Final Report

DOE Award Number: DE-SC0004880
Recipient: Worcester Polytechnic Institute
Project Title: *Anderson Acceleration for Fixed-Point Iterations*
Principal Investigator: Homer F. Walker
Period Covered: 09/01/2010 — 08/31/2015

Accomplishments.

The purpose of this grant was to support research on acceleration methods for fixed-point iterations, with applications to computational frameworks and simulation problems that are of interest to DOE. The particular methods of interest are based on a method introduced by Anderson [1], referred to as *Anderson acceleration* here.[1] For many years, this method has enjoyed considerable success and wide usage within the computational physics, chemistry, and materials communities as a means of accelerating the *self-consistent field* iteration used in electronic-structure computations. However, at the outset of the grant period, this method was very little known in other applications communities, and it remained untried in many important applications in which it seemed likely to be used to advantage. Moreover, while other acceleration methods had been extensively studied within the mathematics and numerical analysis communities, Anderson acceleration had received relatively little attention from those communities, despite there being many significant unanswered mathematical questions.

During the period of this grant, this situation changed considerably. The applications of Anderson acceleration have expanded significantly, and the method is much more widely known among numerical analysts as well as a broader range of applications specialists. In the following, we describe the research accomplishments and other activities supported by this grant and indicate how they have contributed to this change.

Accomplishments. During the period of this grant, two major projects were completed. The first of these was aimed at developing basic theoretical aspects of Anderson acceleration that shed light on its behavior and also at exploring its performance on a broad set of test applications. This work culminated in the paper [19] in *SIAM Journal on Numerical Analysis*, which the PI co-authored with Peng Ni, who was then a PhD student at WPI. This paper establishes fundamental relationships between Anderson acceleration and GMRES and provides promising experimental results in a variety of applications ranging from statistical estimation to simulation of transonic flow. This work has had considerable impact, despite having appeared only in 2011. According to Google Scholar, it has been cited 70 times as of this writing. Of these 70 citations, 26 occured in 2015, and at least 13 involved one or more researchers at DOE national laboratories. Subjects addressed in the citing papers (with numbers of citing papers in

[1]Methods that appear to be mathematically equivalent have been independently proposed by several authors; see [4], [14], [20], [15].

parentheses) include advection/convection-diffusion problems (3), astrophysics (2), eigenvalue problems (1), electronic structure computations (3), financial mathematics (1), flux-correction methods (4), graph theory (1), linear and nonlinear solvers (15), materials modeling (1), multi-physics problems (8), optimal control (1), signal processing (2), social science modeling (1), statistical computing (2), stochastic differential equations (1), transport equations (5), water resources and reservoir simulations (6), wave propagation (3), and other numerical methods and applications (5). These citing papers have in turn already been cited a total of 372 times. For a complete list of citing papers, see

https://scholar.google.com/scholar?hl=en&as_sdt=40000005&cites=12197795612154032936&scipsc=

The second major project completed during the grant period was a joint effort with Carol Woodward, Ulrike Yang, and Aaron Lott in the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL). The primary goal was to apply Anderson acceleration to Picard iteration for solving the partial differential equations that model variably saturated flow in porous media. Picard iteration is widely used in this context but often converges too slowly to be efficient, compared to the major alternative approach, viz., methods based on Newton's method. In our project, guided by results in [19], we developed a preliminary C implementation of Anderson acceleration within the KINSOL nonlinear solver code, a part of the SUNDIALS suite developed at LLNL [6], and also routines for Picard iteration, specifically a widely-used modified form introduced in [5]. We thoroughly tested this modified Picard iteration both with and without acceleration over a range of problem parameters that included those most often appearing in the literature. For comparison, we also included KINSOL's sophisticated Newton-GMRES-linesearch method in our tests.

The figures below show run-time performance profiles[2] and run times (in seconds) for the Newton-GMRES-linesearch method, the modified Picard iteration, and the modified Picard iteration with Anderson acceleration (using a maximum of five saved residual vectors), based on 408 sampled points in a rectangle enclosing most commonly used α-N parameters. All results are for a 1024×1024 grid. The dark-red regions indicate failure to find a solution within 1800 seconds (30 minutes).

It is clear from these figures that, in these extensive tests, the modified Picard iteration with Anderson acceleration is significantly more robust than the other two methods, often more efficient as well, and never much slower than the fastest method. In view of the problem size (over 1M unknowns), it is notable that the method achieved such favorable results with no more than five saved residual vectors.

We feel that this work represents a significant advance in solution methodology for variably saturated subsurface flow problems and also in understanding the capabilities and limitations of Picard iteration and Newton-based methods for these problems. A paper summarizing the results has been published in

[2]These summarize the comparative run-time performance of the three methods on the problem set. In brief, the run-time performance profile for each method is a monotone increasing function such that (1) the leftmost value is the fraction of problems on which the method was fastest; (2) the rightmost value is the fraction of problems on which it succeeded; and (3) at an intermediate point τ, the value of the function is the fraction of problems on which the method required no more than τ times the best run time among the methods. See [8] for details.

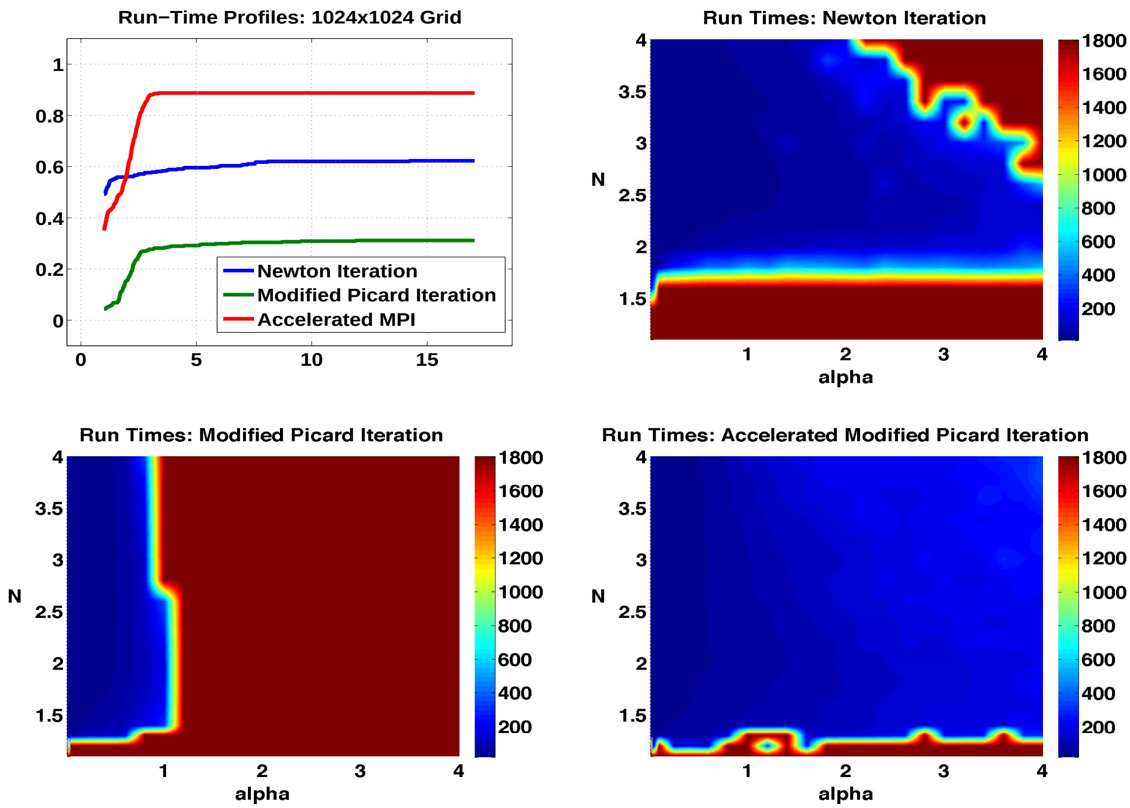

Advances in Water Resources [13]. According to Google Scholar, this paper has been cited 12 times since appearing in 2012.

As part of this project with CASC/LLNL, the PI developed an efficient and numerically sound MATLAB implementation of Anderson acceleration. This refined implementation is described in a WPI Mathematical Sciences Department technical report [18]. It provided the basis for a C implementation of Anderson acceleration in KINSOL [6] and has been useful in a number of applications. The PI gave a talk describing this implementation at a special session on Advances in Nonlinear and Linear Solvers for Water Resources Applications at the XIX Conference on Computational Methods in Water Resources (CMWR2012) in June, 2012.

Another outcome of this collaboration has been the use of Anderson acceleration in other applications at LLNL. In addition to the further development of the KINSOL software [6], these include applying the method to speed implicit integration of dislocation-dynamics equations [10] and exploring the method's possible advantages in communication-minimizing schemes for distributed-memory and GPU-based parallel computations [12].

Two additional projects were undertaken and are nearing completion at this time. The first is a collaboration with S. Olson, a colleague in the WPI Mathematical Sciences Department, and N. Ho, a PhD student in the department. This project is aimed at assessing the effectiveness of Anderson acceleration applied to the *Uzawa algorithm* [2], a widely used and easily implemented iterative procedure for approximately solving saddle-point problems. In this, we considered the Uzawa algorithm with and without preconditioning and compared its performance with and without acceleration on the linear saddle-point Stokes and Oseen systems in several steady incompressible-flow scenarios. The preconditioner for Uzawa on the Oseen problems was taken from [9], in which Sandia National Labs (SNL) researchers played a large collaborative role. For additional perspectives, we included restarted GMRES preconditioned with the state-of-the-art *Relaxed Dimensional Factorization* (RDF) preconditioner developed in [3]. The results were very promising: Accelerated Uzawa greatly outperformed Uzawa without acceleration, converging significantly faster in most cases and succeeding in several cases in which unaccelerated Uzawa failed. Additionally, accelerated Uzawa was competitive with (and occasionally outperformed) GMRES with the RDF preconditioner. Ho gave a well-received talk on preliminary results at the 2014 Copper Mountain Conference on Iterative Methods. An advanced draft of this work is available in [11] and will be submitted for publication after a final polishing.

The second project nearing completion is aimed at exploring the effectiveness of Anderson acceleration applied to the *Expectation-Maximization* (EM) algorithm for determining maximum-likelihood estimates of the parameters in statistical mixtures of normal (Gaussian) probability densities [7, 17]. In this context, the EM algorithm has a number of very attractive properties but can exhibit very slow convergence in some circumstances. Experiments so far have involved trials with randomly generated samples that are *unlabeled*, i.e., have no labels indicating their populations of origin. We have conducted hundreds of trials involving samples of sizes up one million observations drawn on multivariate normal mixtures with dimensions ranging from two to 20. The results have been very encouraging. The accelerated EM algorithm is usually much faster than the EM algorithm without acceleration. Moreover, it is slightly more robust as well. The PI has given talks on these results at the Minisymposium on Anderson Acceleration and Applications, SIAM Conference on Computational Science and Engineering, in March 2015 and at the ICERM Topical Workshop on Numerical Methods for Large-Scale Nonlinear Problems and Their Applications in September 2015. This work has been done jointly with J. Plasse, a former PhD student at WPI who recently transferred to Imperial College London. It has been on hold since Plasse left WPI in summer 2015 but is expected to resume once he has become settled in his new program.

An additional project has been on hold pending completion of the two projects above but will be resumed once those two are finished. This is a software project in which the PI has been developing a MATLAB implementation of the NITSOL Fortran code [16]. A rudimentary implementation has been completed at this time; this needs to be extensively refined and tested. NITSOL is a widely used and influential code implementing Newton-iterative methods for solving large-scale nonlinear systems. The paper [16] has been cited 231 times, according to Google Scholar. The algorithm has influenced the development of nonlinear solvers in MPSalsa and NOX at SNL, KINSOL at LLNL, and PETSc at Argonne National Lab. The MATLAB code will be significantly more capable and flexible than the original Fortran code. It will augment the backtracking globalization in the Fortran code with a trust-region option, accommodate a much broader variety of preconditioners (indeed, almost arbitrary combinations of preconditioning functions and matrices), and allow use of all of the Krylov subspace methods and incomplete-factorization routines offered by MATLAB. It will also be easier to use than the Fortran code because users will not have to define work arrays of specified sizes and will have more flexibility in providing inputs. The MATLAB implementation that will ultimately result will be of considerable interest to the computational science and engineering community.

Publications during the project period relating to the supported research:

- H. F. Walker and P. Ni, *Anderson acceleration for fixed-point iterations*, SIAM J. Numer. Anal., 49 (2011), 1715-1735.
- P. A. Lott, H. F. Walker, C. S. Woodward, and U. M. Yang, *An accelerated Picard method for nonlinear systems related to variably saturated flow*, Advances in Water Resources, 38 (2012), 92-101.
- H. F. Walker, *Anderson Acceleration: Algorithms and Implementations*, WPI Math. Sciences Dept. Report MA 6-15-50, June 2011, revised October 2012.
- N. Ho, S. D. Olson, H. F. Walker, *Accelerating the Uzawa Algorithm*, arXiv:1510.04246~[math.NA], October, 2015.
- J. H. Plasse, H. F. Walker, *Accelerating the EM algorithm for multivariate normal mixture densities*, in preparation.

Talks and other presentations by the PI during the project period relating to the supported research:

- Schlumberger–Tufts University Seminar, September 2011.
- Department of Energy Applied Mathematics Program Meeting, October 2011.
- AMS-SIAM Special Session on the Mathematics of Computation: Differential Equations, Linear Algebra, and Applications, Boston, Massachusetts, January 2012.
- University of Houston Mathematics Department Colloquium, Houston, Texas, April 2012.
- Session on Advances in Nonlinear and Linear Solvers for Water Resources Applications, XIX Conference on Computational Methods in Water Resources (CMWR2012), Urbana, Illinois, June 2012.
- Numerical Methods in PDEs Seminar, MIT, Cambridge, Massachusetts, October 2012.
- Minisymposium on Anderson Acceleration and Applications, SIAM Conference on Computational Science and Engineering, Salt Lake City, Utah, March 2015.
- Topical Workshop on Numerical Methods for Large-Scale Nonlinear Problems and Their Applications, Institute for Computational and Experimental Research in Mathematics (ICERM), Providence, Rhode Island, September 2015.

Synergistic activities by the PI during the project period period relating to the supported research:

- Sabbatical spent visiting CASC/LLNL and collaborating with C. S. Woodward, U. Yang, P. A. Lott, and others in CASC, September 2010 – May 2011.

- Co-organizer, Special sessions on Anderson acceleration and applications, Copper Mountain Conferences on Iterative Methods, Copper Mountain, Colorado, March 2012 (with C. S. Woodward); April 2014 and March 2016 (with C. T. Kelley and C. S. Woodward).
- Co-organizer (with C. T. Kelley and C. S. Woodward), Topical Workshop on Numerical Methods for Large-Scale Nonlinear Problems and Their Applications, Institute for Computational and Experimental Research in Mathematics (ICERM), Providence, Rhode Island, August–September 2015.
- Co-organizer (with C. T. Kelley and C. S. Woodward), Minisymposium on Anderson Acceleration and Applications, SIAM Conference on Computational Science and Engineering, Salt Lake City, Utah, March 2015.

Unexpended funds: There are no unexpended funds.

References

[1] D. G. Anderson. Iterative procedures for nonlinear integral equations. *J. Assoc. Comput. Machinery*, 12:547–560, 1965.

[2] K J Arrow, L Hurwicz, and H Uzawa. *Studies in Linear and Nonlinear Programming*. Stanford University Press, Stanford, CA, 1958.

[3] M Benzi, M Ng, Q Niu, and Z Wang. A relaxed dimensional factorization preconditioner for the incompressible Navier–Stokes equations. *J Comput Phys*, 230:6185–6202, 2011.

[4] N. N. Carlson and K. Miller. Design and application of a gradient-weighted moving finite element code I: in one dimension. *SIAM J. Sci. Comput.*, 19:728–765, 1998.

[5] M. A. Celia, E. T. Bouloutas, and R. L. Zarba. A general mass-conservative numerical solution for the unsaturated flow equation. *Water Resour. Res.*, 26:1483–1496, 1990.

[6] Aaron M. Collier, Alan C. Hindmarsh, Radu Serban, and Carol S. Woodward. User documentation for KINSOL v2.8.2 (SUNDIALS v2.6.2). Technical Report UCRL-SM-208116, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, July 29 2015.

[7] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum-likelihood from incomplete data via the EM algorithm. *J. Royal Statist. Soc. Ser. B (methodological)*, 39:1–38, 1977.

[8] E. D. Dolan and J. J. Moré. Benchmarking optimization software with performance profiles. *Math. Program., Ser. A*, 91:201–213, 2002.

[9] H. C. Elman, V. E. Howle, J. Shadid, R. Shuttleworth, and R Tuminaro. Block preconditioners based on approximate commutators. *SIAM J. Sci. Comput.*, 27:1651–1668, 2006.

[10] D. G. Gardner, C. S. Woodward, D. R. Reynolds, G. Hommes, and S. Aubry. Implicit integration methods for dislocation dynamics. *Modeling Simul. Mater. Sci. Eng.*, 23, 2015.

[11] N. Ho, S. D. Olson, and H. F. Walker. Accelerating the Uzawa algorithm. *arXiv:1510.04246~[math.NA]*, 2015.

[12] J. Loffeld and Carol S. Woodward. Considerations on the implementation and use of anderson acceleration on distributed memory and gpu-based parallel computers. Technical Report LLNL-PROC-675918, Lawrence Livermore National Laboratory, August 5 2015.

[13] P. A. Lott, H. F. Walker, C. S. Woodward, and U. M. Yang. An accelerated Picard method for nonlinear systems related to variably saturated flow. *Advances in Water Resources*, 38:92–101, 2012.

[14] K. Miller. Nonlinear Krylov and moving nodes in the method of lines. *J. Comput. Appl. Math.*, 183:275–287, 2005.

[15] C. W. Oosterlee and T. Washio. Krylov subspace acceleration of nonlinear multigrid with application to recirculating flows. *SIAM J. Sci. Comput.*, 21:1670–1690, 2000.

[16] M. A. Pernice and H. F. Walker. NITSOL: a newton iterative solver for nonlinear systems. *SIAM J. Sci. Comput.*, 19:302–318, 1998.

[17] R. A. Redner and H. F. Walker. Mixture densities, maximum likelihood, and the EM algorithm. *SIAM Rev.*, 26:195–239, 1984.

[18] H. F. Walker. Anderson acceleration: Algorithms and implementations. Technical Report MA 6-15-50, WPI Math. Sciences Dept., June 2011. Revised October 2012.

[19] H. F. Walker and P. Ni. Anderson acceleration for fixed-point iterations. *SIAM J. Numer. Anal.*, 49:1715–1735, 2011.

[20] T. Washio and C. W. Oosterlee. Krylov subspace acceleration for nonlinear multigrid schemes. *Electronic Trans. Numer. Anal.*, 6:271–290, 1997.