Supervisor's Report

Dissertation Title: **Multigrid methods for large-scale problems: approximate coarsest-level solves and mixed precision computation**
Author: **Petr Vacek**

Summary and Evaluation of the Ph.D. Thesis

The Ph.D. thesis presented by Petr Vacek represents a significant contribution to the area of practical multigrid methods for large-scale problems. These methods are important both as stand-alone solvers and as preconditioners for linear systems arising, e.g., from the discretization of partial differential equations. The multigrid method is based on having a hierarchy of grids - in the geometric multigrid variant, these arise from discretization of the continuous problem on multiple nested meshes of different resolutions. The method proceeds by using so-called smoothing on fine levels, which typically consists of a small number of iterations of a stationary iterative method, and the solution of the linear system on the coarsest level, which can be accomplished either by direct or iterative solvers.

Determining the accuracy to which the coarse problem should be solved and how this influences the overall number of multigrid iterations was something that was previously typically done based only on experimentation. The first chapter of Petr Vacek's PhD thesis provides a rigorous mathematical understanding of these effects. He has derived new stopping criteria for iterative solvers used for coarse problems within multigrid such that the overall number of multigrid iterations remains the same as if an exact solve had been performed (i.e., the number of coarse-level iterations is minimized while ensuring that there is no convergence delay). This work was published in the SIAM Journal on Scientific Computing, a top journal in our field. I should also note that this work was done while Petr Vacek was visiting Kirk Soodhalter at Trinity College, Dublin, on an ERASMUS+ stay. Petr Vacek came up with this research topic on his own, and all the analysis, implementation, and experiments are his own work. Kirk and I served in supervisory roles, and our contributions were mainly limited to providing guidance on the scientific presentation of the results.

The second chapter of the thesis focuses on residual-based multilevel a posteriori estimates for the total and algebraic error. This work stemmed from the Master's thesis of the candidate, supervised by Prof. Zdeněk Strakoš. The analysis here shows that the approximation of the term associated with the coarsest level has a significant effect on the overall accuracy and robustness of the estimates. This leads to the development of a new approximation for the coarsest-level term, based on the use of the iterative conjugate gradient method with an appropriate stopping criterion. The stopping criterion here represents an adaptive approach, in which the number of iterations is determined so that the accuracy of the bounds do not depend on the size of the coarsest-level problem. That the resulting estimates retain the desired properties is shown via a rigorous proof. I personally provided only brief comments on the presentation of the results; the work and the collaboration was the candidate's own initiative.

The third chapter of the thesis tackles the problem of the use of mixed precision within multigrid V-cycle methods, which stemmed from a collaboration with Professor Ulrich Rúde during the candidate's visit to FAU Erlangen and with Professor Hartwig Anzt during the candidate's (invited) visit to TU Munich. Driven by the emergence of commercially-available mixed precision hardware over the last decade, and the significant potential performance improvements possible when using low precisions, the study of where and how low precision can be used safely within numerical linear algebra routines is an important open problem. Petr Vacek's work in this chapter is at the forefront of this research, tackling the question of how the used of various precisions within the smoothing and coarse-grid solve steps affect the overall error. The analysis of the candidate indicates that in some practical cases, an incomplete Cholesky-based smoother can be applied in a precision lower than used for the other parts of the computation without detriment. The candidate has collaborated with the authors of the Gingko software library (the group of Hartwig Anzt) to implement this approach and perform mixed precision experiments on modern GPU hardware. It is shown that the mixed precision approach can obtain up to

$2\times$ speedups over the approach that uses double precision throughout. This is a substantial improvement for a sparse problem. A journal paper on this work is currently in preparation.

Overall, the thesis of Petr Vacek represents a significant body of work in the field of numerical mathematics. He has developed several novel approaches related to the practical implementation of multigrid methods, which are based on rigorous theoretical analysis. This work will no doubt be useful to practitioners in the field and will stimulate further research. I again would like to stress that the work in this thesis is the candidate's own, including the conception of the topics, the analysis, the experiments, and the writing (modulo some grammatical edits).

Evaluation of the Candidate

During his time in our program, Petr Vacek was an extremely valuable and productive member of my team. He is fully capable of independent scientific research; I have been very impressed with his ability to develop key research questions and to undertake their solution in a systematic, scientifically-sound manner. He is confident and capable in public speaking, can present his ideas clearly, and has given presentations at many international conferences during his time as a PhD student. He has further gained experience in organizing minisymposia at top conferences in our field. The work in this thesis has thus far resulted in one journal paper (the first chapter); I expect the other two chapters will result in published papers as well, and I expect that the groundwork laid in this thesis will lead to other fruitful research questions.

Fostering international collaborations beyond Charles University has been one of Petr Vacek's major strengths, and he has initiated and completed many visits to other groups during his Ph.D. studies. He has fostered ongoing collaborations with Professors Kirk Soodhalter (Trinity College, Dublin), Ulrich Rüde (FAU Erlangen), and Hartwig Anzt (TU Munich).

Petr Vacek has secured a postdoctoral research position at IFPEN in Paris, as part of the project "EoCoE - Energy Oriented Center of Excellence: toward exascale for energy", working under Ani Anciaux Sedrakian (IFPEN), Alberto Buttari (IRIT), and Emmanuel Agullo (INRIA). This position is set to begin in January. The topic of his work on this project will be efficient multigrid solvers for large-scale problems, which he is well-suited for given his Ph.D. research.

To conclude, it has been an absolute pleasure to know and work with Petr Vacek, and his ingenuity and diligence have been a great inspiration. I was proud to have him as a member of my team, and consider myself very lucky that he chose me as his supervisor. He has an incredibly bright future in mathematical research.

Conclusion

I recommend accepting the presented thesis and awarding the Ph.D. degree.

in Prague, 28.11.2024

Dr. Erin Claire Carson, Ph.D.