1. Overview of my work. My main research area is numerical linear algebra and with more recent work being at the interface between this field and work on large-scale ill-posed problems. Primarily I design and analyze Krylov subspace iterative methods and develop efficient implementations thereof. Krylov subspace methods are a class of iterative methods which play an important role in the solution of large-scale sparse matrix equations and eigenvalue problems in cases where the matrix is too large to be efficiently factored or even stored in memory or in cases in which one only has a procedure encoding the matrix-times-vector multiplication. Such methods are quite versatile and are ubiquitous in the computational sciences. As increasing computational power render the solution of problems of increasing dimension and complexity feasible and as new application areas arise, bringing new and interesting problem structures to light, numerical linear algebra remains a vibrant and changing landscape in which to be working, and Krylov subspace methods have demonstrated themselves to be versatile tools.

My interests are both in the analysis and development of Krylov subspace iterative methods, and this began with the work contained in my doctoral thesis [22] (written under the supervision of Daniel B. Szyld of Temple University and defended in March 2012). In addition to the resulting papers, this work also produced one high-performance computing code [18] in [1], and many MATLAB implementations, which can be found at [23]. In my postdoctoral years, I have continued this work while also branching out to broaden my skills and interests. At times, my project choice has been guided by a desire to gain experience with mathematical tools that will be useful in other contexts. This has allowed for unexpected crossover in some of my current projects and has led to new project ideas.

In my dissertation research, I focused on Krylov subspace methods which have a fixed memory requirement, i.e., there is a hard limit on how many vectors need to be stored at one time during execution [22]. This these was divide into three parts. One part involved studying a residual minimizing iterative method which exploited the structure of "nearly Hermitian" matrices to only require storage of a fixed number of vectors but exhibited unexplained instabilities. A floating point analysis illuminated the precise source of the instability. An alternative more stable fixed-memory algorithm was then proposed. These results were published in [3]. In the course of this work, I also developed a new band Lanczos-based block MINRES solver [26].

The other two parts of my dissertation concerned two extensions of a particular type of Krylov subspace augmentation scheme called subspace recycling. In such methods, in addition to generating a Krylov subspace, one includes in the solution reconstruction additional vectors encoding known information about the true solution or the operator to accelerate convergence. In one project, we proposed a block recycled version of the Krylov subspace method called GMRES [17]. In the other project, we explored the difficulties which arise when extending subspace recycling methods to solve multiple non-Hermitian shifted linear systems, which our analysis showed is not generally possible in that setting. This work was published in [28]. At the beginning of my postdoctoral time, I continued in this direction, achieving better results and allowing for preconditioning [25] and finally by using a Sylvester equations interpretation to robust subspace recycling-based shifted system solver [27]. I have continued to explore problems with underlying Kronecker structure, resulting in one submitted paper [14] and one current project and other ideas for future work. This is also a topic of consideration in my future research plans. Other topics of interest include analyzing Krylov subspace augmentation schemes for solving ill-posed problems for image reconstruction and extending a framework proposed for preconditioned methods for symmetric systems to the nonsymmetric case.

Other work completed during my postdoctoral time includes the development of a new MINRES implementation which allows for the inexpensive evaluation of physically relevant subvector preconditioned residual subvector norms [7] and an analysis of the convergence behavior of block GMRES and block FOM [24].

2. Mathematical Background. Krylov subspaces are used for solving a variety of problems, often of the form

$$\mathcal{A}(\mathbf{X}_0 + \mathbf{T}) = \mathbf{F},$$

where $\mathcal{A}$ is a linear operator defined by the action of a matrix or a linear combination of the left and right actions of multiple matrices (e.g., Sylvester and Lyapunov operators), $\mathbf{X}_0$ and $\mathbf{F}$ are known vectors/matrices of conforming dimensions, and $\mathbf{T}$ is the unknown correction satisfying the equation. For example, my dissertation work concerned solving linear systems of the form

$$\mathbf{A}(\mathbf{x}_0 + \mathbf{t}) = \mathbf{f}$$

where $\mathbf{A} \in \mathbb{C}^{n\times n}$ is sparse, $n \gg 0$ is large, and $\mathbf{f} \in \mathbb{C}^n$. If we take $\mathbf{r}_0 = \mathbf{f} - \mathbf{A}\mathbf{x}_0$ to be the initial residual, then we can define the *Krylov subspace* at iteration j generated by $\mathbf{A}$ and $\mathbf{r}_0$ as the space of polynomials in $\mathbf{A}$ of degree less than j acting on $\mathbf{r}_0$,

$$\mathcal{K}_j(\mathbf{A}, \mathbf{r}_0) = \operatorname{span}\left\{\mathbf{r}_0, \mathbf{A}\mathbf{r}_0, \mathbf{A}^2\mathbf{r}_0, \ldots, \mathbf{A}^{\mathrm{j}-1}\mathbf{r}_0\right\}.$$

This subspace is built one dimension per iteration at the cost of one matrix-vector product and we must only orthogonalize each new vector with respect to the two most recent ones. The resulting vector will be orthogonal to the others by default, in exact arithmetic. This means one need not store the entire basis for $\mathcal{K}_j(\mathbf{A}, \mathbf{r}_0)$, and one can use the symmetric Lanczos process to generate the basis. The Krylov subspace is also invariant with respect to scalar shifts of the identity,

$$\mathcal{K}_j(\mathbf{A}, \mathbf{r}_0) = \mathcal{K}_j(\mathbf{A} + \sigma\mathbf{I}, \mathbf{r}_0), \tag{2.1}$$

and this has been exploited for solving multiple shifted linear systems, which arise often in diverse applications such as Tikhonov regularization, interior point methods, lattice quantum chromodynamics, diffuse optical tomography, etc. This has been a focus of my research [28, 25, 27]

At iteration j, one takes $\mathbf{t}_j \in \mathcal{K}_j(\mathbf{A}, \mathbf{r}_0)$ which is determined according to some constraint on the residual $\mathbf{r}_j = \mathbf{b} - \mathbf{A}(\mathbf{x}_0 + \mathbf{t}_j)$, such as the minimum residual condition $\mathbf{r}_j \perp \mathbf{A}\mathcal{K}_j(\mathbf{A}, \mathbf{r}_0)$ which is the basis for the generalized minimum residual method (GMRES) [20]. Another closely related condition is the Galerkin condition $\mathbf{r}_j \perp \mathcal{K}_j(\mathbf{A}, \mathbf{r}_0)$ upon which the full orthogonalization method (FOM) is based. If $\mathbf{A}$ is Hermitian, the need to store only the most recent vectors can be exploited to derived fixed-storage, short-term recurrence methods such as a Hermitian version of GMRES called the minimum residual method (MINRES) [15] and of FOM called conjugate gradients (CG) [8], in the case that $\mathbf{A}$ is also positive definite.

We also note that for most real-world, large-scale problems, one must also precondition the system to obtain sufficiently rapid convergence. Preconditioning refers to a transformation of the linear system to another with the same solution for which a Krylov subspace method produces more rapidly converging iterates. For example, for a matrix $\mathbf{M} \in \mathbb{R}^{n\times n}$ representing some approximation of $\mathbf{A}$ can be applied on the left, yielding the left-preconditioned problem

$$\mathbf{M}^{-1}\mathbf{A}\mathbf{x} = \mathbf{M}^{-1}\mathbf{b}.$$

If $\mathbf{A}$ is Hermitian, one often requires that $\mathbf{M}$ be an SPD preconditioner.

Krylov subspace methods can be generalized to treat

$$\mathbf{A}(\mathbf{X}_0 + \mathbf{T}) = \mathbf{F} \tag{2.2}$$

where $\mathbf{F} \in \mathbb{C}^{n\times p}$ with $p > 1$ and these are called block Krylov subspace methods. In a block Krylov method, one generates a nondirect sum of Krylov subspaces from the columns of $\mathbf{R}_0$. Block Krylov methods have recently generated renewed interest, as one can generate much larger subspaces using operations with superior data movement and cache use characteristics [2, 19]. They can be built to solve linear systems with a single right-hand side in order to accelerate convergence of the iteration by constraining the residual using larger subspaces.

Block Krylov subspaces are often used in the solution of other matrix equations such as Sylvester equations

$$\mathbf{A}(\mathbf{X}_0 + \mathbf{T}) + (\mathbf{X}_0 + \mathbf{T})\mathbf{B} = \mathbf{F} \quad \text{with} \quad \mathbf{B} \in \mathbb{C}^{p\times p} \tag{2.3}$$

with $p \ll n$. One builds a block Krylov subspace using the $n \times p$ initial residual from which to draw approximations to $\mathbf{T}$. Furthermore, if we denote the action of the Sylvester operator by $\mathcal{S} : \mathbf{X} \to \mathbf{AX} + \mathbf{XB}$, then the block Krylov subspace satisfies a block version of the scalar shift invariance, with

$$\mathbb{K}_j(\mathcal{S}, \mathbf{R}_0) = \mathbb{K}_j(\mathbf{A}, \mathbf{R}_0), \tag{2.4}$$

and this fact was exploited in [27] to successfully combine a minimum residual shifted system solver with the subspace recycling framework.

Krylov subspace recycling methods refer to those in which one begins with a subspace $\mathcal{U}$ which has been precomputed, often during a previous application of an iterative method to a problem similar to the present one. One wishes to construct approximations to $\mathbf{T}$ from $\mathcal{U}$ and from a Krylov subspace. Such a method is sometimes called an augmentation method, but what differentiates subspace recycling methods from other augmentation schemes is that the Krylov subspace is not generated from $\mathbf{A}$ and $\mathbf{r}_0$. Rather, both the operator and starting vector are premultiplied by a projector related to $\mathcal{U}$ so that the Krylov subspace is restricted to prevent the iteration from revisiting directions already represented by the subspace $\mathcal{U}$. The precise details depend upon the method be augmented. In the recycled GMRES method [16], for example,, we set $\mathcal{C} = \mathbf{A} \cdot \mathcal{U}$. The initial residual is projected into $\mathcal{C}^\perp$ and $\mathbf{x}_0$ updated accordingly, and a GMRES iteration is performed which is restricted to the space $\mathcal{C}^\perp$. Using a specific projector onto $\mathcal{U}$, one can map the resulting correction back to the full space in order to construct the residual minimizing correction $\mathcal{U} + \mathcal{K}_j$ where $\mathcal{K}_j$ is the Krylov subspace generated by the GMRES iteration in $\mathcal{C}^\perp$.

3. Current and future work.

3.1. Preconditioned subvector norms in MINRES. I enjoy working with problems that are highly structured, whether it is apparent or hidden. In work with Roland Herzog described in [7], we explored the necessary conditions to be able to progressively compute preconditioned residual subvector norms produced when a MINRES iteration is applied to the symmetric-indefinite saddle-point system

$$\begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \end{bmatrix}$$

with a symmetric positive-definite (SPD) preconditioner. These systems arise in applications such as optimization where the blocks have different physical meaning. In the optimization setting, for example, the problem is set up such that the first block corresponds to optimality and the second to feasibility. They generally represent different quantities and are measured in different physical units. Being able to monitor the residual subvector norms separately also offers more control and flexibility if the MINRES algorithm is being used inner solver contained in an optimization code. One can then relax how stringently either the optimality or feasibility conditions are satisfied. We showed that to do this, the preconditioner being used must be block diagonal. In this setting, this type of preconditioner is called "natural", as the preconditioner-induced norms of the two residual subvector are decoupled, respecting the fact that they represent different physical quantities.

In such problems, the SPD preconditioners, induce the norm in which the residual is measured. There is a school of thought that one should consider the system matrices as maps from a Hilbert space (where the approximate solutions live) to its dual (where the residuals live), and the SPD preconditioner is a Riesz map back to the original space. A Krylov method then requires a preconditioner since the space is generated by repeated applications of the operator; see, e.g., [4]. A question that arises from this interpretation is: can we interpret a nonsymmetric method such as GMRES in this way? A nonsymmetric preconditioner does not induce a norm. However, they may still be considered as a mapping back from the dual space. We will explore what the equivalent framework would be in the nonsymmetric case, for a method such as GMRES and develop methods allowing the user to monitor appropriate subvector norms, as in the MINRES case.

3.2. Problems with Kronecker structure. In [26], we were able to demonstrate that by not fully exploiting the shifted system structure, one can much more easily introduce acceleration strategies, such as subspace recycling as well as general preconditioning. These tools are not fully available when one exploits the shifted system invariance of the Krylov subspace (2.1).

These ideas in [26] also underlie the work in [27]. By not taking direct advantage of (2.1), one sees that by interpreting the shifted system family as a set of Sylvester equations, we can avail ourselves of the already available solvers in the literature. This then allows us to propose a subspace recycling solver for shifted systems which exploits the Sylvester shift-invariance (2.4). The key to proposing a viable GMRES-like algorithm was to understand how to correctly project each shifted system residual such that (when they are taken together) the block residual with respect to the Sylvester operator has columns which are orthogonal to the space $\mathcal{C}$. One cannot simply apply the orthogonal projector onto $\mathcal{C}^{\perp}$ to all residuals because there would be no way to update the associated approximations to the solution. This required the use of specially constructed oblique projectors directly related to the shifts (i.e., to the exact structure of the Sylvester operator in play) which are applied to each shifted system residual individually. The recognition of this fundamental necessity (operator specific projectors) opens the door to extending the subspace recycling strategy to the setting of other matrix equations. The most straightforward of these would be Sylvester equations which do not correspond to a family of shifted systems, i.e., (2.3) in the case that $\mathbf{B}$ is not a diagonal matrix.

A more interesting and less straightforward problem is to combine the subspace recycling framework with solvers for continuous time Lyapunov equations

$$\mathbf{A}\mathbf{X} + \mathbf{X}\mathbf{A}^T = -\mathbf{F} \tag{3.1}$$

where the right-hand side $\mathbf{F}$ and unknown $\mathbf{X}$ have the same dimensions as the coefficient matrix $\mathbf{A}$. For large, sparse $\mathbf{A}$, this would render the problem prohibitively expensive to solve. However, in many applications $\mathbf{F}$ has low-rank structure with $\mathbf{F} = \mathbf{G}\mathbf{G}^T$ such that $\mathbf{G} \in \mathbb{R}^{n\times p}$, and $p \ll n$. In this situation, the solution $\mathbf{X}$ can be well-approximated by a low-rank matrix, and there is a great deal of analysis which relates the rank of $\mathbf{F}$ to the singular value decay of $\mathbf{X}$. This is important as $\mathbf{X}$ is, in general, not sparse. Thus, obtaining an approximate solution with a low-rank representation is crucial for the feasibility of developing solvers in this setting.

Let $\mathbf{W} \in \mathbb{R}^{n\times m}$ be a matrix with columns spanning a subspace $\mathcal{W}$ from which we want to construct an approximation to $\mathbf{X}$ from (3.1). This is accomplished using the low-rank assumption on the right-hand side $\mathbf{F}$. We define the low-rank approximation of $\mathbf{X}$ to be $\mathbf{X}_{\mathcal{W}} = \mathbf{W}\mathbf{Y}_{\mathcal{W}}\mathbf{W}^T$ where $\mathbf{Y}_{\mathcal{W}} \in \mathbb{R}^{m\times m}$. Just as with linear systems and Sylvester equations, one selects $\mathbf{Y}_{\mathcal{W}}$ according to a constraint on the residual $\mathbf{R}_{\mathcal{W}} = \mathbf{A}\mathbf{X}_{\mathcal{W}} + \mathbf{X}_{\mathcal{W}}\mathbf{A}^T + \mathbf{F}$. If one vectorizes the problem, one can consider, for example, the FOM and GMRES residual constraints in the Kronecker products linear system setting. One must then choose $\mathcal{W}$ (or rather a sequence of subspaces which grows at each iteration) such that we obtain a good low-rank approximation of $\mathbf{X}$ and the residual constraints can be easily applied. At iteration j, one could generate the block Krylov subspace $\mathcal{W} = \mathbb{K}_j(\mathbf{A}, \mathbf{G})$. This allows us to compute $\mathbf{Y}_j$ as the solution of a small Lyapunov-type equation (FOM) or least squares problem (GMRES). However, a better space to choose yielding superior convergence would be the extended Krylov subspace, which is the sum of block Krylov subspaces generated by $\mathbf{A}$ and $\mathbf{A}^{-1}$, i.e., $\mathbf{K}_j(\mathbf{A}, \mathbf{G}) = \mathbb{K}_j(\mathbf{A}, \mathbf{G}) + \mathbb{K}_j(\mathbf{A}^{-1}, \mathbf{A}^{-1}\mathbf{G})$. This is the basis for many proposed methods, e.g., [9, 12, 21], and is still an active area of research.

There are a variety of problems in which one must solve a sequence of slowly changing Lyapunov equations, such as in the discretization of time-dependent Ricatti equations. Thus it is useful to extend subspace recycling type strategies to this setting. This is the subject of a current project I am working on together with Hermann Mena and Daniel B. Szyld. This raises a number of interesting challenges beyond those encountered in the shifted system or Sylvester equations setting. As before, one must define an appropriate projection of the initial residual that is compatible with the recycling method and admits a corresponding update of the approximate solution. In addition, one must unambiguously define the notion of an extended Krylov subspace when the underlying operator is the composition of the matrix $\mathbf{A}$ and a projector. Further complicating matters is the fact that the initial residual in this setting is generally taken to be the right-hand side with $\mathbf{X}_0 = \mathbf{0}$. Any initial projection of the residual and updating of $\mathbf{X}_0$ must be done without explicit construction of the updated approximation or residual, as we wish to only store their low-rank representations. All of these challenges are surmountable, and we are currently developing a stable implementation in order to pursue further testing.

3.3. Ill-posed problems. The construction of a low-rank approximation to an operator using a Krylov subspace is also a strategy I have been pursuing in the setting of blind deconvolution when solving ill-posed problems, and Krylov subspace augmentation has the potential to play an important role here as well.

When we discuss ill-posed problems, we mean in the sense described by Hadamard [5]. Consider an operator $T : \mathcal{X} \to \mathcal{Y}$ and the equation $Tx = y$. The problem is called ill-posed if the unknown x does not depend continuously on y. Therein lies the problem. The right-hand side is acquired by taking measurements with scientific instruments or sensors. Therefore, we do not possess y but rather some y^δ with $\left\| y - y^\delta \right\| \leq \delta$ where $\delta > 0$ is called the measurement error. Without measurement error, one could compute the optimal approximation $x^\dagger = T^\dagger y$ using the generalized inverse. However, due to the aforementioned ill-posedness, this frequently leads to unbounded errors. Instead one seeks to compute a regularized solution. In regularizing a problem, we seek to construct an approximate solution x_{reg} which limits the unbounded errors due to measurement while still maintaining some fidelity to the data. Krylov subspace methods with the iteration truncated appropriately represent one class of regularization methods.

In many cases, the operator in question represents a model of image blurring or corruption describing how the wanted, unknown image/signal was transformed into the measured image/signal. When the blurring matrix, or equivalently, the point spread function (PSF), is explicitly known, this task commonly is referred to as deconvolution. It is often the case that one does not actually possess an accurate representation of the underlying operator, but rather some approximation. Furthermore, one may only have a procedure which applies this incorrect operator to an image. This is the case for the PSF, which is simply convolved with the image in Fourier space. In such a case, one often wants to both reconstruct the wanted image and also obtain an improved approximation of the underlying model. This process is called *blind deconvolution*

Following previous work on blind deconvolution by diagonalization of the operator [6, 10], in a project with Lothar Reichel and colleagues, we have been developing a Krylov subspace method for simultaneous image reconstruction and blind deconvolution for symmetric blurring matrices. We propose a family of blind deconvolution methods that allow a user to adjust the blurring matrix used in the computations to achieve an improved restoration. The major computational effort required for large-scale problems is the partial reduction of an available large symmetric approximate blurring matrix by a few steps of the symmetric Lanczos process. As is the case for the Lyapunov solver described above, we reconstruct an approximation of our operator from the Krylov subspace. Here, though, we partially diagonalize the problem, obtaining a few of the dominant eigenvectors of the incorrect operator but replace the eigenvalues with ones better approximating the action of the true operator while also reconstructing the signal according to the data. Because everything is represented in the eigenvector basis, this entire process can be performed per-coordinate, meaning that once the Krylov subspace has been generated, the remaining calculations are with scalars.

The target application for this project is adaptive optics in ground-based telescopes. Here, the image is a section of the night sky in order to view astronomical phenomena (e.g., stars). One cause of distortion of the wanted image is the light passing through Earth's turbulent atmosphere. Using measurement data and simulation software, one can generate approximations of the PSF representing the atmospheric distortion. The goal here is to get an improved approximation of the blurring matrix, represented now as a low-rank matrix rather than a PSF while also obtaining a stable reconstruction of the image. For the case of a PSF whose convolution with the image can be represented as the action of a symmetric matrix, we have been able to obtain high quality image reconstructions with our method using the incorrect operator, where other iterative methods (such as Arnoldi-Tikhonov [11]) produce nonsense.

We are pleased with these results, but there is much work to be done. Due to aberrations caused by measurement noise, the PSF is not represented by a symmetric matrix. Thus we must develop a nonsymmetric version of this procedure which does not rely on the Lanczos process which allows us to instead reconstruct a low-rank approximation of the operator through a truncated singular value decomposition generated by a method such as the Golub-Kahan bidiagonalization method. Krylov subspace augmentation will also play a role in this algorithm. In a distorted image, astronomers often know where certain undistorted astronomical phenomena are located, either from other simultaneously recorded measurements or from general knowledge of the measured region of the sky. One would like to incorporate this prior knowledge into both the full image reconstruction and into the improved reconstruction of the blurring operator. Thus, we will combine this blind deconvolution strategy into the subspace recycling framework in order to design a blind deconvolution method based on a nonsymmetric Krylov subspace iteration which takes advantage of the knowledge and experience of astronomers.

The role of augmentation techniques such as subspace recycling have been widely explored in the well-posed problems community, and there is a great deal of research concerning them, some of which has been discussed in this document. With Eric de Sturler and Misha Kilmer, I am currently writing a review paper which seeks to describe all such techniques in one framework and offer a guide for practitioners in the computational science community who wish to use these methods for their particular application.

In the ill-posed problems community, conversely, these techniques have received some attention, but they have not been fully explored or analyzed. I am currently working on a project developing new augmentation methods based on flexible GMRES for ill-posed problems with Lothar Reichel. In the course of working on this project, we have classified many of the existing augmentation strategies which have been proposed in the large-scale ill-posed problems community in order to create a bridge between what has been done in the well- and ill-posed problem communities, as these connections had not previously been well-known in either community. In the existing literature, in fact, similar methods have often developed in parallel with the subspace recycling techniques.

The goals of the user when augmenting a Krylov subspace are different in the ill-posed setting. One does not generally attempt to accelerate the iteration. The goal instead is to prescribe specific structure to be included in the reconstructed image, such as a-priori known information about the image. This information is encoded in vectors spanning an augmentation subspace. The iteration on the projected problem then resolves the part of the image not determined by the imposed features. In the literature, there are many promising numerical results, but there has been virtually no analysis demonstrating their reliability or underlying regularization properties. The long-term goal is to develop a complete, self-contained regularization theory for augmented Krylov subspace methods and derive convergence rates as the norm of the noise approaches zero. In the first stage, we will consider the finite dimensional setting. The results obtained will be connected to the infinite dimensional case through the work in [13] and then be used as a guide to obtain results in the infinite dimensional Hilbert space setting.

We will analyze specific applications, such as image reconstruction and processing, adaptive optics, and medical imaging, applications for which one often has a-priori information or desired characteristics to include in the image reconstruction. We will also investigate what type of information in combination with an augmented Krylov subspace method will lead to fast and stable image reconstruction. In addition, we will analyze the behavior of these methods when one has inexact information about an image, such as that generated by a preprocessing algorithm.

This work will lead to the development of new iterative algorithms for image reconstruction applications. We will develop both general purpose algorithms which can automatically insert information for augmentation without user input as well as hybrid algorithms in which the augmentation methods are combined with a preprocessing method which reconstructs specific aspects of the image, which can then be fed into the augmented Krylov subspace iteration.

REFERENCES

[1] *The Trilinos Project website.*, Accessed December 10, 2014. `http://trilinos.org/`.
[2] A. H. Baker, J. M. Dennis, and E. R. Jessup, *On improving linear solver performance: a block variant of GMRES*, SIAM Journal on Scientific Computing, 27 (2006), pp. 1608–1626.
[3] M. Embree, J. A. Sifuentes, K. M. Soodhalter, D. B. Szyld, and F. Xue, *Short-term recurrence krylov subspace methods for nearly hermitian matrices*, SIAM. J. Matrix Anal. and Appl., 33-2 (2012), pp. 480–500.
[4] A. Günnel, R. Herzog, and E. Sachs, *A note on preconditioners and scalar products in Krylov subspace methods for self-adjoint problems in Hilbert space*, Electron. Trans. Numer. Anal., 41 (2014), pp. 13–20.
[5] J. Hadamard, *Lectures on Cauchy's problem in linear partial differential equations*, Dover Publications, New York, 1953.
[6] T. A. Hearn and L. Reichel, *Extensions of the Justen-Ramlau blind deconvolution method*, Adv. Comput. Math., 39 (2013), pp. 465–491.
[7] R. Herzog and K. M. Soodhalter, *A modified implementation of MINRES to monitor residual subvector norms for block systems*, ArXiv e-prints (and submitted for publication), (2016).
[8] M. R. Hestenes and E. Stiefel, *Methods of conjugate gradients for solving linear systems*, Journal of Research of the National Bureau of Standards, 49 (1952), pp. 409–436 (1953).
[9] M. Heyouni and K. Jbilou, *An extended block Arnoldi algorithm for large-scale solutions of the continuous-time algebraic Riccati equation*, Electronic Transactions on Numerical Analysis, 33 (2008/09), pp. 53–62.
[10] L. Justen and R. Ramlau, *A non-iterative regularization approach to blind deconvolution*, Inverse Problems, 22 (2006), pp. 771–800.
[11] B. Lewis and L. Reichel, *Arnoldi-Tikhonov regularization methods*, J. Comput. Appl. Math., 226 (2009), pp. 92–102.
[12] Y. Lin and V. Simoncini, *Minimal residual methods for large scale Lyapunov equations*, Appl. Numer. Math., 72 (2013), pp. 52–71.
[13] F. Natterer, *Regularisierung schlecht gestellter Probleme durch Projektionsverfahren*, Numer. Math., 28 (1977), pp. 329–341.
[14] J. Niebsch, R. Ramlau, and K. M. Soodhalter, *Solution of coupled differential equations arising from imbalance problems*, Submitted, (2015).
[15] C. C. Paige and M. A. Saunders, *Solutions of sparse indefinite systems of linear equations*, SIAM Journal on Numerical Analysis, 12 (1975), pp. 617–629.
[16] M. L. Parks, E. de Sturler, G. Mackey, D. D. Johnson, and S. Maiti, *Recycling Krylov subspaces for sequences of linear systems*, SIAM Journal on Scientific Computing, 28 (2006), pp. 1651–1674.
[17] M. L. Parks, K. Soodhalter, and D. B. Szyld, *A block recycled GMRES method with investigations into aspects of solver performance*, Tech. Rep. 16-04-04, Department of Mathematics, Temple University, Apr. 2016.
[18] M. L. Parks and K. M. Soodhalter, *Block GCRO-DR.* in Belos package of the Trilinos C++ Library, 2011.
[19] M. L. Parks, K. M. Soodhalter, and D. B. Szyld, *A block Recycled GMRES method with investigations into aspects of solver performance*, ArXiv e-prints (and submitted for publication), (2016).
[20] Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.
[21] V. Simoncini, *A new iterative method for solving large-scale Lyapunov matrix equations*, SIAM J. Sci. Comput., 29 (2007), pp. 1268–1288.
[22] K. M. Soodhalter, *Krylov subspace methods with fixed memory requirements: Nearly Hermitian linear systems and subspace recycling*, PhD thesis, Temple University, 2012.
[23] ———, *Personal website*, 2014. `http://math.soodhalter.com`.
[24] ———, *Stagnation of block GMRES and its relationship to block FOM*, ArXiv e-prints (and submitted for publication), (2014).
[25] ———, *Two recursive GMRES-type methods for shifted linear systems with general preconditioning*, ArXiv e-prints (and submitted for publication), (2014).
[26] ———, *A block MINRES algorithm based on the banded Lanczos method*, Numerical Algorithms, 69 (2015), pp. 473–494.
[27] ———, *Block krylov subspace recycling for shifted systems with unrelated right-hand sides*, SIAM Journal on Scientific Computing, 38 (2016), pp. A302–A324.
[28] K. M. Soodhalter, D. B. Szyld, and F. Xue, *Krylov subspace recycling for sequences of shifted linear systems*, Applied Numerical Mathematics, 81C (2014), pp. 105–118.