FRIEDRICH-ALEXANDER-UNIVERSITÄT ERLANGEN-NÜRNBERG
INSTITUT FÜR INFORMATIK (MATHEMATISCHE MASCHINEN UND DATENVERARBEITUNG)

Lehrstuhl für Informatik 10 (Systemsimulation)

Efficient hierarchical grid coarsening for the open Poisson problem

Daniel Ritter, Ulrich Rüde

Lehrstuhlbericht 09–6

Abstract

The following document describes the progress made in dealing with Poisson's equation on unbound domains. The main focus is placed on interface treatment in the tradition of Stephen McCormick's Fast Adaptive Composite Grid method. The method itself is described in detail and numerical tests are performed. Results of these experiments are extensively shown and analyzed in terms of convergence rates and error norms. The report shows also a way to derive the FAC from a standard approach of a totally refined grid. The focus of the report is on graphs and data for further research and publications.

Contents

1 Introduction

Hierachically coarsened grids play a major role in simulation applications in cases where problems are given that need to be solved on a global domain. A wide variety of methods has been developed, starting either from a global coarse grid or from a local fine one to enable hierarchical structures with different grid sizes. One of these approaches is the full adaptive composite grid algorithm (FAC) that was introduced by Steve McCormick in [5]. It is not only a method that is very accurate in mathematical terms, but also can be easily implemented in an efficient way. This is due to the fact that the FAC can be executed with regular stencil operations which operate only on one grid level at a time (except prolongation and restriction).
This report reviews a simple, straightforward test environment that was written in the scripting language *python*. This was chosen due to its vectorization capabilities (similar syntax as Matlab), its free availability and portability, and its object–oriented facilities. For this report only 2D tests were performed (3D will be implemented and tested in the near future). For visualization *gnuplot* was used in batch mode.
The report is structured the following way: In section 2 an introduction to the model problem is given, which covers the governing equation as well as the concept of hierarchical grid coarsening that is applied including some new, convenient naming conventions. In section 3 different test cases for the multigrid solver are described and evaluated: Convergence rates and error norms are shown for different grid sizes and numbers of levels. Finally, an overview over the next planned steps is given in section 6. In the appendix the source files are listed.

2 Model description

2.1 Poisson's equation on unbound domains

From potential theory it is known that the acting forces at position x can be derived from a global potential $\Phi(x)$ that is given by the charged particles, i. e. their positions. The governing equation for Φ is also known as Poisson's equation and is given as

$$\Delta\Phi(x) = -\frac{1}{\varepsilon_0}\rho(x), \tag{1}$$

with the dielectric constant ε_0 and charge density $\rho(x)$ at position $x \in \mathbb{R}^2$. $\Delta = \frac{\partial^2}{\partial x_1^2} + \frac{\partial^2}{\partial x_2^2}$ is the Laplace operator. In case we consider particles as points at positions $x_1, x_2, \dots x_N$, $\rho(x)$ is a superposition of the Dirac impulses $\delta_i(x)$: $\rho(x) = \sum_{i=1}^{N} \delta(x - x_i)$. For numerical simulation this right-hand side is impractical to handle since it is not smooth and may cause large errors. Therefore the impulses are usually approximated by a smoothed function with similar properties, e. g. a gaussian distribution or a spline function. As an alternative, Zenger corrections can be used that do not smooth the charges, but distribute them to the grid points that are direct neighbors of the particle. In [4] these corrections were analyzed exhaustively. From here, sufficient smoothness of the right-hand side is assumed and $f(x)$ will denote the right-hand side.
In molecular dynamics the domains are often considered to be unbound. This property is equivalent to the so-called open boundary conditions. These correspond to the case, when particles are moving within a bounded area but with long-range interactions. The charges of the molecules are inside the area, whereas the induced field is unbounded. That fact has to be considered if we want to avoid unpredictable errors. The overall problem (including the b.c.s) can be formulated as

$$\begin{aligned} &\Delta\Phi(x) = f(x),\ x \in \mathbb{R}^3, \\ &\text{with } \Phi(x) \to 0 \text{ for } \|x\| \to \infty, \end{aligned} \tag{2}$$

where $\text{supp}(f) \subset \Omega$ is a bounded subset of $\mathbb{R}^3$. This is a generalization of zero Dirichlet boundary conditions at infinity. An analog discrete formulation of this problem is given as

$$\begin{aligned} &\Delta_h\Phi(x) = f(x),\ x \in \{x | x = h \cdot z, z \in \mathbb{Z}^3\}, \\ &\text{with } \Phi(x) \to 0 \text{ for } \|x\| \to \infty, \end{aligned} \tag{3}$$

where $h > 0$ is the grid size of the discretization and Δ_h is the discrete Laplace operator. Equation (3) still constitutes an infinite number of equations.
This problem can be overcome using a finite hierarchical coarsening approach as described in [1] and used before by [9]. Different from Bolten, we perform the coarsening without decreasing the number of grid points, i. e. the physical diameter of the domain is doubled from level i to $i + 1$ (what corresponds to $\beta = 2$ in Bolten's notation). This means, however, that we need another method for solving the overall linear system on the coarsest grid, since the size of the system is still the same. For the tests we assume a known, exact solution on the coarsest level.

2.2 The fast adaptive composite grid method for hierarchically coarsened grids

The fast adaptive composite grid (FAC) method was introduced in [5]. Its great advantage over the approaches mentioned before ([1], [9]) is that it enables adaptive hierarchical grids without requiring irregular stencils. Basically, it was introduced for adaptive refinement, but it can be applied just the other way round for hierarchical coarsening: Let $\Omega_0 = \{h \cdot z | z \in \{0, \pm 1, \cdots \pm \frac{(n-1)}{2}\}^2\}$, where n is the number of grid points per dimension and h is the meshwidth. Accordingly, coarsenings of the domain are given as $\Omega_k = \{x = h_k \cdot z, z \in \{0, \pm 1 \cdots \pm \frac{(n-1)}{2}\}^2\}$ with $h_k = 2^k h$. Furthermore, let ${}^{+}\Omega_k$ and ${}^{++}\Omega_k$ denote the domain plus 1 resp. 2 outer layers (called *ghost layers* or *halo*) of grid points and ${}^{-}\Omega_k$ the interior of Ω_k. Let $\Omega_k^{h_{i+1}}$ denote the restriction of Ω_k to grid size h_{k+1}. A two level example is given in figure 1. In chapter 4 of [5] the FAC is introduced with a slightly different notation. The basic idea of it is quite simple: For solving equation equation 3 a hierarchical grid structure is created. The interfaces between different grid levels are easy to handle with the introduced approach.

- Presmoothing (using a red-black Gauss-Seidel smoother) is done on the interior ${}^{-}\Omega_k$ on the according level k,

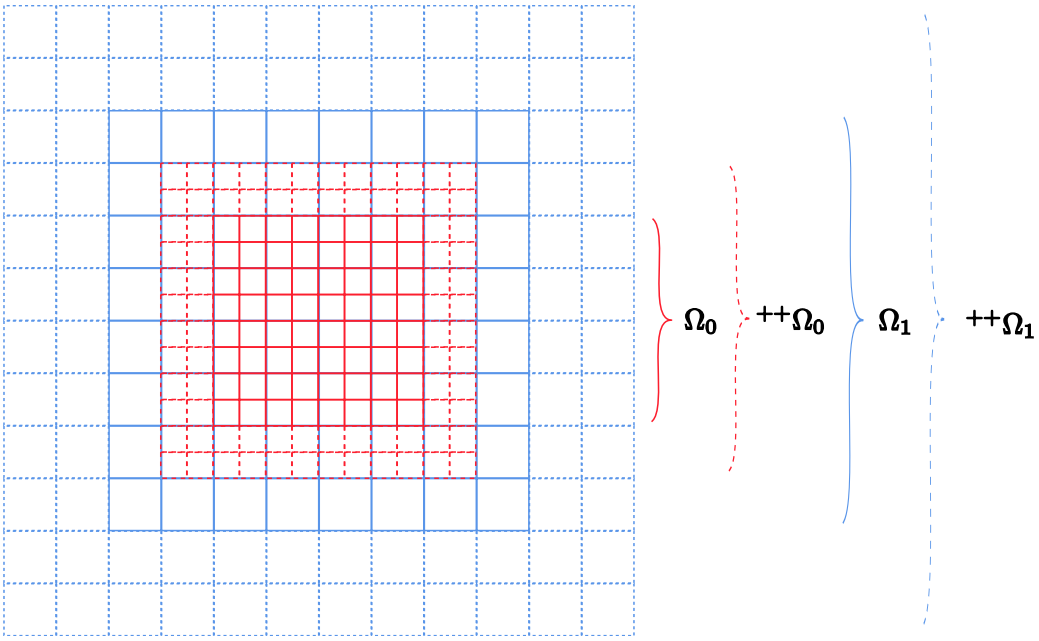

Figure 1: Two level coarsened grid

- the residual is calculated on the domain and the first border, i. e. ${}^{+}\Omega_k$ and
- is transfered to the coarser ${}^{+}\Omega_k^{h_{k+1}}$. Note, that the boundary of the latter is identical to the boundary of Ω_k.
- After solving the problem on the coarse grid, correction terms are prolongated from ${}^{++}\Omega_k^{h_{k+1}}$ to the fine domain and its first and second border ${}^{++}\Omega_k$.
- Finally, postsmoothing is done on the interior ${}^{-}\Omega_k$ again. A full approximation scheme is required for the posed problem since the boundary conditions have to be set on the coarsest level.

Let $\tilde{I}_k^{k+1}$ be the full weighting restriction operator from level k to level $k+1$, $\hat{I}_k^{k+1}$ the according direct injection operator and I_{k+1}^k the bilinear interpolation operator (from level $k+1$ to k). Δ_h is the regular five-point stencil with mesh size h and Δ_h^{-1} its inverse, which can be established for given boundary conditions. Then, a formal definition of the FAC method is given in algorithm 1.

2.3 The FAC as restricted refinement scheme

The FAC can also be seen from another point of view that was developed implicitly by [5] and [8] and was formulated as adaptive relaxation in [7] (but for triangle meshes). Imagine a setup where the domain $\Omega_{k_{\max}}$ is given with mesh size $2^{k_{\max}}h$. The domain is now refined globally, i. e. in the previous notation the grids $\Omega_{k_{\max}}^{2^{k_{\max}-1}h}, \Omega_{k_{\max}}^{2^{k_{\max}-2}h}, \dots \Omega_{k_{\max}}^{h}$ are used for solving the problem. This scheme leads to a classical multigrid scheme with Dirichlet boundary conditions. This approach is referred to as *fully refined scheme*.
Imagine now, that the relaxation (i. e. pre- and postsmoothing) is restricted compared to the previous scheme: On level k it is now only performed on the subdomain that was defined as ${}^{-}\Omega_k$ before. This inner-point relaxation is called *smoothed interior scheme*.
Now it is only one small step to the previously introduced *FAC*: All the remaining operations within the multigrid cycle — computation of the residual, restriction, prolongation and correction — are also only done on the regarding subsets of the domain. Note that the smoothed interior and the FAC schemes are numerically equivalent.
The convergence rates for grid resolutions 33^2, 65^2 and 129^2 are shown in table 1. The convergence rates for iteration i are computed as the quotient $\frac{||r^i||_2}{||r^(i-1)||_2}$ of the L2-Norms of the residual r^i for iteration i on the finest grid. The performance of FAC scheme and smoothed interior scheme is almost identical, while the

Algorithm 1 FAC method

FAC$(k, k_{\max}, \nu_{\text{pre}}, \nu_{\text{post}}, h)$:

1: **if** $(k = k_{\max})$ **then**
2: Calculate u_k exactly.
3: **return**
4: **else**
5: smooth u_k for ν_{pre} times on ${}^{-}\Omega_k$
6: compute residual on ${}^{+}\Omega_k$: $r_k \leftarrow f_k - \Delta_h u_k$
7: restrict residual $r_{k+1} \leftarrow \tilde{I}_k^{k+1} r_k$ for $r_{k+1} \in^{+} \Omega_k^{2h}$
8: restrict $u_{k+1} \leftarrow \begin{cases} \hat{I}_k^{k+1} u_k & u_{k+1} \in^{++} \Omega_k^{2h} \\ u_{k+1} & \text{otherwise} \end{cases}$
9: set $f_{k+1} \leftarrow \begin{cases} r_{k+1} + \Delta_{2h} u_{k+1} & u_{k+1} \in^{++} \Omega_k^{2h} \\ f_{k+1} & \text{otherwise} \end{cases}$
10: call FAC$(k+1, k_{\max}, \nu_{\text{pre}}, \nu_{\text{post}}, 2h)$
11: compute error $e_{k+1} \leftarrow \Delta_{2h}^{-1} f_{k+1} - u_{k+1}$ on Ω_{k+1}
12: prolongate $e_k \leftarrow I_{k+1}^k e_{k+1}$ for $e_k \in^{++} \Omega_{k+1}$
13: correct $u_k \leftarrow u_k + e_k$ on Ω_k
14: smooth u_k on ${}^{-}\Omega_k$ for ν_{post} times
15: **end if**

convergence for the fully refined multi grid is slightly better than for the other cases.
The executed test case will be described just in the next section and is refered to as test 0. It is a noise-based test using so-called *power iteration*. Note that for this case 0 the residual norms are equivalent to the relative errors, i. e. the errors induced by the interfaces are considerably small.

3 Description of numerical tests

3.1 Test setups

For testing the quality of the approximate solution two-, three- and four-level schemes were used. Exact boundary conditions were set on the coarsest grid level each and the FAC was used to get an approximate solution on the finest one. The following general setups were tested:

0. Gaussian noise in initial $u_h(x)$ with boundary conditions $u_h(x_\delta) = 0$ for boundary points $x_\delta \in \delta\Omega_{k_{\max}}$[1] and right-hand side $f(x) = 0$. The exact solution is $u(x) = 0$ and the approximate solution converges against zero fast. Therefore, only a few V-cycles are possible, before round-off errors dominate the residual norms. One way to prevent that is the so-called "Power iteration", i. e. the elements of the approximate solution vector are divided by the L2 norm of the error vector after each V-cycle.

1. Impose the boundaries $u_h(x_\delta) = c$ (constant) and set right-hand side $f_h(x) = 0$. The analytical solution is $u(x) = c$.

2. Impose the boundaries $u_h(x_\delta) = x_1 + x_2$ (bilinear function) and set the right-hand side $f_h(x) = 0$. The analytical solution is $u(x) = x_1 + x_2$.

3. Impose the boundaries $u_h(x_\delta) = x_1^2 - x_2^2$ (biquadratic function) and set right-hand side $f_h(x) = 0$. The analytical solution is $u(x) = x_1^2 + x_2^2$.

Additionally, setups were tested that reflect a typical situation in molecular dynamics: A charged particle is placed at the origin and is modeled in three different ways: It is assumed to be a point charge in test case 4, what leads to a singularity in the solution, whereas it is shielded by a cubic function in test case 5 and by a gaussian function in test case 6:

[1] Define the boundary as $\delta\Omega_{k_{\max}} := \Omega_{k_{\max}} \setminus^{-} \Omega_{k_{\max}}$

(a) Grid resolution 33^2

	Convergence rate of i-th V-Cycle		
i	Fully refined	Smoothed interior	FAC scheme
2	0.014597	0.016352	0.015739
3	0.020572	0.020045	0.020434
4	0.023102	0.022276	0.023039
5	0.024276	0.023473	0.024140
6	0.024865	0.024115	0.024624
7	0.025195	0.024485	0.024860
8	0.025398	0.024712	0.024987
9	0.025532	0.024861	0.025062
10	0.025624	0.024960	0.025108
11	0.025690	0.025030	0.025139
12	0.025738	0.025078	0.025159
13	0.025774	0.025113	0.025173
14	0.025802	0.025137	0.025182
15	0.025823	0.025155	0.025189

(b) Grid resolution 65^2

	Convergence rate of i-th V-Cycle		
i	Fully refined	Smoothed interior	FAC scheme
2	0.015568	0.016208	0.016197
3	0.020282	0.020473	0.020437
4	0.022617	0.022938	0.022888
5	0.023994	0.024267	0.024215
6	0.024849	0.024999	0.024988
7	0.025399	0.025446	0.025474
8	0.025765	0.025753	0.025808
9	0.026018	0.025984	0.026057
10	0.026202	0.026168	0.026251
11	0.026342	0.026321	0.026408
12	0.026452	0.026451	0.026538
13	0.026542	0.026564	0.026647
14	0.026619	0.026663	0.026741
15	0.026686	0.026749	0.026822

(c) Grid resolution 129^2

	Convergence rate of i-th V-Cycle		
i	Fully refined	Smoothed interior	FAC scheme
2	0.018199	0.017298	0.017745
3	0.022172	0.021364	0.021829
4	0.023908	0.023588	0.023766
5	0.024829	0.024788	0.024817
6	0.025381	0.025470	0.025441
7	0.025747	0.025899	0.025846
8	0.026007	0.026199	0.026129
9	0.026204	0.026425	0.026341
10	0.026359	0.026604	0.026506
11	0.026485	0.026752	0.026638
12	0.026589	0.026878	0.026748
13	0.026677	0.026987	0.026841
14	0.026753	0.027083	0.026921
15	0.026820	0.027168	0.026990

Table 1: Convergence rates for different relaxation schemes and sizes

4. Point charge at $x = 0$, approximated with discrete Dirac impulse, right-hand side $f_h(x) = \begin{cases} 0 & x \neq 0 \\ \frac{1}{h^2} & x = 0. \end{cases}$

   Since the analytical solution is unbounded, the convergence properties of the numerical solution are not trivial. However, it holds that for

   $$x \in \Omega \setminus 0 \lim_{h \to 0} |u_h(x) - u(x)| = 0$$

   and furthermore the asymptotic behavior $|u_h(x) - u(x)| = ch^2$ (with c depending on the distance of x to 0, see [3] and [6]).

5. Point charge at $x = 0$, approximated by cubic polynom:

   $$f(x) = \begin{cases} 0 & ||x|| > r_{\text{cut}} \\ \frac{2}{\pi r_{\text{cut}}^4} ||x^3|| - \frac{3}{\pi r_{\text{cut}}^3} ||x^2|| + \frac{1}{\pi r_{\text{cut}}} & \text{else.} \end{cases}.$$

   In that case, we are not aware of an analytical solution for $u(x)$, only the far-field behavior approximates $u_{\text{far}}(x) = -\frac{1}{2\pi} \log ||x||$ again, but the right-hand side is smooth. Therefore, this test case is left out in section 5.

6. Point charge at $x = 0$, approximated by Gaussian function:

   $$f(x) = \begin{cases} 0 & ||x|| > r_{\text{cut}} \\ \frac{4}{r_{\text{cut}}^2} (\frac{||x||}{r_{\text{cut}}} - 1) e^{-\frac{||x||}{r_{\text{cut}}}} & \text{else.} \end{cases}.$$

   In this case the exact solution is given as $u(x) = \frac{1}{r_{\text{cut}}} e^{-\frac{||x||}{r_{\text{cut}}}}$. This solution is also smooth, so that the near field is also approximated with sufficient accuracy.

In test cases 4–6 the parameter r_{cut} was chosen such that $\int_\Omega u(x) d\Omega = 1$, i. e. the charge of the particle was normalized. For the tests, 10 V-cycles were run each. Pre- and postsmoothing is performed by a fixed number red-black Gauss-Seidel sweeps, if not mentioned explicitely, $\nu = \nu_{\text{pre}} = \nu_{\text{post}} = 3$. The coarse grid solution is determined by a sufficient number of these sweeps, i. e. $\nu_{\text{tot}} = n^2$ for domain size n. The restriction is executed by direct injection for the unknown u and by full-weighting for the residual. The prolongation operator is bilinear interpolation. The next subsection shows how we approximate the boundary conditions on the coarse grid.

3.2 Setting the boundary conditions on the coarse grid

An approximation formula for the boundary condition was developed by [2], but this operator needs a summation over all of the N fine grid function values, i. e. it has complexity $\mathcal{O}(N)$ if and only if the number of boundary points is assumed to be constant. However, the complexity constant is quite big, so this computation needed more than 90% of the computation time. Therefore, another cheaper approach was chosen to speed up the algorithm. First, the total charge q_{tot} is calculated as well as the center of charge x_{tot} and then the boundary (far field) points can be calculated using the approximation $u_{\text{far}}(x) = -\frac{q_{\text{tot}}}{2\pi} \log ||x - x_{\text{tot}}||$. An error estimate for that formula and a general setup maybe subject to further research (we assume that it is well suited for a fine grid that is small compared to the coarse one (sufficient number of levels)). For a symmetric setup as in our test cases this approximation is sufficient.

4 Numerical results for convergence rates

Within this section numerical results of the convergence rates are provided for the FAC scheme. The convergence rates are defined the following way: The convergence rate after iteration i is defined as $\frac{||r^i||_2}{||r^{i-1}||_2}$ with the discrete L2 Norm $||r||_2 := \frac{1}{nm} \sqrt{\sum_{j=1}^{n} \sum_{k=1}^{m} r_{j,k}^2}$ for a $n \times m$ domain. Analogous, we refer to the convergence rate in terms of the maximum norm, which is defined as $||r||_\infty := \max_{1 \leq j \leq n, 1 \leq k \leq m} |r_{j,k}|$.

4.1 Convergence rates for test case 0

Test Case 0 was executed as described above for two grid levels with three pre- and postsmoothing sweeps. The convergence rates after each iteration are displayed in figure 2 for different domain sizes.

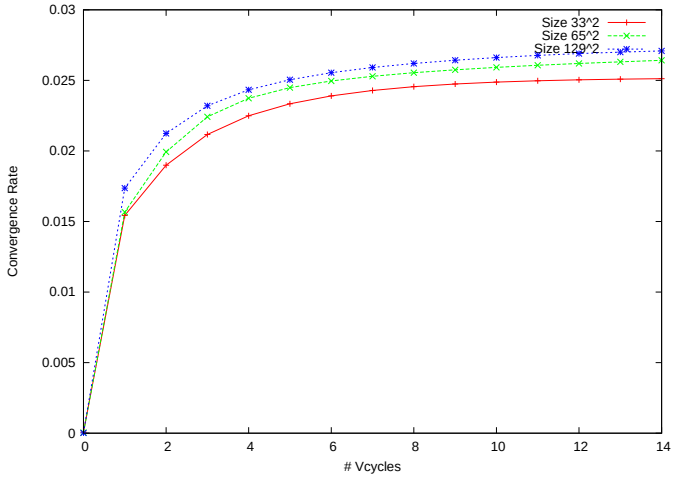

Figure 2: Convergence rates over iterations for test case 0 and two grid levels

One can see that the convergence rate stabilizes after a few iterations and that it is independent of the problem size. The asymptotic value for the convergence rate is about 0.025. This convergence rate should also be observable in the other test cases.

4.2 Convergence rates for test case 4

Within this overview subsection only test case 4 will be of interest. The convergence rates for this test case are shown in table 2. The convergence is given for two-, three- and four level schemes.

	Two-level scheme			Three-level scheme			Four-level scheme		
Size	33^2	65^2	129^2	33^2	65^2	129^2	33^2	65^2	129^2
It. 2	0.0164	0.0164	0.0164	0.0230	0.0231	0.0231	0.0248	0.0250	0.0250
3	0.0210	0.0210	0.0210	0.0276	0.0282	0.0282	0.0290	0.0301	0.0302
4	0.0231	0.0233	0.0233	0.0295	0.0311	0.0311	0.0310	0.0331	0.0332
5	0.0242	0.0246	0.0246	0.0304	0.0329	0.0329	0.0319	0.0350	0.0351
6	0.0314	0.0254	0.0254	0.0378	0.0340	0.0341	0.0380	0.0367	0.0364
7	0.1201	0.0301	0.0261	0.1469	0.0366	0.0349	0.1236	0.0543	0.0373
8	0.1809	0.1929	0.0479	0.4034	0.1344	0.0377	0.2629	0.2734	0.0463
9	0.1739	0.3559	0.2863	0.4142	0.3729	0.1255	0.3314	0.3806	0.2193

Table 2: Overview convergence rates (w.r.t. the L2 norm of the residual) in test case 4

The convergence shows almost the expected behavior: After an initial phase it reaches rates of about 0.024 for two grids (0.035 for three and four grids) before it slows down because round-off errors dominate. The convergence rate is identical to that in the previous subsection. What is unexpected is the fact that the round-off errors come into play not at one point, but gradually. The results will be shown graphically and in more detail within the next three subsections in order to investigate where they result from.

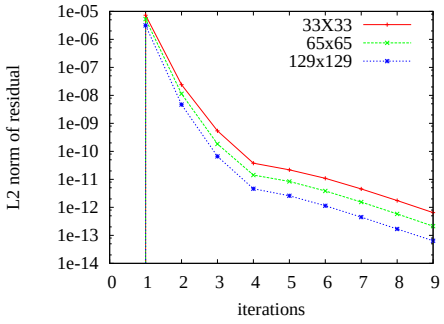

(a) Test case 1

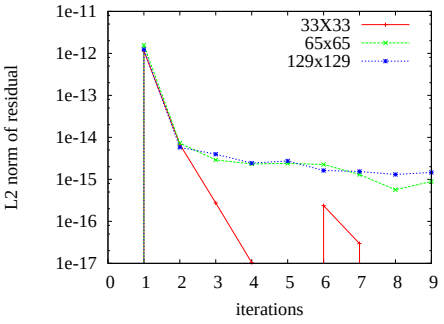

(b) Test case 2

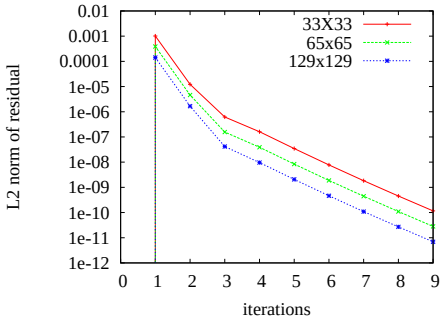

(c) Test case 3

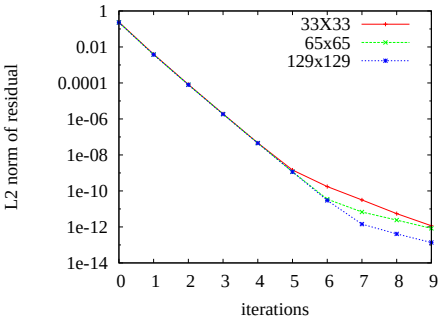

(d) Test case 4

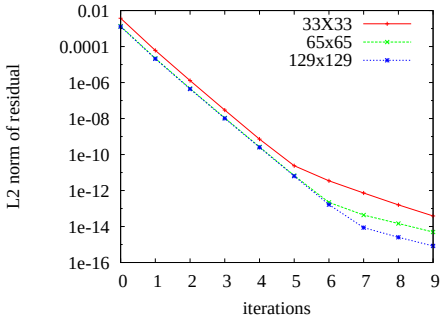

(e) Test case 5

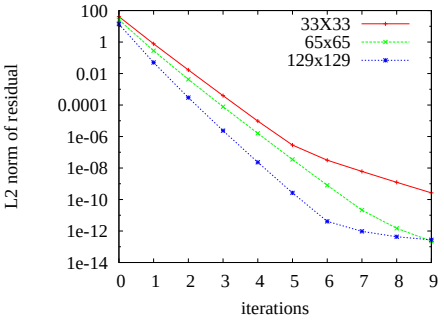

(f) Test case 6

Figure 3: Residuals norms over iterations for a two-level grid

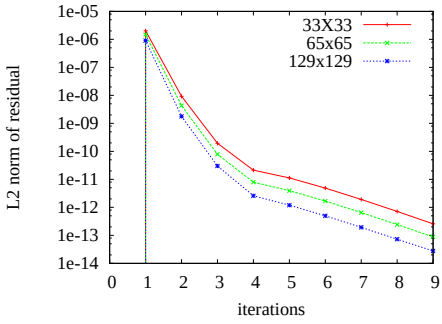

(a) Test case 1

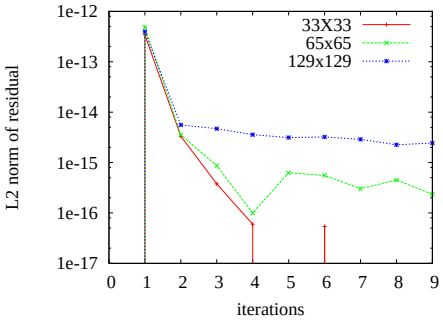

(b) Test case 2

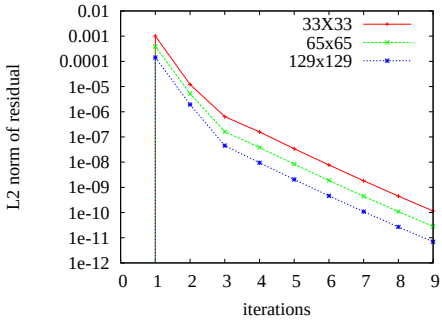

(c) Test case 3

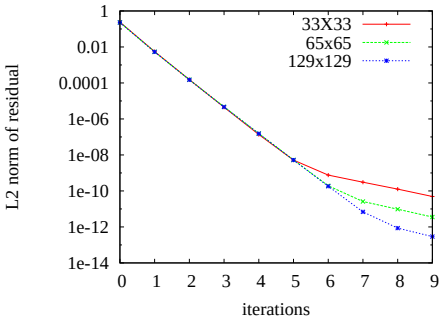

(d) Test case 4

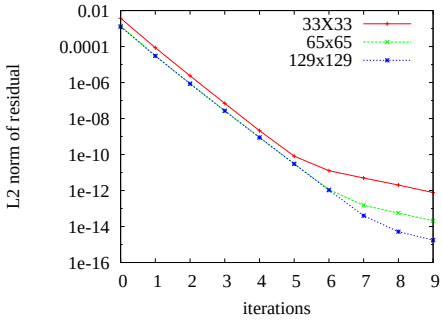

(e) Test case 5

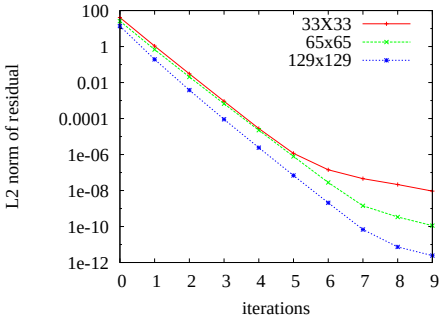

(f) Test case 6

Figure 4: Residuals norms over iterations for three-level grid

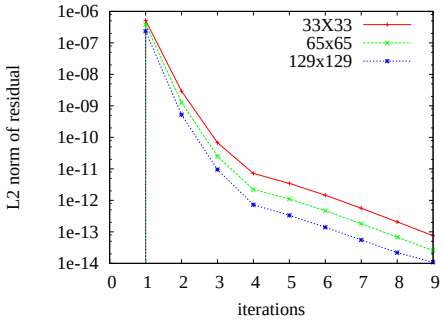

(a) Test case 1

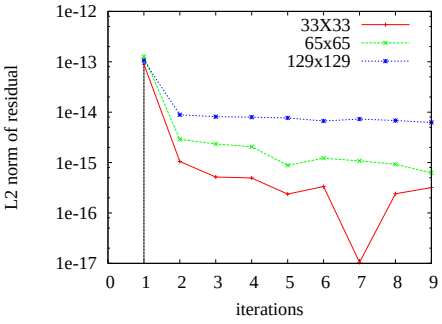

(b) Test case 2

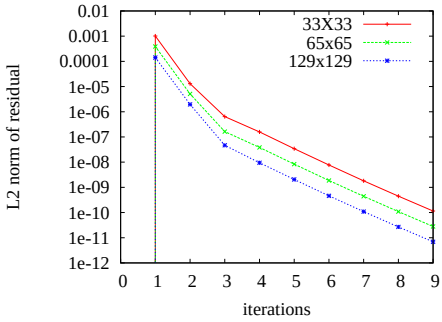

(c) Test case 3

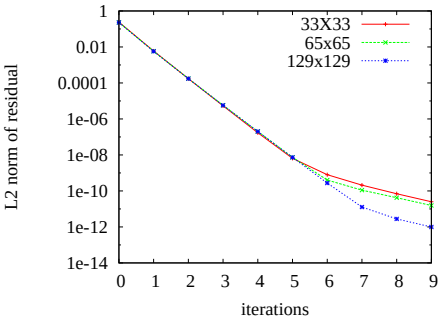

(d) Test case 4

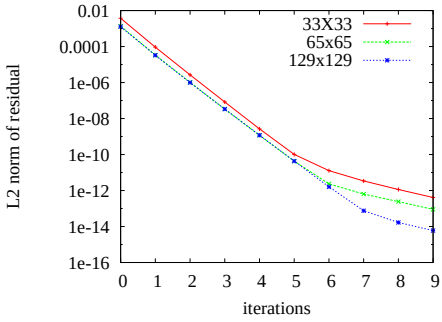

(e) Test case 5

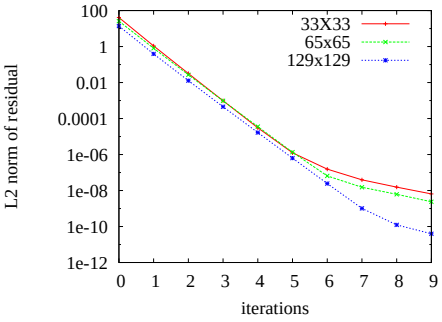

(f) Test case 6

Figure 5: Residuals norms over iterations for four-level grid

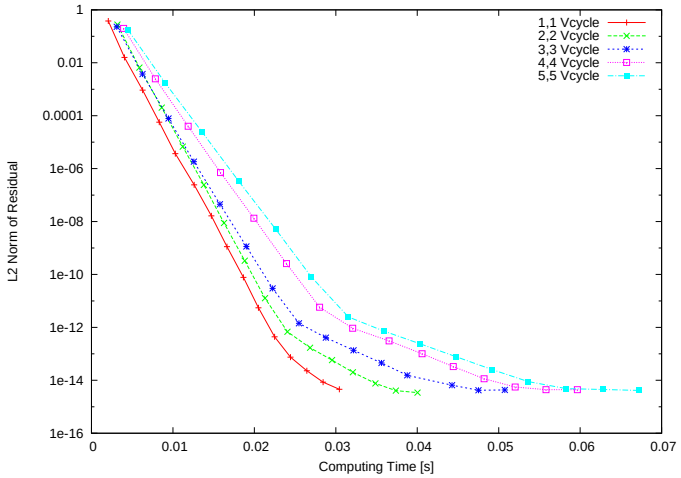

Figure 6: Computational time for a 129^2 grid and different numbers of pre- and postsmoothing.

4.3 Detailed results for two-, three- and four-level schemes

In the following the convergence rates, i. e. the L2 norms of the residuals are shown for 10 V-cycles each. Test runs were executed for all the test cases, with different spatial resolutions and a varying number of grid levels. The results are in total consistent with the ones in the previous subsection. There are several interesting developments shown in the graphs below that need to be explained:

1. In the test cases 1–3 the residual is zero after the first V-cycle since the initial guess solves the governing equation (without the b.c.s). After the first V-cycle the residual that is induced by the b.c.s spreads over the whole domain. This leads to a jump in the residual norm.

2. The norm of the residuals decreases by a factor of roughly 0.02 in the first few iterations. Then it is slowed down to factors of around 0.2.

3. The norm of the residual can be decreased to 10^{-16} – 10^{-14} before round-off errors dominate and it stops decreasing.

To draw a conclusion about the convergence rate: The convergence rate was observed to be independent of the problem size, the test case and the number of grid levels. For the number of pre- and post-smoothing sweeps being 3 each it is around 0.025. In the next subsection the results will be shown for a varying number of smoothing sweeps.

4.4 Residual against computational time

Besides the previous test cases, an experiment was done with different numbers of pre- and postsmoothing sweeps. The idea of this is that an increased number of smoothing steps on each level is leading to a decreased residual. On the other hand, the smoother is dominating the computation time, which is therefore increasing with the number of smoothing sweeps. In figure 6 the L2-Norm of the residual is plotted against the runtime for $\nu_{\text{pre}} = \nu_{\text{post}} \in \{1, 2, 3, 4, 5\}$ and test case 4. The test run shows that the efficiency according to this measure is best for one pre- and postsmoothing step, a larger number of smoothing steps is less efficient. (This measure, however, only is valid for the python code and might be very different on other, more sophisticated implementations.)
In table 3 the convergence rates for the different smoothing sweeps are displayed. These results apply for test case 0 with resolution of 129^2 . The convergence is faster for more smoothing sweeps of course.

	Number of smoothing sweeps ν				
	1	2	3	4	5
Conv. rate	0.0709	0.0395	0.027	0.0208	0.0167

Table 3: Convergence rates dependent on $\nu = \nu_{\text{pre}} = \nu_{\text{post}}$

5 Numerical results for errors

Besides the residual, the error is an important performance measure. For test cases 1, 2, 3, 4 and 6 analytical continuous solutions are known and were compared to the numerical, discrete one. In the first subsection the L2 and maximum norms for different numbers of grid levels and test cases are shown, further below graphical results for all test cases are provided. The error results from the discretization of the governing equation, from inexact boundary conditions and from the interfaces between fine and coarse grid.

5.1 Overview

A first impression of the L2 and maximum error norms is given in table 4 and table 5. The errors were computed for test case 6 here. The plots show different phenomena: after three V-cycles, the error norms stay constant. However, as shown in the previous subsection the residual continues to converge. This means that the discretization error is already reached after three V-cycles. A surprising effect is that the L2 error norms are smallest after the first cycle in some cases. This holds also for the maximum norms and we currently have no explanation for this behavior.
If the error order is investigated it shows in all cases that the error quarters roughly if the meshwidth is doubled, i. e. an error order of $\mathcal{O}(h^2)$ is induced. In some cases the error is getting even smaller.
The influence of the number of grid levels is not so strong. This may come from the good interface approximation of the FAC method and also from the fact that the influence of the boundary values is neglectable. The error is slightly bigger for more grid levels what may also result from a small error that is added by each of the interfaces.

	Two-level scheme			Three-level scheme			Four-level scheme		
Size	33	65	129	33	65	129	33	65	129
It. 1	0.00227	0.000513	7.81e-05	0.00565	0.00213	0.000439	0.00706	0.00359	0.0012
2	0.00259	0.000311	3.93e-05	0.0028	0.000277	3.63e-05	0.00252	0.000409	3.51e-05
3	0.00265	0.000318	3.96e-05	0.00292	0.00032	4.09e-05	0.00265	0.000458	4.55e-05
4	0.00265	0.000318	3.96e-05	0.00292	0.000321	4.09e-05	0.00265	0.00046	4.6e-05
5	0.00265	0.000318	3.96e-05	0.00292	0.000321	4.09e-05	0.00265	0.00046	4.61e-05
6	0.00265	0.000318	3.96e-05	0.00292	0.000321	4.09e-05	0.00265	0.00046	4.61e-05
7	0.00265	0.000318	3.96e-05	0.00292	0.000321	4.09e-05	0.00265	0.00046	4.61e-05
8	0.00265	0.000318	3.96e-05	0.00292	0.000321	4.09e-05	0.00265	0.00046	4.61e-05
9	0.00265	0.000318	3.96e-05	0.00292	0.000321	4.09e-05	0.00265	0.00046	4.61e-05
10	0.00265	0.000318	3.96e-05	0.00292	0.000321	4.09e-05	0.00265	0.00046	4.61e-05

Table 4: Overview error L2 norms for test case 6

Within the following subsections, graphs of the error will be shown for different test cases, meshwidths and numbers of grid levels. These shall just illustrate the behavior of the error. Due to the symmetrical nature of the problem also the errors are symmetric, so only one quarte of the domain is plotted. Since there is a singularity at the origin in test case 4 only the outer part of the domain is shown in the according plots.

5.2 Results for a two-level scheme

Figure 7, figure 8 and figure 9 show the errors for test cases 1–3 (i. e. the ones with zero right-hand side). The errors were drawn after 10 V-cycles. In test case 1, the error is almost constant over the whole domain. This is plausible since the solution of the potential equation is also constant there. The error is largest for the

	Two-level scheme			Three-level scheme			Four-level scheme		
Size	33	65	129	33	65	129	33	65	129
It. 1	0.195	0.199	0.068	0.59	0.803	0.394	0.741	1.22	0.968
2	0.606	0.15	0.0387	0.552	0.127	0.0335	0.579	0.0919	0.00967
3	0.62	0.154	0.0391	0.582	0.15	0.0397	0.618	0.132	0.0343
4	0.62	0.154	0.0391	0.583	0.15	0.0398	0.619	0.133	0.0351
5	0.62	0.154	0.0391	0.583	0.15	0.0398	0.619	0.133	0.0351
6	0.62	0.154	0.0391	0.583	0.15	0.0398	0.619	0.133	0.0351
7	0.62	0.154	0.0391	0.583	0.15	0.0398	0.619	0.133	0.0351
8	0.62	0.154	0.0391	0.583	0.15	0.0398	0.619	0.133	0.0351
9	0.62	0.154	0.0391	0.583	0.15	0.0398	0.619	0.133	0.0351
10	0.62	0.154	0.0391	0.583	0.15	0.0398	0.619	0.133	0.0351

Table 5: Overview error maximum norms for test case 6

largest grid size and smallest for the smallest one as it was expected. The error is growing slightly towards the boundary.

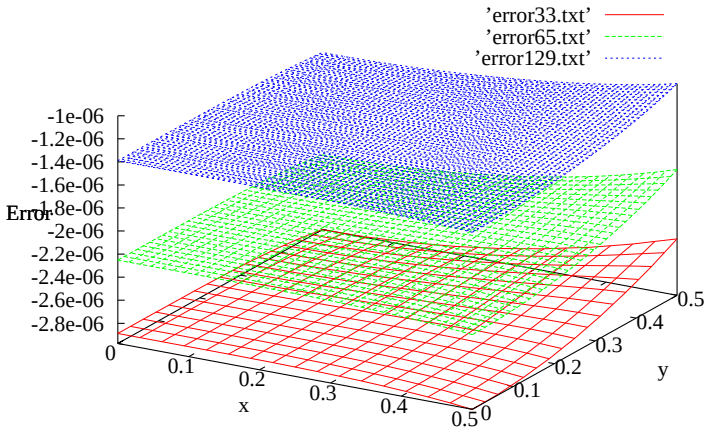

Figure 7: Error for test case 1 and different resolutions (two levels)

In test case 2 the error is zero at the center point and increases bilinearily in direction of the interfaces. Since the solution itself is a bilinear function the error is increasing with the absolute value of the function. The error is here largest for the intermediate resolution of 64^2. Compared to the previous test case the error is smaller by 5–6 orders of magnitude.

Test case 3 shows a different behavior: Here the interface effects dominate by far. This means that the error is oscillating with high frequency at the boundary points. The error is largest for the lowest resolution and smallest for the highest as expected. It can also be seen nicely that the error is jumping with high frequency independent of the grid size.

Figure 10 shows the error for test case 4. Again the error is decreasing with increased resolution and there is a high-frequency component near the interfaces. The center is not displayed since there is a sigularity and the error is unbounded there.

Figure 11 displays test case 6 where the solution is a Gaussian distribution. The error near the center is much larger, since the magnitude of the right-hand side is large there. The charge is placed at $(0., 0.0)$, the boundary of the fine grid is defined by the corner points $(\pm 0.5, \pm 0.5)$.

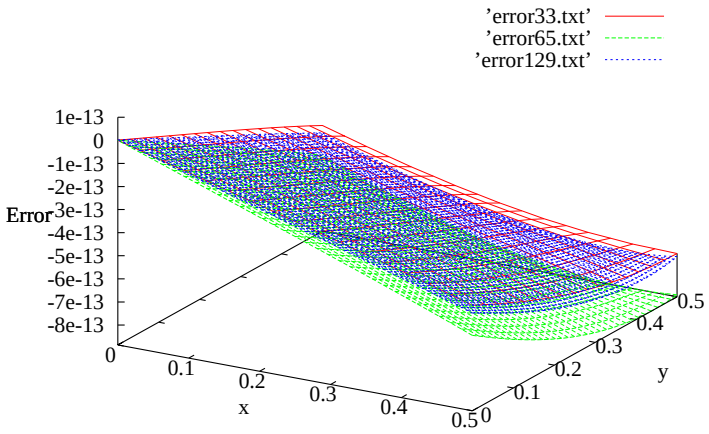

Figure 8: Error for test case 2 and different resolutions (two levels)

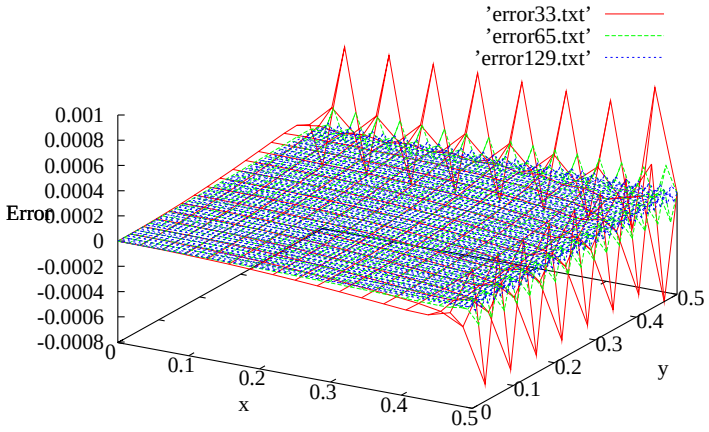

Figure 9: Error for test case 3 and different resolutions (two levels)

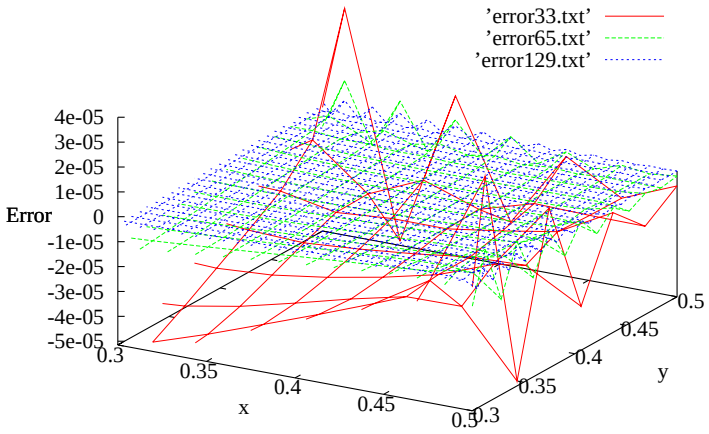

Figure 10: Error for test case 4 and different resolutions (two levels)

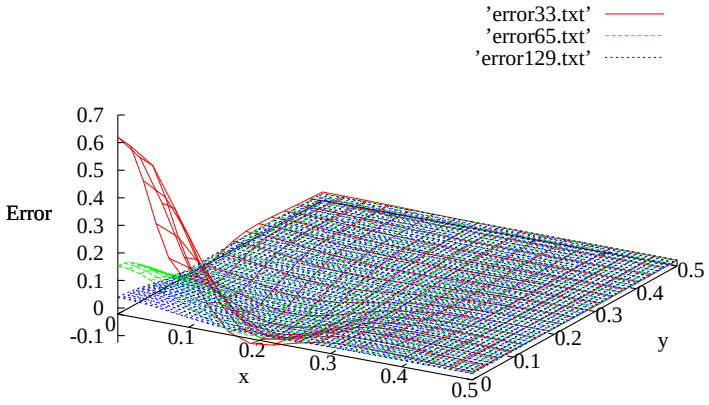

Figure 11: Error for test case 6 and different resolutions (two levels)

Figure 10–11 imply that the error declines with increasing resolution. Besides the large errors near the center (due to the influence of the right-hand side there), the errors are also large at the interfaces.
In total, the experiments demonstrate that the errors decrease with increased grid size. In quantitative terms, they also show the expected $O(h^2)$ behavior, i. e. they are quartered if h is halved.

5.3 Results for a three-level scheme

The error norms are also analyzed for the three-level grid scheme. Analogous to the previous subsection, the errors are plotted in figures 12–16.

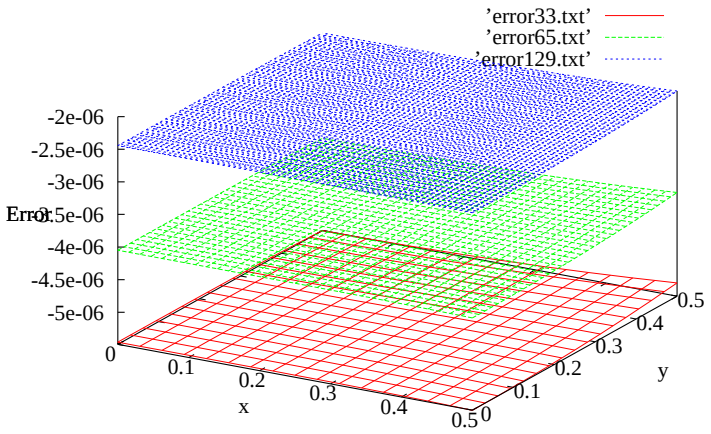

Figure 12: Error for test case 1 and different resolutions (three levels)

The tests show a very similar behavior as for the two-level scheme. The perimeter of the whole domain does not play an important role. The error is slightly larger compared to the two-level grid. This results from an additional interface that has to be treated and introduces an additional error component.

5.4 Results for a four-level scheme

Within this subsection the errors for the four-level tests are shown. In principle, those results are very similar to the previous ones. Again, the errors are slightly larger for four grid levels than for three ones, since another interface has to be considered in the four-level case.

5.5 Comparison of errors for different levels

To give a short comparison of the errors depending on the number of grids used for computation, these are displayed for two three and four grids for a resolution of 65^2 and all the different test cases in figure 22, 23 and 24. Again, we can see that the error is slightly increasing with the number of grid levels.

6 Conclusion and next steps

Within this report a variant of McCormick's FAC to Poisson's equation on a hierarchically coarsened grids with free boundary conditions was introduced. The focus of application was molecular dynamics simulation. The algorithm was implemented as a python program for 2D in different variations. The program proved that the FAC is equivalent to a classic multigrid algorithm where pre- and postsmoothing is restricted to the inner part of the domain and that the introduced additional error is relatively small. It could also be shown

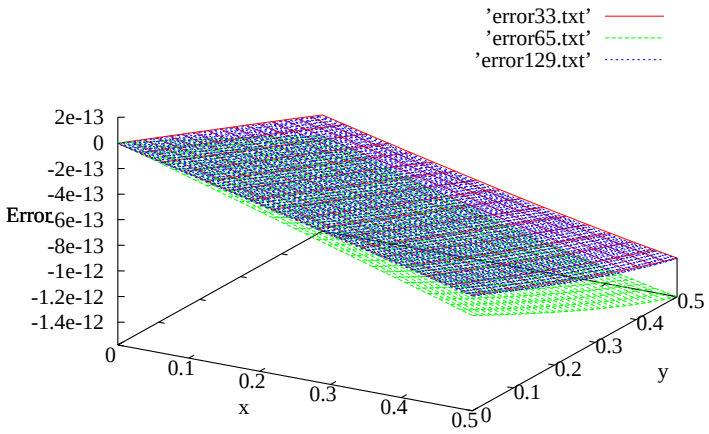

Figure 13: Error for test case 2 and different resolutions (three levels)

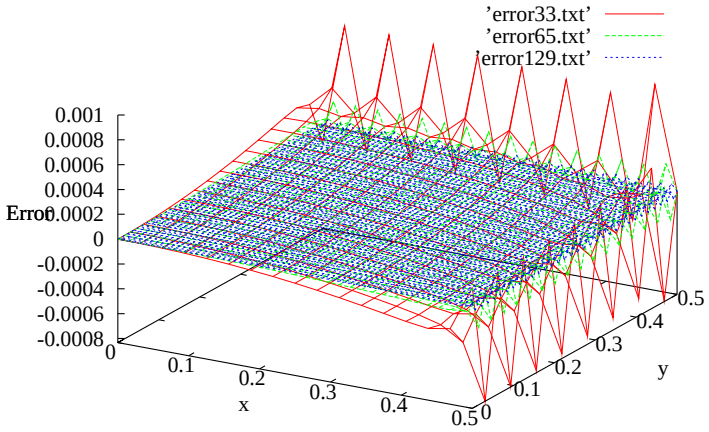

Figure 14: Error for test case 3 and different resolutions (three levels)

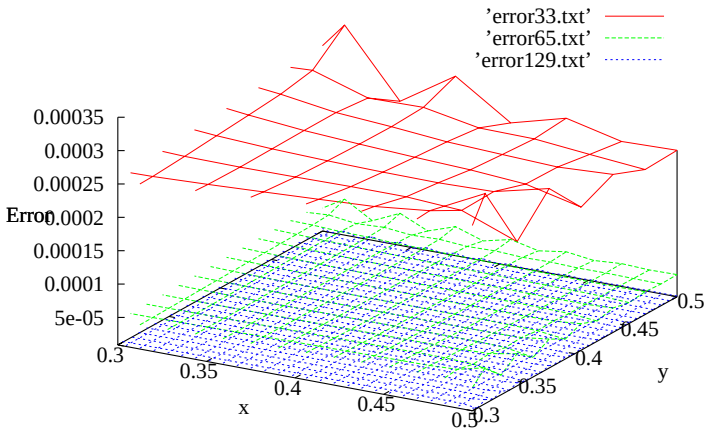

Figure 15: Error for test case 4 and different resolutions (three levels)

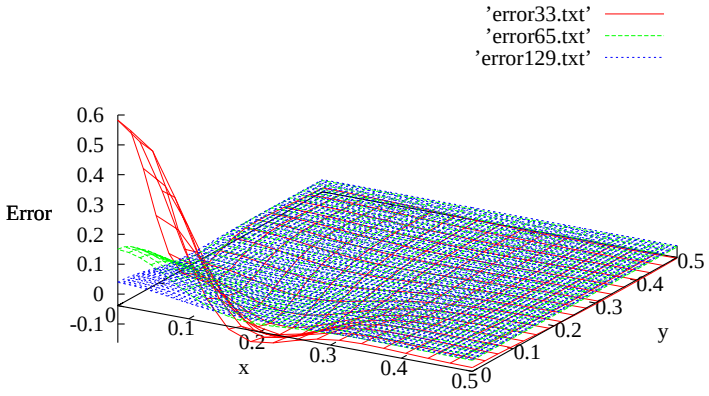

Figure 16: Error for test case 6 and different resolutions (three levels)

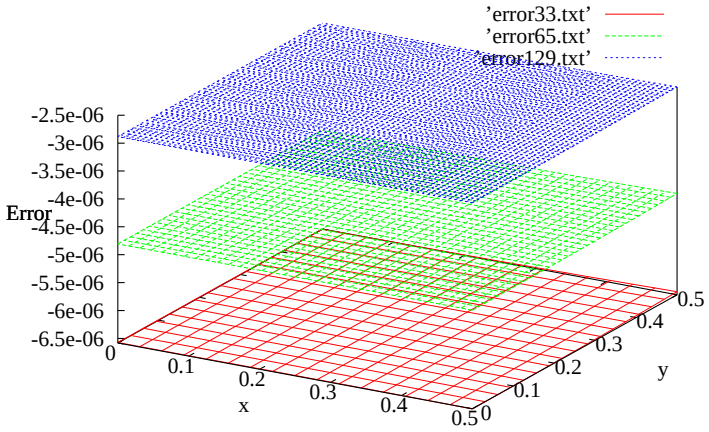

Figure 17: Error for test case 1 and different resolutions (four levels)

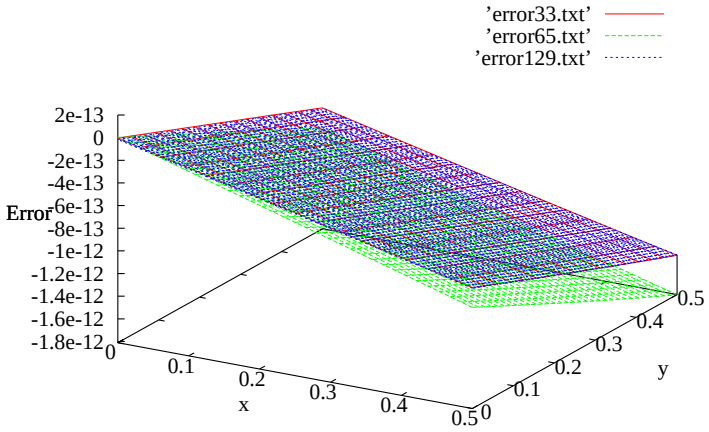

Figure 18: Error for test case 2 and different resolutions (four levels)

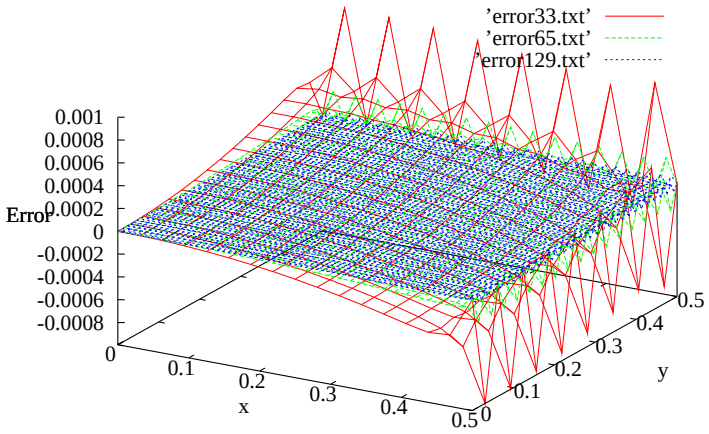

Figure 19: Error for test case 3 and different resolutions (four levels)

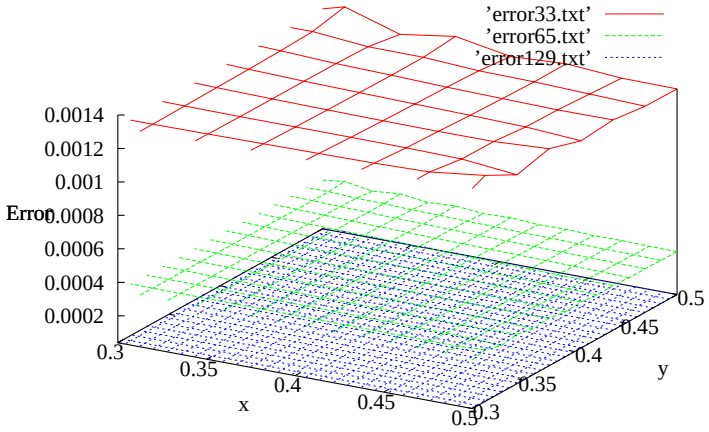

Figure 20: Error for test case 4 and different resolutions (four levels)

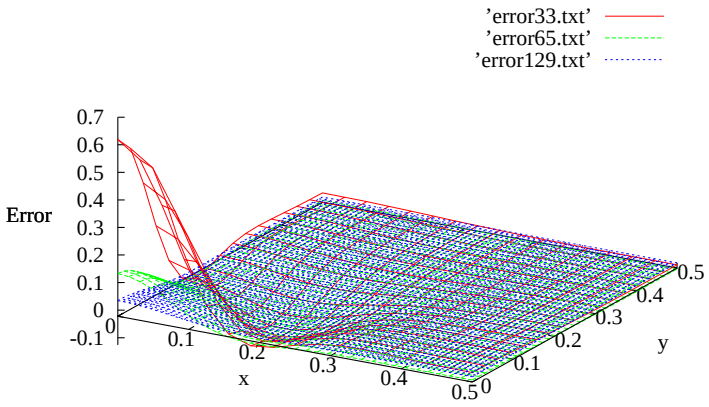

Figure 21: Error for test case 6 and different resolutions (four levels)

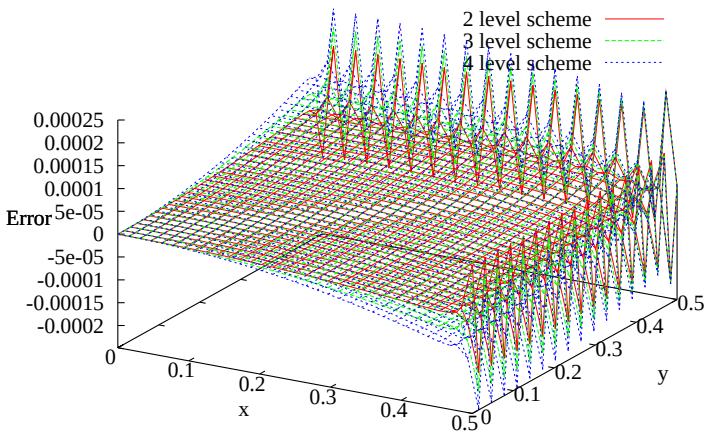

Figure 22: Errors for test case 3 for 2, 3 and 4 grid levels, resolution 65^2

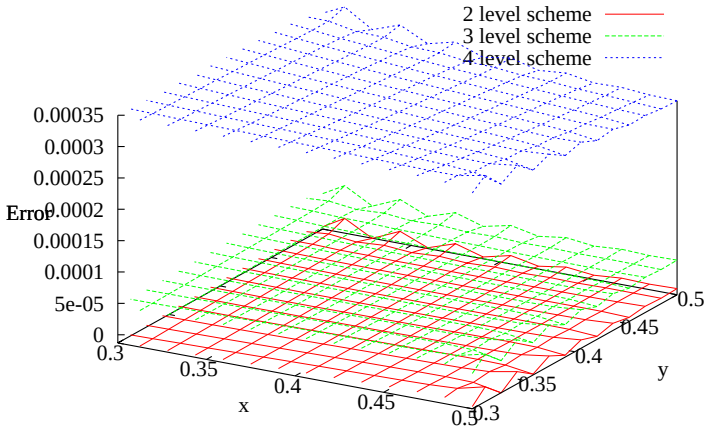

Figure 23: Errors for test case 4 for 2, 3 and 4 grid levels, resolution 65^2

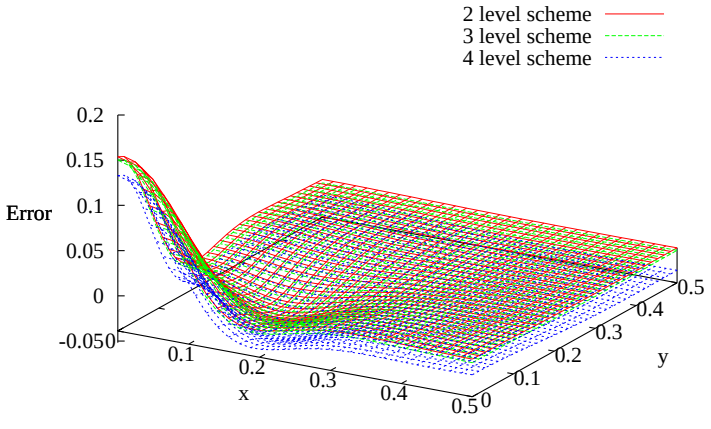

Figure 24: Errors for test case 6 for 2, 3 and 4 grid levels, resolution 65^2

in experiments that the method has overall error order $\mathcal{O}(h^2)$ for meshwidth h. Also the convergence rate only changes very slightly if local refinement is done.
The influence of the boundary conditions was also of interest and it turned out that it is comparably small for molecular dynamics, what means that approximation formulas (or even Dirichlet zero conditions) could be applied instead of an exact integral form. The prove for this is however still pending.
The influence of the number of grid levels to the quality of the solution is smaller than we expected before. This implies that the interface discretization is quite accurate since more interface levels introduce only small errors that have local effects only.
For the future there are some areas of research left:

1. The experiments we performed shall be repeated in 3D to find out whether the same behavior shows there.
2. The implementation shall be integrated in a more performant framework, possibly written in C++.
3. The solver shall be parallelized efficiently.
4. The FAC method shall be generalized to higher error orders.
5. A more general problem, i. e. a governing equation with a jumping dielectric coefficient shall be solved.
6. Other boundary conditions, especially periodic ones shall be integrated in the solver.

References

[1] M. Bolten. Hierarchical grid coarsening for the solution of the poisson equation in free space. *Electronic Transactions on Numerical Analysis*, 29:70–80, 2008.

[2] R. H. Burkhart. Asymptotic expansion of the free-space green's function for the discrete 3-d poisson equation. *SIAM Journal on Scientific Computing*, 18(4):1142–1162, 1997.

[3] H. Köstler. Akkurate Behandlung von Singularitäten bei partiellen Differentialgleichungen. Master's thesis, 2003.

[4] H. Köstler. An accurate multigrid solver for computing singular solutions of elliptic problems. In *Abstracts Of the 12th Copper Mountain Conference on Multigrid Methods*, pages 1–11. SIAM, SIAM, Apr 2005.

[5] Stephen F. McCormick. *Multilevel Adaptive Methods for Partial Differential Equations*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1989.

[6] U. Rüde. On the Accurate Computation of Singular Solutions of Laplace's and Poisson's Equation. In S. F. McCormick, editor, *Multigrid Methods: Theory, Applications, Supercomputing: Proceedings of the Third Copper Mountain Conference on Multigrid Methods, April 5-10, 1987*, 1988.

[7] U. Rüde. Fully Adaptive Multigrid Methods. 30(1):230–248, February 1993.

[8] U. Rüde. *Mathematical and computational techniques for multilevel adaptive methods*, volume 13 of *Frontiers in Applied Mathematics*. 1993.

[9] T. Washio and C. Oosterlee. Error analysis for a potential problem on locally refined grids. *Numerische Mathematik*, 86(3):539–563, 2000.