

Automated Performance Modeling of the UG4 Simulation Framework

Andreas Vogel, Alexandru Calotoiu, Arne Nägel, Sebastian Reiter, Alexandre Strube, Gabriel Wittum, and Felix Wolf

Abstract Many scientific research questions such as the drug diffusion through the upper part of the human skin are formulated in terms of partial differential equations and their solution is numerically addressed using grid based finite element methods. For detailed and more realistic physical models this computational task becomes challenging and thus complex numerical codes with good scaling properties up to millions of computing cores are required. Employing empirical tests we presented very good scaling properties for the geometric multigrid solver in Reiter et al. (*Comput Vis Sci* 16(4):151–164, 2013) using the UG4 framework that is used to address such problems. In order to further validate the scalability of the code we applied automated performance modeling to UG4 simulations and presented how performance bottlenecks can be detected and resolved in Vogel et al. (10,000 performance models per minute—scalability of the UG4 simulation framework. In: Träff JL, Hunold S, Versaci F (eds) *Euro-Par 2015: Parallel processing, theoretical computer science and general issues*, vol 9233. Springer, Springer, Heidelberg, pp 519–531, 2015). In this paper we provide an overview on the obtained results, present a more detailed analysis via performance models for the components of the geometric multigrid solver and comment on how the performance models coincide with our expectations.

A. Vogel (✉) • A. Nägel • S. Reiter • G. Wittum
Goethe Universität Frankfurt, Frankfurt, Germany
e-mail: andreas.vogel@gcsc.uni-frankfurt.de; arne.naegel@gcsc.uni-frankfurt.de;
sebastian.reiter@gcsc.uni-frankfurt.de; wittum@gcsc.uni-frankfurt.de

A. Calotoiu • F. Wolf
Technische Universität Darmstadt, Darmstadt, Germany
e-mail: calotoiu@cs.tu-darmstadt.de; wolf@cs.tu-darmstadt.de

A. Strube
Jülich Supercomputing Center, Germany
e-mail: a.strube@fz-juelich.de

1 Introduction

The mathematical description for many important scientific and industrial questions is given by a formulation in terms of partial differential equations. Numerical simulations of the modeled systems via finite element and finite volume discretizations (e.g., [6, 10, 16]) can help to better understand the physical behavior by comparing with measured data and ideally provide the possibility to predict physical scenarios. Using detailed computational grids the discretization thereby leads to large sparse systems of equations and these matrix equations can be resolved using advanced methods of optimal order—such as the multigrid method (e.g., [6, 15]).

Looking at the variety of applications and the constantly growing computing resources on modern supercomputers the efficient solution of partial differential equations is an important challenge and it is advantageous to address the numerous problems with a common framework. Ideally, the framework should provide scalable and reusable components that can be applied in all of the fields of interest and serve as a common base for the construction of applications for concrete problems. To this end the UG software framework has been developed and a renewed implementation has been given in the current version 4.0 [35, 37] that pays special attention to parallel scalability.

In order to validate the scaling properties of the software framework on such architectures we carried out several scalability studies. Starting with a hand-crafted analysis we presented close to optimal weak scaling properties of the geometric multigrid solver in [27]. However, the study focused only on a few coarse-grained aspects leaving room for potential performance bottlenecks, that are not visible at current scales due to a small execution constant, but may become dominant at largest scales due to bad asymptotic behavior. Therefore, in a subsequent study we analyzed entire *UG 4* runs in [38] applying an automated performance modeling approach by Calotoiu et al. [7] to *UG 4* simulations. The modeling approach creates performance models at a function level granularity and uses few measurement runs at smaller core counts in order to predict the asymptotic behavior of each code kernel at largest scales. By detecting bad asymptotic behavior for code kernels in the grid setup phase we were able to detect and remove a performance bottleneck.

In this paper we focus on more detailed models for the geometric multigrid solver and explain how the observed performance models meet our expectations. Since the geometric multigrid solver is one of the crucial aspects for simulation runs in terms of scalability, we have evaluated in more depth the models for fine-grained kernels of the employed geometric multigrid solver and compare the observed behavior to the intended implementation.

The main aspects of this report are:

- Summarize the automated modeling approach and obtained results for its application to the simulation framework *UG 4*.
- Provide a detailed analysis for components of the geometric multigrid solver.
- Validation of the scaling behavior for the multigrid solver components.

The remainder of this paper is organized as follows. In Sect. 2, the *UG 4* simulation environment is presented with focus on the parallelization aspects of the parallel geometric multigrid (Sect. 2.2) and the skin permeation problem (Sect. 2.3) used in the subsequent studies. Section 3 outlines the performance modeling approach. In Sect. 4 we briefly summarize previously obtained analysis results for entire simulation runs and then present a detailed performance modeling for the geometric multigrid solver used in a weak scaling study for the skin problem. Sections 5 and 6 are dedicated to related work and concluding remarks.

2 The UG4 Simulation Framework

As a real world target application code for the performance modeling approach we will focus on the simulation toolbox *UG 4* [37]. The software framework is written in C++ and uses grid-based methods to numerically address the solution of partial differential equations via finite element or finite volume methods. With the main goal to address questions from biology, technology, geology and finance with one common effort, several components are reused in all types of applications and thus the performance modeling for those program parts provides insight into the performance of all these applications. In the following we give a brief overview on the used numerical methods and especially comment on the parallelization aspects.

2.1 Concepts and Numerical Methods

In order to construct the required geometries the meshing software ProMesh [25, 26] is used that shares code parts with the *UG 4* library. Meshes can be composed of different element types (e.g., tetrahedron, pyramid, prism and hexahedron in 3d) and subset assignment is used to distinguish parts of the domain with different physics or where boundary conditions are to be set. Once loaded in *UG 4* meshes are further processed to create distributed, unstructured, adaptive multigrid hierarchies with or without hanging nodes. Implemented load-balancing strategies [25] range from simple but fast bisection algorithms to more advanced strategies including usage of external algorithms such as ParMetis [22]. In this study, however, we restrict ourselves to a 3d hexahedral grid hierarchy generated through globally applied anisotropic refinement (cf. [38]). A study for adaptive hierarchies with hanging nodes is work in progress and will be considered in a subsequent study.

A flexible and combinable discretization module allows to combine different kinds of physical problems discretized by finite element and finite volume methods (e.g., [6, 10, 16]) and boundary conditions in a modular way to build a new physical problem selecting from basic building blocks [36, 37]. As algebraic structures for the discretized solutions and associated matrices, block vectors and a *CSR* (compressed sparse row) matrix implementation are provided. For the parallel solution of such

matrix equations several solvers are implemented, including Krylov methods such as CG and BiCGStab and preconditioners such as Jacobi, Gauss-Seidel, incomplete LU factorization, ILUT and block versions of these types (e.g., [28]). In addition, a strong focus is on multigrid methods (e.g. [15]) and geometric and algebraic multigrid approaches [17, 27].

The parallelization for the usage on massively parallel computing clusters with hundred thousands of cores is achieved using MPI. The separate library called PCL (parallel communication layer, [27]) builds on top of MPI and is used to ease the graph-based parallelization. Both, the parallelization of the computing grid—assigning a part of the multigrid hierarchy to each process—and of the algebraic structures are programmed based on the PCL. By storing parallel copies on each process in a well-defined order in interface containers identification is performed in an efficient way [25, 27, 37], and global IDs are dispensable.

In order to hide parallelization aspects and ease the usage for beginners the scripting language Lua [18] is used as end-user interface. A flexible plugin system allows to add additional functionality if required.

2.2 Parallel Hierarchical Geometric Multigrid

The multigrid method [16] is used to solve large sparse systems of equations that arise typically by the discretization of some partial differential equation. We briefly recap the idea of the algorithm and our modifications and implementation [27] for the parallel version. Given the linear equation system $\mathbf{A}_L \mathbf{x}_L = \mathbf{b}_L$ on the finest grid level L , the desired solution \mathbf{x}_L is computed iteratively: starting with some arbitrary initial guess \mathbf{x}_L , in every iteration the defect $\mathbf{d}_L = \mathbf{b}_L - \mathbf{A}_L \mathbf{x}_L$ is used to compute a multigrid correction $\mathbf{c}_L = \mathbf{M}_L(\mathbf{d}_L)$, where \mathbf{M}_L is the multigrid operator, that is added to the approximate solution $\mathbf{x}_L := \mathbf{x}_L + \mathbf{c}_L$. In order to compute the correction \mathbf{c}_L not only the fine grid matrix \mathbf{A}_L is used but several auxiliary coarse grid matrices $\mathbf{A}_l, L_B \leq l \leq L$, are employed, where L_B denotes the base level. The multigrid cycle is then defined in a recursive manner: given a defect \mathbf{d}_l on a certain level l the correction is first partly computed via a smoothing operator (e.g. Jacobi iteration). Then the defect is transferred to the next coarser level, where the algorithm is applied to the restricted defect \mathbf{d}_{l-1} . The thereby computed coarse grid correction \mathbf{c}_{l-1} is then prolonged to the finer level and added to the correction on level l , followed by some postsmoothing. Once the algorithm reaches the base level L_B , the correction is computed exactly as $\mathbf{c}_l = \mathbf{A}_l^{-1} \mathbf{d}_l$ by, e.g., using LU factorization. Algorithm 1 summarizes this procedure.

The matrix equations for complex problems can easily grow beyond the size of billions of unknowns. In order to solve such problems, massively parallel linear solvers with optimal complexity have to be used. The multigrid algorithm only depends linearly on the number of unknowns and therefore good weak scaling properties are to be expected. As demonstrated in [27] geometric multigrid solvers

Algorithm 1 $\mathbf{c}_l = \mathbf{M}_l(\mathbf{d}_l)$ [16, 27]

```

Requirement:  $\mathbf{d}_l = \mathbf{b}_l - \mathbf{A}_l \mathbf{x}_l$ 
if  $l = L_B$  then
  Base solver:  $\mathbf{c}_l = \mathbf{A}_l^{-1} \mathbf{d}_l$ 
  return  $\mathbf{c}_l$ 
else
  Initialization:  $\mathbf{d}_l^0 := \mathbf{d}_l, \mathbf{c}_l^0 := 0$ 
  (Pre-)Smoothing for  $k = 1, \dots, \nu_1$ :
     $\mathbf{c} = S_l(\mathbf{d}_l^{k-1}),$ 
     $\mathbf{d}_l^k = \mathbf{d}_l^{k-1} - \mathbf{A}_l \mathbf{c}, \quad \mathbf{c}_l^k = \mathbf{c}_l^{k-1} + \mathbf{c}$ 
  Restriction:  $\mathbf{d}_{l-1} = \mathbf{P}_l^T \mathbf{d}_l^{\nu_1}$ 
  Coarse grid correction:  $\mathbf{c}_{l-1} = \mathbf{M}_{l-1}(\mathbf{d}_{l-1})$ 
  Prolongation:
     $\mathbf{c}_l^{\nu_1+1} = \mathbf{c}_l^{\nu_1} + \mathbf{P}_l \mathbf{c}_{l-1},$ 
     $\mathbf{d}_l^{\nu_1+1} = \mathbf{d}_l^{\nu_1} - \mathbf{A}_l \mathbf{P}_l \mathbf{c}_{l-1}$ 
  (Post-)Smoothing for  $k = 1, \dots, \nu_2$ :
     $\mathbf{c} = S_l(\mathbf{d}_l^{\nu_1+k}),$ 
     $\mathbf{d}_l^{\nu_1+1+k} = \mathbf{d}_l^{\nu_1+k} - \mathbf{A}_l \mathbf{c}, \quad \mathbf{c}_l^{\nu_1+1+k} = \mathbf{c}_l^{\nu_1+k} + \mathbf{c}$ 
  return  $\mathbf{c}_l^{\nu_1+1+\nu_2}$ 

```

can exhibit nearly perfect weak scalability when employed in massively parallel environments with hundred thousands of computing cores.

To this end, the components of the algorithm must be parallelized. The basic idea is to construct a distributed multigrid hierarchy as follows:

1. Start with a coarse grid on a small number of processes.
2. Refine the grid several times to create additional hierarchy levels.
3. Redistribute the finest level of the hierarchy to a larger set of processes.
4. Repeat at (2) until the desired grid resolution is obtained. At this point all active processes should contain a part of the finest level of the multigrid hierarchy.

Refining the grid, new levels of the multigrid hierarchy are created and after some refinements the finest grid level is distributed to a larger set of processes and communication structures (called *vertical interfaces*) are established. This process can be iterated, successively creating a tree structure of processes holding parts of the hierarchical grid. Figure 1 shows a process hierarchy for a distributed multigrid hierarchy on four processes (cf. [25, 27]). The communication structures in vertical direction are used to parallelize the transfer between the grid levels, i.e. to implement the transfer of data between grid levels at restriction and prolongation phases within a multigrid cycle. However, if no vertical interface is present the transfer operators act completely process-locally. For the communication within multigrid smoothers on each grid level additional *horizontal interfaces* are required. These interfaces will be used to compute the level-wise correction in a consistent way. An illustration for the resulting hierarchy distribution and interfaces is given in Fig. 2 (cf. [25, 27, 36, 38]). In order to compute the required coarse grid matrices, each process calculates the contribution of the grid part assigned to the process itself.




Fig. 1 Illustration for a 1d process hierarchy on 4 processes. Ghost elements (*red*) are sent during redistribution. Data is communicated between ghosts and actual elements through vertical interfaces (*orange*) (cf. [25, 27])





Fig. 2 Illustration for a 1d parallel multigrid hierarchy distributed onto two processes. Parallel copies are identified via horizontal (*blue*) and vertical interfaces (*orange*) (From [38], cf. also [27])

Thus, the matrices are stored in parallel in an additive fashion and no communication is required for this setup.

A Jacobi smoother has very good properties regarding scalability, however it may not be suitable for more complicated problems (e.g. with anisotropic coefficients or anisotropic grids). To handle this issue for anisotropic problems, we use anisotropic refinement in order to construct grid hierarchies with isotropic elements from anisotropic coarse grids: refining only those edges in the computing grid that are longer than a given threshold, and halving this threshold in each step, the approach yields a grid hierarchy which contains anisotropic elements on lower levels and more and more isotropic elements on higher levels. An illustration for a resulting hierarchy is shown in Fig. 3. The used refinement strategy produces non-adaptive grids, i.e. meshes that fully cover the physical domain. This eases the load-balancing compared to adaptive meshes where huge differences in the spatial resolution and thereby element distribution may occur during refinement and redistribution is necessary. In this work we focused on the non-adaptive strategy only, however, plan to report on the adaptive case in future works.

Reconsidering the hierarchical distribution approach described above, lower levels of the multigrid hierarchy are only contained on a smaller number of processes. This is well suited for fast parallel smoothing, prolongation and restriction operations thanks to maintaining a good ratio between computation and communication costs on all levels. A smoothing operation on coarser levels with

Fig. 3 Grid hierarchy created by anisotropic refinement for the 3d brick-and-mortar model (in exploded view). The aspect ratio of the grid elements improve with every refinement step



only few inner unknowns would be dominated by the communication and thus the work is agglomerated to fewer processes resulting in idle processes on coarse levels. However, the work on finer grid levels is dominating the overall runtime.

2.3 Application: Human Skin Permeation

As an exemplary application from the field of computational pharmacy we focus on the permeation of substances through the human skin. These simulations consider the outermost part of the epidermis, called stratum corneum, and are used to estimate the throughput of chemical exposures. An overview on different descriptions of the biological and geometric approaches to simulate such a setting can be found in [20] and the references therein. For this study, we use the same setup as used in [38]: the transport in two subdomains $s \in \{cor, lip\}$ (corneocyte, lipid) is described by the diffusion equation

$$\partial_t c_s(t, \mathbf{x}) = \nabla \cdot (\mathbf{D}_s \nabla c_s(t, \mathbf{x})) ,$$

using a subdomain-wise constant diffusion coefficient \mathbf{D}_s . We use a 3d brick-and-mortar model consisting of highly anisotropic hexahedral elements with aspect ratios as bad as 1/300 in the coarse grid. Employing anisotropic refinements we construct a grid hierarchy with better and better aspect ratios on finer levels. The resulting grid hierarchy is displayed in Fig. 3. For a more detailed presentation we refer to [38].

3 Automated Performance Modeling

The automated modeling approach used to analyze the *UG* 4 framework has been presented by Calotoiu et al. in [7, 8]. Here, we give a brief overview on the procedure and ideas. For further details please refer to [7, 30, 38, 40].

Automated performance modeling is used to empirically determine the asymptotic scaling behavior for a large number of fine-grained code kernels. These scaling models can then be inspected and compared to the expected complexity: a discrepancy indicates a potential *scalability bug* that can be addressed and hopefully removed by the code developers. If no such scalability issues are found this can be taken as a strong evidence that no unexpected scalability problems are present. In addition the created models can also be used to predict the resource consumption at larger core counts if required.

In order to create the models, the simulation framework *UG* 4 has been instrumented to measure relevant metrics such as time, bytes sent/received or floating-point operations in program regions at a function level granularity. Running simulations at different core counts now offers the opportunity to determine via cross-validation [24] which choice of parameters in the *performance model normal form* (PMNF, [7])

$$f(p) = \sum_{k=1}^n c_k \cdot p^{i_k} \cdot \log_2^{j_k}(p),$$

with $i_k, j_k \in I, J \subset \mathbb{Q}$, best fits the measurements. The approach is applicable to strong and weak scaling. In this study, however, we have focused on weak scaling only. In order to account for jitter, several runs for every core count have to be executed. The required effort for this approach therefore is to run the application a few times at a few core counts.

For the correct analysis of the multigrid algorithm in weak scaling studies, a more careful approach than just analyzing the code kernels directly has to be taken [38]. This is due to the following observation: within a weak scaling study the problem size has to be increased and for multigrid approaches this leads to an increase in the number of grid levels. The multigrid algorithm—traversing the multigrid hierarchy top-down and then again bottom-up, applying smoothers and level transfers for every level—will thus create a different call tree at different core counts for multigrid functions due to the varying numbers of grid levels. However, the performance modeling approach usually assumes the same call tree for all core counts. Therefore, we preprocess the call tree: only kernels present in all measurements remain in the modified call tree and measurement data of code kernels not present in modified call trees is added to the parent kernel. This approach is not only limited to multigrid settings but may be useful for all codes that use recursive calls whose invocation count increases within a scaling study.

4 Results

The automated performance modeling (Sect. 3) has been applied to entire simulation runs of the *UG 4* simulation framework (Sect. 2) and several aspects of the analysis have been reported in [38]. In order to analyze and validate the code behavior the proceeding and reasoning is as follows: we run several simulations at different core counts p and measure detailed metric information (times, bytes sent) at a function level granularity. By this, we receive fine grained information for small code kernels. For all of these kernels and all available metrics we create performance models and then rank these by their asymptotic behavior with respect to the core count. All code kernels with constant or only logarithmical dependency are considered optimal. However, if some code kernel, e.g., in the multigrid method would show a linear or quadratic dependency, this would not match our performance expectations and we consider it a scalability bug that has to be addressed and removed. Inspecting all measured code kernels thus provides us with a fine grained information for different parts of the simulation code. Given that all code kernels show an optimal dependency we finally obtain a validation of the expected scaling properties.

Here, we first briefly give an overview on the results presented in [38] and then show more detailed results focusing on the multigrid kernels and their scaling properties.

4.1 Analysis for Grid Hierarchy Setup and Solver Comparison

In a first test, we analyzed entire runs in a weak scaling study for the human skin permeation simulating the steady-state concentration distribution on a three-dimensional brick-and-mortar skin geometry. A scalability issue has been detected by the performance modeling that can be explained and resolved [38]: at the initialization of the multigrid hierarchy an `MPI_Allreduce` operation for an array of length p was used to inform each process about its intended communicator group membership. The resulting $p \cdot \mathcal{O}(\text{MPI_Allreduce})$ dependency has been addressed by using `MPI_Comm_split` instead that can be implemented with a $\mathcal{O}(\log^2 p)$ [31] behavior. By this, we were able to remove this potential bottleneck at this stage of code development.

In a second study, we provided a comparison for two different types of solvers: the geometric multigrid solver has been compared in a weak scaling study to the unpreconditioned conjugate gradient (CG) method. The unpreconditioned conjugate gradient method is known to have unpleasant weak-scaling properties due to the increase by a factor of two for the iteration count resulting in a $\mathcal{O}(\sqrt{p})$ dependency (see [38] for a detailed theoretical analysis). Due to the created models we confirmed that the theoretical expectations are met by our implementation of the parallel solvers.

4.2 Scalability of Code Kernels in the Geometric Multigrid

In this section we give a more detailed view on the code kernels for the geometric multigrid. For the analysis of the multigrid solver we consider the human skin permeation model: we compute the steady-state of the concentration distribution for the brick-and-mortar model described in Sect. 2.3 and choose the diffusion parameter to $\mathbf{D}_{cor} = 10^{-3}$ and $\mathbf{D}_{lip} = 1$. For the solution of the arising linear system of equations, the geometric multigrid solver is used. As acceleration an outer conjugate gradient method is applied. For the smoothing a damped Jacobi is employed with three smoothing steps. As cycle type the V-cycle is used and as base solver we use a LU factorization. The stopping criterion for the solver is the reduction of the absolute residuum to 10^{-10} . The anisotropic refinement as laid out in Sect. 2.2 is used to enhance the aspect ratios of the hierarchy from level to level. Once satisfactory ratios are reached, this level is used as base level for the multigrid algorithm.

In Fig. 4 we present the accumulated wallclock times for exemplary coarse-grain kernels of the multigrid method and provide information on the number of used cores and the size of the solved matrix system (degrees of freedom). Please note that the iteration counts are bounded as expected for a multigrid method. Since the assembling for the matrix is an inherent parallel process without any communication it can be performed—given an optimal load-balancing—with constant wall-clock time in the weak scaling. This is confirmed by the generated performance model. All other shown aspects of the multigrid method show a logarithmical dependency. This is due to the fact that the number of involved coarse grid levels $L = \mathcal{O}(\log p)$ depends on the number of processes in a weak scaling. We consider this logarithmical dependency still as optimal since even allreduce operations implemented in a tree-like fashion will show the same behavior and are used to check for convergence.

The performance models for several code kernels are shown in Table 1. Please note that all code kernels in our measurements have shown constant or logarithmical dependency with respect to the number of processes. Here, we show some selected kernels in order to give more details on the parallelization aspects of the multigrid method. For a more detailed description on the mathematical algorithm and parallelization aspects we refer to Sect. 2 and [27].

The presmoothing is performed in a two step fashion: first, the Jacobi iteration is applied on process-wise data structures resulting in no data transfer (CG \rightarrow GMG \rightarrow PreSmooth \rightarrow Jacobi \rightarrow apply \rightarrow step). In a second phase update information is exchanged between nearest neighbors in order to gain a consistent update resulting in data transfer (PreSmooth \rightarrow Jacobi \rightarrow apply \rightarrow AdditiveToConsistent \rightarrow MPI_Isend). All behaviors are found to depend logarithmically due to the increase in grid levels that are using this method.

The grid transfer is performed process-wise as well (Restrict \rightarrow apply). No communication is needed unless vertical interfaces are present. The setup phase (Restrict \rightarrow init) simply assembles the transfer operators into a matrix structure on each process and a constant time within a weak scaling is thus observed.




Fig. 4 Measured accumulated wallclock times (*marks*) and models (*lines*) for the skin 3d problem (self time and subroutines). **(a)** (cf. [38]) Initialization of the multigrid solver and time spent in the multigrid cycle. **(b)** Times for coarse matrix assembling, smoothing and prolongation. **(c)** (From [38]) Number of grid refinements (L), degrees of freedom (DoF) and number of iterations of the solver (n_{gmg}). **(d)** Performance models for the kernels shown in the graphs

Finally, we show some kernels for the outer CG iteration. In order to check for convergence, the norm of a defect vector is computed in each iteration step. After a nearest neighbor communication in order to change the storage type of the vector (CG \rightarrow norm \rightarrow AdditiveToUnique \rightarrow MPI_Isend), the norm is first computed on each process (CG \rightarrow norm) and then summed up globally (CG \rightarrow norm \rightarrow allreduce \rightarrow MPI_Allreduce).

This way our expectations for the code kernels of the multigrid solver are confirmed and we have strong evidence that only logarithmical complexity with respect to the core count (or better) occurs.

Table 1 Skin 3d study. Models for self-time and bytes sent for selected kernels of the geometric multigrid method with outer CG iteration. $|1 - R^2|$, the absolute difference between R^2 and the optimum scaled by 10^{-3} , can be considered a normalized error

Kernel	Time		Bytes sent	
	Model	$ 1 - R^2 $ [10^{-3}]	Model	$ 1 - R^2 $ [10^{-3}]
CG → GMG → PreSmooth → Jacobi → apply → step	time = $f(p)$ [ms] $18.9 + 0.4 \cdot \log p$	42.6	bytes = $f(p)$ 0	0.0
CG → GMG → PreSmooth → Jacobi → apply → AdditiveToConsistent → MPI_Isend	$1.51 + 1.12 \cdot \log p$	36.0	$5.77 \cdot 10^5 + 0.95 \cdot 10^5 \cdot \log p$	53.2
CG → GMG → Restrict → init	510	0.0	0	0.0
CG → GMG → Restrict → apply	$51.0 + 0.05 \cdot \log p$	378	0	0.0
CG → norm	$3.52 + 0.002 \cdot \log^2 p$	544	0	0.0
CG → norm → AdditiveToUnique → MPI_Isend	$0.52 + 0.45 \cdot \log p$	38.9	$1.95 \cdot 10^5 + 0.34 \cdot 10^5 \cdot \log p$	45.5
CG → norm → allreduce → MPI_Allreduce	$1.67 + 0.92 \cdot \log^2 p$	7.5	$\mathcal{O}(\text{MPI_Allreduce})$	0.0

5 Related Work

Numerous analytical and automated performance modeling approaches have been proposed and developed. The field ranges from manual models [5, 23], over source-code annotations [34] to specialized languages [32]. Automated modeling methods are developed based on machine-learning approaches [19], and via extrapolating trace measurements in [42] (extrapolating from single-node to parallel architectures), in [41] (predicting communication costs at large core counts) and in [9] (extrapolating based on a set of canonical functions). The Dimemas simulator provides tools for performance analysis in message-passing programs [13].

Various frameworks to solve partial differential equations use multigrid methods. Highly scalable multigrid methods are presented in [4, 14, 29, 33], and [39] for geometric multigrid, and in [1, 2], and [3] for algebraic multigrid methods. Work on performance modeling for multigrid can be found in [11, 38] for geometric and in [12] for algebraic multigrid. For an overview for the numerical treatment of skin permeation, we refer to [21] and the references therein.

6 Conclusion

The numerical simulation framework UG4 consists of half a million lines of code and is used to address problems formulated in terms of partial differential equations employing multigrid methods to solve arising large sparse matrix equations. In order to analyze, predict and improve the scaling behavior of UG4 we have conducted a performance modeling at code kernel granularity. Inspecting automated performance models we validated the scalability of entire simulations and presented the close to optimal weak scaling properties for the components of the employed geometric multigrid method that only depend logarithmically on the core count.

Acknowledgements Financial support from the DFG Priority Program 1648 *Software for Exascale Computing* (SPPEXA) is gratefully acknowledged. The authors also thank the Gauss Centre for Supercomputing (GCS) for providing computing time on the GCS share of the supercomputer JUQUEEN at Jülich Supercomputing Centre (JSC).

References

1. Baker, A., Falgout, R., Kolev, T., Yang, U.: Multigrid smoothers for ultra-parallel computing. *SIAM J. Sci. Comput.* **33**, 2864–2887 (2011)
2. Baker, A.H., Falgout, R.D., Gamblin, T., Kolev, T.V., Schulz, M., Yang, U.M.: Scaling algebraic multigrid solvers: on the road to exascale. In: *Competence in High Performance Computing 2010*, pp. 215–226. Springer, Berlin/New York (2012)
3. Bastian, P., Blatt, M., Scheichl, R.: Algebraic multigrid for discontinuous Galerkin discretizations of heterogeneous elliptic problems. *Numer. Linear Algebra* **19**(2), 367–388 (2012)

4. Bergen, B., Gradl, T., Rude, U., Hulsemann, F.: A massively parallel multigrid method for finite elements. *Comput. Sci. Eng.* **8**(6), 56–62 (2006)
5. Boyd, E.L., Azeem, W., Lee, H.H., Shih, T.P., Hung, S.H., Davidson, E.S.: A hierarchical approach to modeling and improving the performance of scientific applications on the KSR1. In: *Proceedings of the 1994 International Conference on Parallel Processing*, St. Charles, vol. III, pp. 188–192. IEEE (1994)
6. Braess, D.: *Finite Elemente*. Springer, Berlin (2003)
7. Calotoiu, A., Hoefler, T., Poke, M., Wolf, F.: Using automated performance modeling to find scalability bugs in complex codes. In: *Proceedings of the ACM/IEEE Conference on Supercomputing (SC13)*, Denver. ACM (2013)
8. Calotoiu, A., Hoefler, T., Wolf, F.: Mass-producing insightful performance models. In: *Workshop on Modeling & Simulation of Systems and Applications*, Seattle, Aug 2014
9. Carrington, L., Laurenzano, M., Tiwari, A.: Characterizing large-scale HPC applications through trace extrapolation. *Parallel Process. Lett.* **23**(4), 1340008 (2013)
10. Ciarlet, P.G., Lions, J.: *Finite Element Methods (Part 1)*. North-Holland, Amsterdam (1991)
11. Gahvari, H., Gropp, W.: An introductory exascale feasibility study for FFTs and multigrid. In: *International Symposium on Parallel & Distributed Processing (IPDPS)*, pp. 1–9. IEEE, Piscataway (2010)
12. Gahvari, H., Baker, A.H., Schulz, M., Yang, U.M., Jordan, K.E., Gropp, W.: Modeling the performance of an algebraic multigrid cycle on HPC platforms. In: *Proceedings of the International Conference on Supercomputing*, pp. 172–181. ACM, New York (2011)
13. Girona, S., Labarta, J., Badia, R.M.: Validation of dimemas communication model for MPI collective operations. In: *Proceedings of the 7th European PVM/MPI Users' Group Meeting on Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pp. 39–46. Springer, London (2000). <http://dl.acm.org/citation.cfm?id=648137.746640>
14. Gmeiner, B., Köstler, H., Stürmer, M., Rude, U.: Parallel multigrid on hierarchical hybrid grids: a performance study on current high performance computing clusters. *Concurr. Comput.: Pract. Exp.* **26**(1), 217–240 (2014)
15. Hackbusch, W.: *Multi-grid Methods and Applications*, vol. 4. Springer, Berlin/New York (1985)
16. Hackbusch, W.: *Theorie und Numerik elliptischer Differentialgleichungen: mit Beispielen und Übungsaufgaben*. Teubner (1996)
17. Heppner, I., Lampe, M., Nägel, A., Reiter, S., Rupp, M., Vogel, A., Wittum, G.: Software framework ug4: parallel multigrid on the hermit supercomputer. In: *High Performance Computing in Science and Engineering '12*, pp. 435–449. Springer, Cham (2013)
18. Ierusalimsky, R., De Figueiredo, L.H., Celes Filho, W.: Lua-an extensible extension language. *Softw. Pract. Exp.* **26**(6), 635–652 (1996)
19. Lee, B.C., Brooks, D.M., de Supinski, B.R., Schulz, M., Singh, K., McKee, S.A.: Methods of inference and learning for performance modeling of parallel applications. In: *Proceedings of the 12th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '07)*, pp. 249–258. ACM, New York (2007)
20. Nägel, A., Heisig, M., Wittum, G.: A comparison of two- and three-dimensional models for the simulation of the permeability of human stratum corneum. *Eur. J. Pharm. Biopharm.* **72**(2), 332–338 (2009)
21. Nägel, A., Heisig, M., Wittum, G.: Detailed modeling of skin penetration—an overview. *Adv. Drug Deliv. Rev.* **65**(2), 191–207 (2013)
22. ParMetis (Nov 2015), <http://glaros.dtc.umn.edu/gkhome/metis/parmetis/overview>
23. Petrini, F., Kerbyson, D.J., Pakin, S.: The case of the missing supercomputer performance: achieving optimal performance on the 8,192 processors of ASCI Q. In: *Proceedings of the ACM/IEEE Conference on Supercomputing (SC'03)*, pp. 55ff. ACM, New York (2003)
24. Picard, R.R., Cook, R.D.: Cross-validation of regression models. *J. Am. Stat. Assoc.* **79**(387), 575–583 (1984)
25. Reiter, S.: *Efficient algorithms and data structures for the realization of adaptive, hierarchical grids on massively parallel systems*. Ph.D. thesis, University of Frankfurt, Germany (2014)

26. Reiter, S.: Promesh (Nov 2015), <http://www.promesh3d.com>
27. Reiter, S., Vogel, A., Heppner, I., Rupp, M., Wittum, G.: A massively parallel geometric multigrid solver on hierarchically distributed grids. *Comput. Vis. Sci.* **16**(4), 151–164 (2013)
28. Saad, Y.: *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia (2003)
29. Sampath, R., Biros, G.: A parallel geometric multigrid method for finite elements on octree meshes. *SIAM J. Sci. Comput.* **32**, 1361–1392 (2010)
30. Shudler, S., Calotoiu, A., Hoefler, T., Strube, A., Wolf, F.: Exascalng your library: will your implementation meet your expectations? In: *Proceedings of the International Conference on Supercomputing (ICS)*, Newport Beach, pp. 1–11. ACM (2015)
31. Siebert, C., Wolf, F.: Parallel sorting with minimal data. In: *Recent Advances in the Message Passing Interface*, pp. 170–177. Springer, Berlin/New York (2011)
32. Spafford, K.L., Vetter, J.S.: Aspen: a domain specific language for performance modeling. In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC '12*, pp. 84:1–84:11. IEEE Computer Society Press, Los Alamitos (2012)
33. Sundar, H., Biros, G., Burstedde, C., Rudi, J., Ghattas, O., Stadler, G.: Parallel geometric-algebraic multigrid on unstructured forests of octrees. In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, p. 43. IEEE Computer Society Press, Los Alamitos (2012)
34. Tallent, N.R., Hoisie, A.: Palm: easing the burden of analytical performance modeling. In: *Proceedings of the International Conference on Supercomputing (ICS)*, pp. 221–230. ACM, New York (2014)
35. UG4 (Nov 2015), <https://github.com/UG4>
36. Vogel, A.: Flexible und kombinierbare Implementierung von Finite-Volumen-Verfahren höherer Ordnung mit Anwendungen für die Konvektions-Diffusions-, Navier-Stokes- und Nernst-Planck-Gleichungen sowie dichtegetriebene Grundwasserströmung in porösen Medien. Ph.D. thesis, Universität Frankfurt am Main (2014)
37. Vogel, A., Reiter, S., Rupp, M., Nägel, A., Wittum, G.: UG 4: a novel flexible software system for simulating PDE based models on high performance computers. *Comput. Vis. Sci.* **16**(4), 165–179 (2013)
38. Vogel, A., Calotoiu, A., Strube, A., Reiter, S., Nägel, A., Wolf, F., Wittum, G.: 10,000 performance models per minute—scalability of the UG4 simulation framework. In: Träff, J.L., Hunold, S., Versaci, F. (eds.) *Euro-Par 2015: Parallel Processing, Theoretical Computer Science and General Issues*, vol. 9233, pp. 519–531. Springer, Heidelberg (2015)
39. Williams, S., Lijewski, M., Almgren, A., Straalen, B.V., Carson, E., Knight, N., Demmel, J.: s-step Krylov subspace methods as bottom solvers for geometric multigrid. In: *28th International Parallel and Distributed Processing Symposium*, pp. 1149–1158. IEEE, Piscataway (2014)
40. Wolf, F., Bischof, C., Hoefler, T., Mohr, B., Wittum, G., Calotoiu, A., Iwainsky, C., Strube, A., Vogel, A.: Catwalk: a quick development path for performance models. In: *Euro-Par 2014: Parallel Processing Workshops. Lecture Notes in Computer Science*, pp. 589–600. Springer, Cham (2014)
41. Wu, X., Mueller, F.: ScalaExtrap: trace-based communication extrapolation for SPMD programs. In: *Proceedings of the 16th ACM Symposium on Principles and Practice of Parallel Programming (PPoPP '11)*, pp. 113–122. ACM, New York (2011)
42. Zhai, J., Chen, W., Zheng, W.: Phantom: predicting performance of parallel applications on large-scale parallel machines using a single node. *Sigplan Not.* **45**(5), 305–314 (2010)