0 & 4 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & -1 & 4 & -1 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & -1 & 4 \end{bmatrix}$$

$$\begin{split} & \boldsymbol{M_1} = [-6, \, 0, \, 0, \, 0, \, 0, \, 0, \, 0, \, 0] \text{e-}2 \ , \boldsymbol{M_2} = [-9, \, -6, \, 0, \, 0, \, 0, \, 0, \, 0, \, 0, \, 0] \\ & 0] \text{e-}2 \ , \boldsymbol{M_3} = [-3, \, -9, \, -6, \, 0, \, 0, \, 0, \, 0] \text{e-}2 \ , \boldsymbol{M_4} = [-9, \, 0, \, 0, \, -6, \, 0, \, 0, \, 0] \\ & 0, \, 0, \, 0] \text{e-}2 \ , \boldsymbol{M_5} = [-5, \, -9, \, 0, \, -9, \, -6, \, 0, \, 0, \, 0] \text{e-}2 \ , \boldsymbol{M_6} = [-2, \, -5, \, -9, \, -2, \, -9, \, -6, \, 0, \, 0, \, 0] \text{e-}2 \ , \boldsymbol{M_7} = [-3, \, 0, \, 0, \, -10, \, 0, \, 0, \, -6, \, 0, \, 0] \text{e-}2 \ , \boldsymbol{M_8} = [-2, \, -3, \, 0, \, -5, \, -10, \, 0, \, -10, \, -6, \, 0] \text{e-}2 \ , \boldsymbol{M_9} = [-1, \, -2, \, -3, \, -2, \, -5, \, -10, \, -3, \, -10, \, -6] \text{e-}2. \end{split}$$

Scalable iteration matrices?!



Comparsion of iteration matrices





Optimize pattern for parallelization!

Non-sparse and assymetric system





A simple preconditioner based on inv(A)!

Acknowledgements

• Funded by

- Bundesministerium für Bildung und Forschung
- KONWIHR. Bavarian project
- DFG SPP 1648/1 Software for Exascale computing

DFG Deutsche Forschungsgemeinschaft

Industry

н









des Bundesministeriums für Bildung und Forschung von Richtlinien zur Förderung von Forschungsvorhaben auf dem Gebiet "HPC-Software für skalierbare Parallertechener" im Rahmen des Förderprogramms "IKT 2020 - Forschung für Innovationen"

Bundesministerium für Bildung und Forschung







itenci

Thank you for your Attention!

Questions?



