Towards a Physics-friendly Algebraic Multigrid Method for Systems of PDEs

Peter Thum[†], **Tanja Clees**[†]

†*Fraunhofer Institute SCAI, Schloss Birlinghoven, 53754 St. Augustin, Germany*

Abstract

Algebraic multigrid methods (AMG) offer an efficient solution technique for many industrial applications. However, the employment of AMG for very complex applications which deal with convection, diffusion, migration and reaction terms is rather new. This paper focuses on developing a reordering framework which is used to improve the physical-friendliness of the employed AMG approach. The presented method is predicated on point-based AMG and incorporates level-dependent alternating smoothing. Numerical results are presented for a range of electrochemical test cases with scientific and industrial relevance. The results indicate the usefulness of basing the AMG components upon the physics of the underlying system.

1 Introduction

In today's industrial processes, simulation plays an increasingly important role. Due to the high computational complexity of traditional linear solvers, optimal solvers which provide linear complexity, are necessary to run the simulations in an appropriate time. Algebraic multigrid methods promise optimality and are additionally very robust and very simple to plug into existing numerical simulation software.

In the past decades, AMG for systems of PDEs was applied with great success to many areas, in particular, in the area of fluid dynamics. However, the applicability of AMG for very complex applications which deal with convection, diffusion, migration and reaction terms is rather new. An interesting example for such applications is the electrochemical simulation.

We focus on the so-called **M**ulti-**i**on **T**ransport and **Re**action **M**odel (MiTReM) which is employed in electrochemical plating and etching processes. We want to investigate the applicability of point-based algebraic multigrid (PAMG) to MiTReM, see e.g. [1] for details on PAMG. PAMG was developed in the context of device simulation, especially, to solve drift-diffusion systems and has also been applied with great success in the context of Navier-Stokes simulation, see e.g. [4]. Hence, it is obvious to transfer the experience of PAMG with Navier-Stokes and drift-diffusion systems to the cognate field of electro-chemistry which, in principle, combines the convection of Navier-Stokes and the migration/drift of device simulation.

We present a reordering framework which makes it possible to use different measures for strength of connectivity and derived point orderings within the smoothing process. The idea of alternating smoothing, often found in geometric multigrid methods, see e.g. [12], are thereby revitalized. Especially, the fact that the convection and migration propagate into different directions, shall be addressed by this smoothing approach. The development of measures and orderings shall be inspired by physics and shall also be used in the development of coarse level correction.

The gained knowledge from the reordering framework shall also help to improve the PAMG solution process in many areas such as device simulation, electrophoresis, and reactor simulation.

The remainder is organized as follows. Section 2 gives a brief description of the underlying mathematical problem. Section 3 then introduces the solution approaches which are employed for the numerical experiments described in Section 4. The paper closes with Section 5 where the results are summarized and conclusions are given, as well as an outlook for further research.

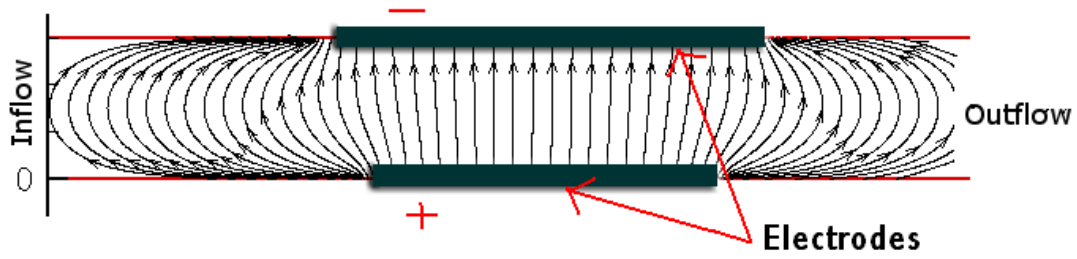

Figure 1: Illustration of migration field (black arrows). The convection field proceeds from inflow to outflow

2 Mathematical Model

The considered electrochemical plating and etching processes make use of dilute ionic solutions in a container bounded by electrodes, membranes and insulators, see Figure 1. The solution is flowing at steady state. A current is applied between the electrodes. By the current the ions in the solution are influenced and attracted or rejected by the electrodes. We consider two different types of reactions. Homogeneous reactions occur between different ions, these reactions are typically reversible. The second type, the so-called heterogeneous reactions, take place at the electrodes. The electrode reactions are captured by the boundary conditions. Due to the fact that not all solved ions react on the electrodes, we have different boundary conditions for different ions on the electrodes.

The resulting PDE system is given by

$$\frac{\partial c_i}{\partial t} = -\mathrm{div}(\underbrace{-F z_i u_i c_i \nabla\Phi}_{\text{migration}} - \underbrace{D_i \nabla c_i}_{\text{diffusion}} + \underbrace{c_i \mathbf{v}}_{\text{convection}}) + \underbrace{R_i}_{\text{reaction}} \tag{1}$$

$$\Delta\Phi = -\frac{F}{\epsilon_0}\sum_i z_i c_i \tag{2}$$

where Equation 1 denotes the transport equation for ion i and Equation 2 represents the electroneutrality constraint. c_i denotes the concentration of ion i, t the time, F the Faraday constant, ϵ_0 the dielectricity constant, Φ the potential and z_i the load, u_i the mechanical mobility and D_i the diffusion coefficient of ion i. R_i represents the term for the homogeneous reactions, which is also dependent upon the other ions of the system. $\mathbf{v}$ denotes the velocity field which describes the steady state flow and is computed with help of the Navier-Stokes equations.

A detailed description of the model can be found in [7]. It is discretized using a combined residual distribution (see [6]) and finite element approach and linearized applying a Newton method.

3 Solution Approach

The employed PAMG approach was originally developed for drift-diffusion systems which can be found in device simulation. There are two new aspects here. First, we have to deal with convection and reaction terms (at least in another form than in device simulation). Second, migration and convection directions are nearly orthogonal to one another (see Figure 1). Both aspects have a similar impact on the character of the linear systems, which makes it difficult to configure the components of PAMG in an optimal way.

It is well known that the ordering of points has a strong influence on the convergence of iterative methods. In the case of MiTReM it is not clear what a sufficiently good ordering should look

like. Especially, the fact that convection and migration propagate in nearly orthogonal directions makes a decision for one ordering very difficult. We will develop an ordering framework which can easily be adjusted to tackle various physical effects and analyze the influence of the resulting orderings on the convergence rate of PAMG. In particular, we will investigate the benefit of an alternating smoother employing "online" reordering on all levels.

In the following, we want to distinguish between the expressions points, unknowns and variables, see [1]. Points denote geometric positions where the values of the solution of our linear system are computed. Here, Points are identified with the grid points. Unknown functions (in the following referenced as unknowns) represent physical functions to be solved for; in the case of the multi-ion model the unknowns are the concentrations and the potential. Finally, we denote with variables the entries of the solution vector.

3.1 Point-Based Algebraic Multigrid (PAMG)

Algebraic multigrid methods (AMG) have their roots in the geometric multigrid, thus, we speak of grids rather than levels, which would be more related to the algebraic context. This is due to the fact that AMG methods commonly do not rely on the grids of the underlying discretized system. They only work with the resulting linear system and typically use no extra information.

The employed point-based AMG has been developed in [1] and is based on the variable-based AMG (VAMG) described by Stüben in [10]. In the following sections we will only introduce the main components of the employed PAMG method and their relation to components of VAMG.

Point-based algebraic multigrid was, in particular, developed for strong coupled systems of PDEs which, among other things, also arise when solving the multi-ion model. It is based on the observation, that for many PDE systems the different unknowns are discretized on (principally) the same grid, so that it appears to be quite natural to create the same hierarchy for all unknowns. PAMG sticks to the concept of creating an approach, which does not need any additional information, besides naturally available ones. To be more specific, it only makes use of matrix entries and the variable-unknown and variable-point mappings.

PAMG uses couplings between points, which are represented by so-called point-coupling matrices. Each of these matrices is then assigned a scalar value. The scalar values represent the so-called primary matrix, which shall sufficiently represent the couplings between points. On the basis of this matrix the C-points, which exist at the coarser level, and the F-points, which only exist at the finer level, are set up. Of course, these couplings can also be used for interpolation.

In order to obtain an overall converging PAMG approach which is also efficient in memory consumption and cpu time, we accelerate the multigrid method with BICGSTAB. On the coarsest level we use the direct solver ParDiSo [9].

3.1.1 Smoothing

Although, at least formally, any relaxation method could be used as a smoother in a point-based approach, a distinctly point-oriented smoothing often is a prerequisite for the success of a point-based approach. Since it is often most appropriate for handling strong unknown cross couplings and for producing an algebraically smooth error which allows for a point-coarsening.

Considering ordinary convection-diffusion systems, the smoother has to capture the direction of convection appropriately. In the case of the multi-ion model, however, this is not that obvious. Actually, additionally to the convective term in our PDE system, we have to deal with migration. The migration term also has properties which are very similar to convection. The smoother shall thus handle the convective and as well the migrative part appropriately. However, finding a

compromise is very difficult since the migration field is more or less orthogonal to the direction of convection, see Figure 1. Another difficulty is that the migration field develops during the simulation process, thus in the first iterations one is not able to detect it.

Commonly used point-based smoothing approaches are Block-Gauss-Seidel and Block-ILU. The PAMG method employing the Block-Gauss-Seidel smoother does not converge, this is due to the fact that BGS does not smoothen the error. A reordering of points improves the smoothing behavior of BGS, however it does nonetheless not lead to an overall converging approach. Hence, we chose the BILU(0) smoother for our PAMG method.

In Section 4.2, we will show that the PAMG approach with BILU(0) smoothing is not efficient for complicated geometries which deal with recirculation areas. Therefore, we propose the use of alternating smoothing. Alternating smoothing is well known from the context of geometric multigrid methods. This idea strikes a little bit the common philosophy of AMG since the original idea was to improve the interpolation and coarsening process and use very simple Gauss Seidel smoothers. However, BILU(0) together with the reordering framework introduced in Section 3.2 helps to capture the physics easily within this alternating smoothing. Hence, we make a step towards a universally applicable AMG method. Also one can think of extending this idea to the interpolation and coarsening further on, see Section 5.

3.1.2 Coarsening

In PAMG, the coarsening is done with the help of a primary matrix which has dimension $n \times n$, if n is the number of points. The primary matrix has to reasonably represent the coupling structure of the linear system. Since the linear systems of our electrochemical simulations can have totally different coupling characters depending on the reactive terms, it comes in hand to choose the primary matrix based on norms. We will make use of the maximum norm of the coupling blocks $p_{kl} = ||A(k,l)||_{max}$. p_{kl} enters in the primary matrix. $A(k,l)$ denotes the point-coupling matrix of dimension "number of unknowns in point k" $\times$ "number of unknowns in point l".

After its setup, the primary matrix is treated as if it were given into classical AMG. The algorithm creates the coarsening strategy for the primary matrix and then transfers it to the original matrix, which we want to solve.

3.1.3 Interpolation

The employed PAMG method makes use of the single-unknown interpolation (SU). The SU-interpolation uses variable-wise interpolation formulas which are the "same" for the variables belonging to the same point, so that unknown cross interpolation is avoided. The procedure is similar to the coarsening process. Namely, classical AMG is applied to the primary matrix, and then the interpolation scheme is transferred to the variables belonging to each point. Note that the primary matrix for interpolation and the one for coarsening do not have to be the same. However, it is obvious that this often is the best choice. In our case of MiTReM we take the same primary matrix for obtaining the interpolation and coarsening operator.

3.2 The Reordering Framework

In this section, we give a short introduction into the reordering framework employed in our numerical experiments in Section 4. The framework consists of two main components. In the first step, we construct a primary matrix, already known from the PAMG framework. Concrete variants are described in Section 3.2.1. This primary matrix is then interpreted as a graph.

Based on the graph a reordering is applied, see Section 3.2.2.

3.2.1 Creating a Primary Matrix

Since we deal with a point-based AMG approach it is obvious to base the new ordering on a matrix which only represents the couplings between (grid) points. Hence, we base our reordering again on a primary matrix, see Section 3.1.2. Note, that the primary matrix for reordering does not have to be the same as the one used to determine the fine and coarse grid points in the context of the PAMG framework.

The couplings in the resulting primary matrix will usually build up many cycles, if one interprets the primary matrix as a graph representation. Therefore, it is useful to further modify the primary matrix before determining a new ordering. There are many possibilities to modify the primary matrix. One can for example think of erasing small entries, or of using only strong couplings. Choosing a modification based on strong couplings does not lead to significant changes for the arising primary matrices from MiTReM. Hence, we decided to "strengthen the asymmetry" of the primary matrix. Namely, we delete all entries a_{pq} with $|a_{pq}| < factor * |a_{qp}|$. This modification was proposed in [5] and also in [3]. In order to get rid of all small cycles consisting of only two nodes we choose the factor higher than one.

3.2.2 Set up an Ordering

Once we have found a primary matrix which represents the main direction of information propagation, we apply an ordering algorithm. One can think of applying e.g. depth-first search [11], Reverse Cuthill-McKee [2] or heuristic feedback vertex set [5]. Also it is possible to interpret the graph as a weighted graph, see [8]. Doing so the resulting ordering will usually be nearly independent of the initially given one. This is very important, since in industries one is often not able to prescribe the initial ordering given by the simulator.

We order the points with the help of a block-triangular ordering algorithm using a weighted graph. With this method we gain a block-block-structure of the matrix. Each outer block represents a strongly connected component and each inner block represents a point-coupling matrix. The use of weights is very important for our method, otherwise, we obtain orderings which do not differ much from the originally given one. This can be expected, since we do not have many strongly connected components in the linear systems because of the two directions of information propagation.

Since the reordering has only to be done once in the setup phase, the method is very cheap.

4 Numerical Experiments

In the following, we want to analyze the results of our experiments and examine the usefulness of the AMG components chosen. The section begins with a description of the model configuration 4.1. We introduce different geometries and grid sizes of industrial relevance. Also, we describe the investigated ion-models which have different physical properties and thus also different numerical properties. In Section 4.2 we describe problems occurring for the original PAMG approach using an ordinary BILU(0) smoother. In Section 4.4 we then investigate how these problems can be remedied with help of alternating smoothing, employing physically based ordering described in Section 4.3.

4.1 Model Configuration

Name	small	big	Geometry
channel	30123	49693	channel with length 300 and height 30
bfs	31868	54624	backward facing step with length 300, inlet height 10 and outlet height 30, step after 50 electrodes behind the recirculation area at 110
crevice	36611	54011	channel with length 300 and height 30 and a crack at the bottom after 145 with width 10 and height 100

Table 1: Description of the grids, electrode positions and grid sizes

For our numerical experiments we use two grid sizes for each geometry. The geometries with the associated grid sizes are shown in Table 1. Namely, we make use of a simple channel geometry, a more challenging crevice geometry and a backward facing step.

In Table 2 the ion systems which we will deal with are shown. The numbering of the ion-systems represents the rank of difficulty of each system. Namely, system 1 is the easiest to solve. Note that systems 1 and 2 do not make use of homogeneous reactions, which means that there exists no reactive term $R_i \neq 0$ in Equation 1 in Section 2. System 3 is a realistic silver-ion model, the other systems are model cases designed to display the arising difficulties very well.

4.2 Multi-Level Analysis

In this section, we investigate the performance of the PAMG approach employing standard BILU(0) smoothing. For this method, we discover difficulties if dealing with complicated geometries. Hence, we perform a convergence analysis with varying numbers of levels. With the help of this analysis, we are able to point out scaling problems and other problems concerning the coarse level correction, which stem from the coarsening or interpolation.

It appears that the two-level method does not work for all considered models. This is due to the fact that the matrix for these cases on the second level is numerically singular, which causes a failure in the direct solver ParDiSo which we use on the coarsest level. The singularity is most likely caused by the very small diagonal entries in the lines of the conservation equation (Equation 2 in Section 2).

Considering Table 3, we find a significant increase of the convergence rates after level 4 in case of systems 1 and 2. In case of system 3 this increase already appears between levels 3 and 4. These significant increases in the convergence rates point to a scaling problem on coarser levels, which can be caused by the smoother as well as by an insufficient coarsening or interpolation. Note, that the described effect only appears for complicated geometries and big grid sizes. Considering the channel geometry up to 50,000 grid points no problems with the hierarchy can be detected.

If we take a look at the coupling structure of the coarse level matrices, we see a significant change in the structure between level 2 and 3 for all geometries employing systems 1 and 2. Namely, the share of the concentration to concentration couplings in the primary matrix rises. Another

No.	Homog. Reactions	Ions
1	-	Ag^+, NO_3^-
2	-	$Ag^+, NO_3^-, K+$
3	x	$NaS_2O_3^-, Na^+, S_2O_3^{2-}, NO_3^-, AgS_2O_3^-, Ag(S_2O_3)_2^{3-}$

Table 2: Description of the ion systems

Ion-	Level						
system	2	3	4	5	6	7	8
1	x	0.13	0.16	0.27	0.56*	0.74*	
2	x	0.12	0.18	0.27	0.36	0.62	0.69
3	x	0.27	0.49	0.55	0.80	0.89*	

Table 3: Average convergence rates for bfs-big with different numbers of levels at 80 per cent of the limiting current; (x): these methods did not work due to a singular coarse level matrix; (*): For these methods the Newton process diverged

observance is that for level 5, on the models where we observe the problems with the hierarchy, the BILU(0) smoother does not converge anymore. This also points to a change in the matrix character.

Via the above analysis, the observed problems within the hierarchy cannot be distinctly ascribed to either the smoother or the coarse grid correction. It might be, that the problems of the smoother are caused by an insufficient coarse grid correction, which then leads to a deficient behavior of the smoother. However, it might also be vice-versa. Since it is pretty easy to create an ordering which reasonably represents the coupling structure for each level within the reordering framework. We stick to attuning the smoothing process in this paper. The next step will then be transferring the ideas to the coarse grid correction, see Section 5.

In the following, we show that the problems within the multigrid hierarchy can be captured by a modified smoothing which follows the physical properties of the underlying systems.

4.3 Physically Based Orderings

In this section, we want to investigate which components of the reordering framework we should use to obtain an ordering which reasonably represents the physical character of the underlying model. The specific ordering is then used within an alternating BILU(0) smoothing method.

The choice of the primary matrix within the reordering framework has the highest impact on the resulting ordering. Therefore, we shall choose the primary matrix in a way that it represents the physical effects of the underlying system reasonably well. One possibility to find a good compromise between all effects for our systems is the choice of a norm-based primary matrix. Which specific norm we choose to create the primary matrix, here, is less important, since one unknown coupling is always very dominant.

We will have a look at a primary matrix based on the Schur norm. Another possibility to choose the primary matrix is to base it on so-called primary variables. If we choose the c_i to c_i coupling matrix as our primary matrix, the convection character is represented. On the other hand, if we choose the primary matrix based on the c_i to potential coupling matrix we end up with a primary matrix which represents, somehow, the character of the migration. This is shown in Figure 2 where the vectors which point to the strongest coupled point is visualized for each point. Note that the migration field is not given initially. Rather, it changes during the whole simulation. Therefore, the c_i to concentration couplings only indirectly represent the migration field.

Another needful parameter to find a reasonable ordering is the possibility to reduce the number of cycles in the graph which is represented by the primary matrix. We already mentioned in Section 3.2 that we "strengthen the asymmetry" of our matrix. If one does not reduce the primary matrix, one most probably gets a very unstructured ordering as it is shown in Figure 3. In particular, if there are many different effects in the matrix as for system 3 which is the one with reactive terms. Note, if we "strengthen" the asymmetry of the Schur norm based primary

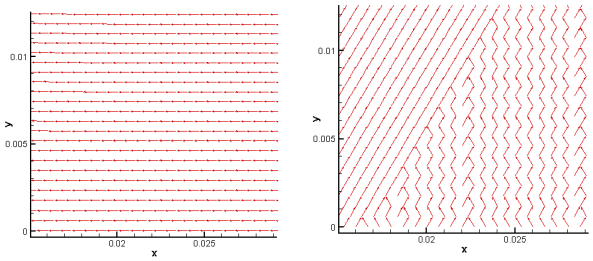

Figure 2: Illustration of directions of the strongest coupling. Left hand side shows the directions for c_i to c_i primary matrix, on the right hand side the directions for the c_i to potential primary matrix are displayed.

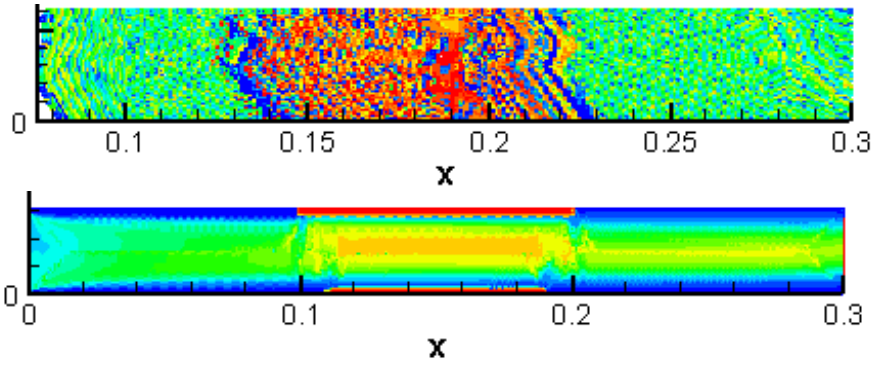

Figure 3: Illustration of resulting orderings on the channel geometry. Top: Ordering for system 3 based on Schur norm without reducing the primary matrix. Bottom: Ordering for system 1 based on Schur norm with "strengthened" asymmetry.

matrix of system 3, we will get an ordering structure which is nearly orthogonal to the ordering of system 1 shown in Figure 3.

Since our systems can have very different properties depending on the underlying electrochemical model, it is usually very useful to employ a norm-based approach. The different characters of the systems become very clear if comparing system 3, which has dominating reactive terms and system 1, which has dominating convective terms, see Figure 3 for an illustration.

We should note that it is of course also possible to choose a different ordering algorithm. For example one could choose the heuristic feedback vertex set algorithm proposed in [5]. However, this algorithm has a higher computational demand and does in the case of MiTReM not lead to significantly better results. Hence, we stick to the block-triangularization, which shows the best overall performance in aspects of cpu time and convergence rates of the AMG approach for the chosen primary matrices.

Considering all the described phenomena, we conclude that a choice of a Schur norm based primary matrix which is modified by "strengthening of the asymmetry" and applying a block-triangularization based on weights, we obtain a very good compromise of the underlying physical properties.

4.4 Numerical Results

In this section, we present an PAMG approach employing alternating BILU(0) smoothing which leads to convergence rates independently of the problem size and geometry. In contrast to the PAMG approach using ordinary BILU(0) smoothing.

Grid-	Ion-	channel			crevice			bfs		
size	system	RCM	ALT	5-Lev	RCM	ALT	5-Lev	RCM	ALT	5-Lev
small	1	0.20	0.12	0.22	0.30	0.29	0.24	0.23	0.19	0.18
	2	0.22	0.08	0.15	0.31	0.27	0.23	0.45	0.27	0.22
	3	0.23	0.19	0.09	0.70	0.65	0.35	0.38	0.66	0.21
big	1	0.23	0.25	0.24	div.	0.25	0.26	div.	0.49	0.27
	2	0.29	0.17	0.23	0.44	0.17	0.17	0.69	0.65	0.27
	3	0.21	0.26	0.09	div.	0.26	0.26	div.	0.66	0.55

Table 4: Average convergence rates at 80 percent of the limiting current. Problem sizes are ranging from 90k (small channel with system 1) to 385k (big bfs with system 3) DoF; div.: For these methods the Newton process diverged

We compare the PAMG approach with an ordinary BILU(0) smoother and initial reverse-Cuthill-McKee ordering (RCM) given by the simulation code, to the PAMG approach employing alternating BILU(0) smoothing (ALT). The third method displayed in Table 4, denoted with "5-Lev", is the same approach as RCM, despite that the number of levels is restricted to five, which means that this method is not of linear complexity anymore.

The alternating BILU(0) smoother of ALT makes use of the original RCM ordering alternating with the ordering resulting form the reordering framework if using a Schur norm based primary matrix, "strengthening of the asymmetry" (see Section 3.2.1) and block-triangularization, see Section 3.2.2. The convergence rates shown in Table 4 are the averages over a whole simulation run, containing 10 to 50 linear systems. "div." denotes that for this approach the Newton iteration did not converge due to the poor residual reduction of the AMG method.

Considering Table 4, we recognize a significant improvement of the convergence rates for the ALT approach compared to RCM for systems 1 and 2. In particular, this is the case for complicated and big geometries, where the RCM method does not work properly. The method 5-Lev is shown in order to visualize the possible convergence rates of RCM if no hierarchy problems occur. Therefore, we are very satisfied with the performance of ALT for systems 1 and 2.

However, the results for system 3 indicate that there is some potential to improve our approach. The ordering resulting form the Schur norm based primary matrix of the reordering framework for system 3 is very different to the ones of systems 1 and 2, because of the dominating reactive terms. This possibly leads to an insufficient representation of the convection. Note, that we could, of course choose, a primary matrix based on the c_i to c_i couplings which would mainly represent the convection. However, in this case the reactive terms are underrepresented so that the convergence rates are also not satisfying. Hence, we note again that it is very important to incorporate all physical properties of the underlying system.

The price of the reordering for the considered industrial relevant problem sizes is less than 2 multigrid cycles. Since setting up the orderings for each level is only done once during the setup phase, the method shows a pretty good performance in runtime.

Note that the PAMG approach will not converge if using solely the reordering of the framework without alternating BILU(0) smoothing. This can be most likely explained by the fact that not all physical effects are represented sufficiently well in the resulting ordering. BILU(0) one level iterative solvers are known to solve well employing RCM ordering. Hence, a heuristic explanation for the efficiency of the applied PAMG approach might be, that the RCM ordering is advantageous for solving a possible wrong scaling of the error on the coarser grids, and the primary matrix works well for smoothing the error.

5 Conclusion and Outlook

We investigated a PAMG approach employing alternating smoothing for solving linear systems arising in electrochemical applications. The presented method has a level independent convergence rate, which is also independent of the problem size. The performance in cpu runtime is also satisfying, since the orders are only created once during the setup phase and cost only few multigrid cycles. Additionally, we demonstrated the benefits of a reordering framework based on the point-structure to the smoothing process. The reordering framework can easily be adjusted to tackle various physical effects. Our numerical experiments recommend, in particular, to consider the physics of the underlying problem in the multigrid algorithm.

One focus of our future research will be finding primary matrices which represent the real migration part of the underlying systems algebraically to further enhance the approach. A general goal is to consider all physical effects of the underlying system within the smoothing process. Therefore, we have to find a measure for these effects, which is then used to define the respective primary matrices. This will, especially, improve the results for systems with dominating reaction terms.

The primary matrices used in the reordering framework should also be beneficial if used in the coarsening and interpolation process, since optimality of the multilevel approach always needs a good interplay between the different methods.

Another focus is building up the theoretical roots which help analyzing the physical relevance of an ordering. Our goal is to create an algorithm offering an automatic reordering for various industrially relevant PDE systems.

References

[1] T. Clees. *AMG Strategies for PDE Systems with Applications in Industrial Semiconductor Simulation.* PhD thesis. Dissertation, University of Cologne, Germany, November 2004; ISBN: 3-8322-4497-2.

[2] E. Cuthill and J. McKee. Reducing the bandwidth of sparse symmetric matrices. In *Proceedings of the 1969 24th national conference*, pages 157–172, New York, NY, USA, 1969. ACM.

[3] H. Kim, J. Xu, and L. Zikatanov. A multigrid method based matching in graph for convection diffusion equations. *Num. Lin. Alg. and Appl.*, 10:181–195, 2003.

[4] G. Larson, D. Synder, D.V. Abeele, and T. Clees. Application of single-level, pointwise algebraic, and smoothed aggregation multigrid methods to direct numerical simulations of incompressible turbulent flows. *Comput. Visual. Sci.*, 11:27–40, 2008. ISSN 1432-9360 (Print) 1433-0369 (Online), DOI 10.1007/s00791-006-0055-4, Online First: SpringerLink Date Tuesday, January 09, 2007.

[5] S. Le Borne. *Multigrid Methods for Convection-Dominated Problems.* PhD thesis, Dissertation, Christian-Albrechts-University of Kiel, 1999.

[6] G. Nelissen. *Simulation of Multi-Ion Transport in Turbulent Flow.* PhD thesis, Fakulteit Toegepaste Weteneschappen, Vrije Universiteit Brussel, 2003.

[7] J. Newman and K. E. Thomas-Alyea. *Electrochemical Systems.* Wiley-Interscience, 3rd edition, 2004.

[8] B. Pollul and A. Reusken. Numbering Techniques for Preconditioners in Iterative Solvers for Compressible Flows. *Journal for Numerical Methods in Fluids*, 55(3):241 – 261, 2006.

[9] O. Schenk. *Scalable Parallel Sparse LU Factorization Methods on Shared Memory Multiprocessors.* PhD thesis, ETH Zurich, 2000.

[10] K. Stüben. An introduction to algebraic multigrid. In [12], 413–532.

[11] R. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972.

[12] U. Trottenberg, C.W. Oosterlee, and A. Schüller. *Multigrid.* Academic Press, London, 2001. With appendices by K. Stüben, P. Oswald, and A. Brandt.