A fast algorithm for parabolic PDE-based inverse problems based on Laplace transforms and flexible Krylov solvers

Tania Bakhos ^{*†} Arvind K. Saibaba [‡] Peter K. Kitanidis [§]

May 16, 2015

Abstract

We consider the problem of estimating parameters in large-scale weakly nonlinear inverse problems for which the underlying governing equations is a linear, time-dependent, parabolic partial differential equation. A major challenge in solving these inverse problems using Newton-type methods is the computational cost associated with solving the forward problem and with repeated construction of the Jacobian, which represents the sensitivity of the measurements to the unknown parameters. Forming the Jacobian can be prohibitively expensive because it requires repeated solutions of the forward and adjoint time-dependent parabolic partial differential equations corresponding to multiple sources and receivers. We propose an efficient method based on a Laplace transform-based exponential time integrator combined with a flexible Krylov subspace approach to solve the resulting shifted systems of equations efficiently. Our proposed solver speeds up the computation of the forward and adjoint problems, thus yielding significant speedup in total inversion time. We consider an application from Transient Hydraulic Tomography (THT), which is an imaging technique to estimate hydraulic parameters related to the subsurface from pressure measurements obtained by a series of pumping tests. The algorithms discussed are applied to a synthetic example taken from THT to demonstrate the resulting computational gains of this proposed method.

1 Introduction

Consider linear parabolic partial differential equations (PDEs) of the form,

$$\mathcal{S}(x; s) \frac{\partial \phi}{\partial t} - \mathcal{A}(x; s) \phi = f(x, t) \quad \phi(t = 0) = \phi_0 \quad (1)$$

along with the associated boundary conditions, where s are the parameters to be estimated, $f(x, t)$ represents the source term, $\phi(x, t)$ is the state variable and $\mathcal{S}$ and $\mathcal{A}$ are time-independent operators. We are interested in fast solutions to Equation (1) particularly in the context of inverse problems [28], i.e. using measurements ϕ to estimate parameters s , whose governing forward equations are described by Equation (1). Problems of this kind arise in many applications such as Diffuse Optical Tomography [1, 11] ($\mathcal{A} = \nabla \cdot (D \nabla \cdot) - \nu \mu_a$), electromagnetic inversion ($\mathcal{A} = \nabla \times \mu^{-1} \nabla \times$) [35] and Transient Hydraulic Tomography ($\mathcal{A} = \nabla \cdot (\kappa \nabla \cdot)$) [2]. For example, in Diffuse Optical Tomography, measurements of photon fluence are used to “invert” for diffusivity D and the absorption coefficient μ_a .

We will focus our attention on Transient Hydraulic Tomography (THT). THT is a method of imaging of the subsurface that uses a series of pumping tests to estimate important hydrological parameters, such as conductivity and storativity. In this method, water is injected continuously at a constant rate in injection wells and the resulting response in pressure change is recorded at measurement wells, until steady-state is approximately reached. The drawdown curves recorded by each measurement location are stored and from this data, a few key time sampling points are identified and the corresponding measured pressure data are used in an inversion algorithm

^{*}taniab@stanford.edu

[†]Institute for Computational and Mathematical Engineering, Stanford University

[‡]Department of Electrical and Computer Engineering, Tufts University

[§]Department of Civil and Environmental Engineering, Stanford University

to recover the aquifer parameters of interest (e.g. hydraulic conductivity and specific storage). Instead of solving for the groundwater equations in the traditional time-stepping formulation, we propose a Laplace transform-based solver. We emphasize that while we describe methods as related to THT, they can be extended to a large class of linear parabolic PDEs.

We focus our attention on methods based on successive linearization such as Gauss-Newton to solve the inverse problem. However, we note that other possible approaches include Ensemble Kalman Filters, particle filters, rejection samplers, Markov Chain Monte Carlo, for a good review please see [10]. Popular approaches for solving inverse problems based on time-dependent parabolic PDEs rely on traditional time-stepping methods (such as Crank-Nicolson) that require the solutions of linear systems of equations at each time step in order to implicitly march the solution in time. The time step is computed based on accuracy and stability requirements. This process is inherently sequential and is computationally expensive. A second challenge in the context of inverse problems solved using Newton-type methods is that the computation of the Jacobian or sensitivity using the adjoint state method requires storing the entire time history of the forward problem. To alleviate the memory costs, previous work has considered checkpointing methods to trade computational efficiency for memory [5, 27, 31]. Therefore, there is a need to develop a method that is accurate and efficient, both in terms of memory and computational costs.

For solving the problem that arises in the discretization of the forward problem, previous works have noted that applying the contour integral representation of Laplace transforms leads to a sequence of shifted systems of linear equations which can then each be solved for independently. This observation can be exploited to develop a parallel algorithm for time integration for parabolic PDEs [24] or an iterative Richardson method [25]. In [9], the application of Krylov subspace methods that take advantage of the shift-invariance of the Krylov subspace to simultaneously solve the shifted systems arising from the contour integral representation was considered. However, to the best of our knowledge, the application of these methods to inverse problems has not been performed.

Contributions: In this paper, we propose a Laplace transform-based exponential time integrator combined with a flexible Krylov subspace approach to speed up the computation of the forward and adjoint problems, thus yielding significant speedup in total inversion time. The modified Talbot contour is chosen to invert the Laplace transform and is discretized using the Trapezoidal rule to yield an exponentially convergent quadrature rule [30]. The resulting shifted systems of equations are solved using a flexible Krylov subspace solver that allows for multiple preconditioners to efficiently precondition the systems across the entire range of shifts [22]. This yields significant speedup in solving the forward problem. We propose heuristics to pick preconditioners and demonstrate the robustness of the proposed solver with several parameters (variance of the underlying conductivity field, measurement time samples and grid discretization). We show that for a certain range of measurement times, the corresponding inverse Laplace transforms can be sped up simultaneously using shared preconditioners across all shifts allowing for additional speedup.

Additionally, we derive a method for computing the Jacobian using the adjoint state technique based on the Laplace transform. Computing the Jacobian this way requires several solutions of shifted systems of equations. Using the solver developed in our previous work [22] allows for solving the forward (and adjoint) problem for multiple shifts at a cost that is comparable to the cost of solving a system of equations with a single shift. It should be noted that the fast computation of the Jacobian can be used not only to accelerate the geostatistical approach but also other techniques for inversion and data assimilation such as 3D Var, 4D Var filtering and Extended Kalman filter. An additional advantage for using the proposed solver to compute the Jacobian is that the rows corresponding to difference time points can be computed independently without the need to store the entire time history. We demonstrate results using our method to accelerate an example problem in THT.

Outline: The paper is organized as follows. We describe the governing equations for groundwater flow and the choice of contour used to approximate the inverse Laplace transform in Section 2. In Section 3, we discuss the flexible Krylov subspace solver for solving the resulting shifted linear systems of equations. In Section 4 we demonstrate the performance and robustness of the proposed solver under various test conditions. We follow this with a description of the geostatistical method for solving inverse problems and show an example taken from THT in Section 5. Finally, in Section 6 we conclude with a summary and a discussion on future work.

2 Forward Problem

2.1 Transient hydraulic tomography

In this work we will assume that flow in the aquifer of interest can be modeled as confined (with standard, linear elastic storage) or, if unconfined, can be treated appropriately using a saturated flow model with the standard linearized water table approximation [18]. The equations governing groundwater flow through an aquifer for a given domain Ω with boundary $\partial\Omega$ is composed of the union of three non-intersecting regions - $\partial\Omega_D$, $\partial\Omega_N$ and $\partial\Omega_w$ referring to Dirichlet, Neumann and linearized water table boundaries respectively.

$$\begin{aligned} S_s(x) \frac{\partial \phi(x, t)}{\partial t} - \nabla \cdot (\kappa(x) \nabla \phi(x, t)) &= f(x, t), & x \in \Omega \\ \phi(x, t) &= 0, & x \in \partial\Omega_D \\ \nabla \phi(x, t) \cdot n &= 0, & x \in \partial\Omega_N \\ \nabla \phi(x, t) \cdot n &= -S_y \frac{\partial \phi(x, t)}{\partial t}, & x \in \partial\Omega_w \end{aligned} \quad (2)$$

where $\kappa(x)$ [LT^{-1}] is the hydraulic conductivity, S_s [L^{-1}] is the specific storage, $f(x, t)$ [T^{-1}] is the pumping source, $\phi(x, t)$ is the hydraulic head (pressure) and S_y $[-]$ is the specific yield. In 2D aquifers, the terms $\kappa(x)$ [L^2T^{-1}] and $S_s(x)$ $[-]$ are known as transmissivity and storativity respectively. In this work, the pumping source is modeled as $f(x, t) = q(t)\delta(x - x_s)$, where x_s corresponds to the location of the pumping source and $q(t)$ to the pumping rate. To derive the weak formulation, we multiply by appropriately chosen test functions ψ . Integrating by parts

$$(S_s \partial_t \phi, \psi)_\Omega + (S_y \partial_t \phi, \psi)_{\partial\Omega_w} + A(\phi, \psi) = q(t)(\delta(x - x_s), \psi)_\Omega \quad (3)$$

where $(u, v)_\Omega \stackrel{\text{def}}{=} \int_\Omega uv dx$, $(u, v)_{\partial\Omega_w} \stackrel{\text{def}}{=} \int_{\partial\Omega_w} uv dx$ and $A(u, v) \stackrel{\text{def}}{=} \int_\Omega \kappa(x) \nabla u \cdot \nabla v dx$. In our work, we have used the standard linear finite element approach but extensions to higher order finite elements are possible. The above weak form is then discretized using finite elements as $\phi_h(t) = \sum_{j=1}^N \hat{\phi}_j(t) u_j(x)$ where $u_j(x)$ are the finite basis functions, to obtain the following semi-discrete system of equations

$$M \partial_t \phi_h + K \phi_h = q(t)b \quad \phi_{h,0}\{t=0\} = \mathcal{P}_h \phi_0 \quad (4)$$

where the matrices K and M have entries $K_{ij} = A(u_j, u_i)$, $M_{ij} = (S_s u_j, u_i)_\Omega + (S_y u_j, u_i)_{\partial\Omega_w}$ and the vector b has entries $b_j = (\delta(x - x_s), u_j)_\Omega$. Furthermore, $\mathcal{P}_h$ is the orthogonal projector with respect to the $(\cdot, \cdot)_\Omega$ inner product. Taking the Laplace transform $\hat{\phi}(\cdot, z) = \int_0^\infty \phi(\cdot, t) e^{-zt} dt$, we have

$$(K + zM) \hat{\phi}_h(z) = \hat{q}(z)b + M \phi_{h,0} \quad (5)$$

The field $\phi_h(t)$ can be recovered by applying the inverse Laplace transform on a carefully chosen contour using the formula

$$\phi_h(t) = \frac{1}{2\pi i} \int_\Gamma e^{zt} (K + zM)^{-1} (\hat{q}(z)b + M \phi_{h,0}) dz \quad (6)$$

where $\hat{q}(z)$ denotes the Laplace transform of $q(t)$ and Γ is an appropriately chosen contour. We will now discuss possible choices for Γ .

2.2 Choice of contour

To compute the inverse Laplace transform, we will employ the modified Talbot contour

$$\Gamma(\theta) : \quad z(\theta) = \sigma + \mu(\theta \cot \theta + \nu i \theta) \quad -\pi \leq \theta \leq \pi \quad (7)$$

where θ is the contour discretization parameter. This contour was analyzed in [30] with the contour parameters σ , μ and ν optimally chosen to improve the convergence of the quadrature scheme. The contour is a simple, closed curve that encloses both the eigenvalues of the generalized eigenvalue problem $Kx = \lambda Mx$ and the singularities of $\hat{q}(z)$.¹ In [33], it was shown that

¹In this work, we assume that $\hat{q}(z)$ does not have any singularities on the positive real axis. Otherwise, the contour has to be adapted to account for these singularities. This is discussed in [9].

optimal convergence as $N_z \rightarrow \infty$ keeping t fixed is achieved with σ and μ proportional to the ratio N_z/t .

As in [30], we approximate the integral using the Trapezoidal rule on a uniform grid on $[-\pi, \pi]$ using evenly spaced points θ_k spaced by $2\pi/N_z$,

$$\phi_h(t) = \frac{1}{2\pi i} \int_{-\pi}^{\pi} e^{z(\theta)t} z'(\theta) F(z(\theta)) d\theta \approx \sum_{k=1}^{N_z} e^{z(\theta_k)t} z'(\theta_k) F(z(\theta_k)) \quad (8)$$

where we defined $F(z) = (K + zM)^{-1} (M\phi_0 + \hat{q}(z)b)$. Further, we also define the weights of the quadrature $w_k = -\frac{i}{2N} e^{z(\theta_k)t} z'(\theta_k)$. Since (K, M) is a Hermitian definite pencil and the contour is symmetric, we need only use half the points,

$$\phi_h(t) \approx \sum_{k=1}^{N_z} w_k F(z(\theta_k)) = 2 \operatorname{Real} \left\{ \sum_{k=1}^{N_z/2} w_k F(z_k) \right\} \quad (9)$$

Figure 1: Plot of contours corresponding to $N_z = 20$ (left) and $N_z = 40$ (right). Because of symmetry, only half the contour is shown. Here N_z is the number of quadrature points in Equation (9)

An example of the contours for three different measurement times is shown in Figure 1. As mentioned earlier, in this paper we only consider sources $f(x, t) = q(t)\delta(x - x_s)$ but this method is applicable for all sources that are separable in space and time. This assumption is crucial since we intend to use Krylov subspace methods for shifted systems. This requires the right-hand sides to be independent of the shift (except, perhaps by a multiplicative factor). If the Laplace transform of $q(t)$ is known explicitly, it can be used. Otherwise, one may consider the use of Prony's method to fit an exponential sum $q(t) \approx \sum_j a_j \exp(-\lambda_j t)$ on $[0, T]$, for which the Laplace transform can be computed easily [25]. Although in this case, the right hand side is dependent on z_k , the Krylov subspace method for shifted systems can still be applied.

Apart from the Talbot contour, other contours such as parabolas and hyperbolas have been proposed and analyzed. In addition, the connection between the Trapezoidal rule applied to approximate the integral and rational approximations to the exponential function have been pointed out in [30]. We would like to emphasize that our fast algorithm does not depend on the specific choice of contour.

2.3 Accuracy and computational cost

The error between $\phi_{N_z, h}(t)$ and the true solution $\phi(x, t)$ in the $L_2(\Omega)$ norm can be bounded using the triangle inequality

$$\|\phi(t) - \phi_{N_z, h}(t)\| \leq \|\phi(t) - \phi_h(t)\| + \|\phi_h(t) - \phi_{N_z, h}(t)\|$$

The first term contributing to the error is due to the finite element discretization. As shown in Lemma 3.1 in [29], this error is

$$\|\phi(t) - \phi_h(t)\| \leq Ch^2 (\|\phi_0\| + \|\hat{f}\|_\Gamma)$$

where, $\|\hat{f}\|_\Gamma \stackrel{\text{def}}{=} \sup_{z \in \Gamma} |\hat{f}(z)|$ and C is a constant independent of h . The second contributing term is the error incurred by the discretization of the inverse Laplace transform and is dependent on the choice of parameters σ, μ and ν . In summary, the total error is bounded by $\mathcal{O}(e^{-cN_z} + h^2)$. For the parameters chosen, $c = \log 3.89$ [30].

We now analyze the computational cost of the Laplace transform method and compare it to a standard time-stepping scheme. Let us assume that we need to compute the solution $\phi(x, t)$ at N_T time-steps where $N_T = \mathcal{O}(1)$. For a given discretization, let the number of grid points be N and let the cost of solving a linear system of equation be $\mu(N)$. The cost of computing the solution at N_T points in time using the Laplace transform methods is $N_z N_T \mu(N)$ and using time-stepping schemes is $N_t \mu(N)$. Let us now analyze the cost for a desired tolerance ε . Assuming second-order discretization technique in time we have $C \Delta t^2 \sim \varepsilon$, so that $N_t = T/\Delta t \sim C' \varepsilon^{-1/2}$. However, for Laplace transform method the convergence is exponential, so the number of systems is $N_z \sim C'' |\log \varepsilon|$. Thus, for the same accuracy we expect that the Laplace transform-based method is much more efficient. Furthermore, as has been noted in [33, 30, 32] and several others, the system of equations can be solved independently for each shift and parallel computing resources can be leveraged for efficient solution of the Laplace transform since the computations are embarrassingly parallel.

Efficient methods have been proposed for solving the sequence of shifted systems arising from the discretization of the contour integral using the Trapezoidal rule. In [9], the authors propose a Krylov subspace method for solving the shifted system of equations. However, there is a rich literature concerning the Krylov subspace methods for shifted systems (for a good review see [26]). In practice, efficient preconditioners are needed to ensure convergence in a reasonable number of iterations. For this purpose, in this paper, we adapt the flexible Krylov subspace methods proposed in [6] and analyzed in [22].

3 Fast solvers for shifted systems

Recall the approximation of the inverse Laplace transform as shown in Equation (8). It leads to a sequence of shifted systems at each time step, which can then each be solved independently. In particular, for each time step, we need to solve two shifted systems of equations,

$$(K + z_k M) X_k = [b, \quad M \phi_0], \quad k = 1, \dots, N_z/2 \quad (10)$$

where z_k are the (complex) shifts corresponding to the inverse Laplace transform contour. $F(z_k)$ is thus found by adding the solution to the first system of shifted systems to the second (times the corresponding $\hat{q}(z_k)$). In this section, we consider Krylov subspace methods for the solution of the first shifted systems of equations,

$$(K + z_k M) x_k = b, \quad k = 1, \dots, N_z/2 \quad (11)$$

It should be noted that the second system of equations can be solved for in a similar fashion. We assume that the systems are non-singular for all $k = 1, \dots, N_z/2$. Recall that because of Equation (9) only half the systems need to be solved. In Equation (11), the matrices K and M and right-hand side b are independent of the shifts z_k . Krylov based methods are particularly attractive for this class of problems because of their shift-invariant property. The main strategy of Krylov solvers is to build a Krylov basis and then to search for an approximate solution in the resulting reduced space. The expensive step of constructing the Krylov basis needs to be performed only once since an approximation space is generated independently of the shifts. Once the basis is built, the subproblem for each shift can then be solved at a much reduced cost. We consider Arnoldi-based Krylov solvers, in particular the Full Orthogonalization Method (FOM) and the General Minimum RESidual Method (GMRES) [21].

The number of iterations required by the Krylov subspace solvers for convergence can be quite large, particularly for problems arising from realistic applications. To reduce the number of iterations, we use right preconditioners of the form $K_\tau \stackrel{\text{def}}{=} (K + \tau M)$ factorized and inverted using a direct solver. For preconditioners of this form and for $k = 1 \dots N_z/2$ it can be readily verified that,

$$(K + z_k M)(K + \tau M)^{-1} = I + (z_k - \tau)M(K + \tau M)^{-1} \quad (12)$$

It can be shown that $\mathcal{K}_m(MK_\tau^{-1}, b)$ is equivalent to the Krylov subspace $\mathcal{K}_m((K + z_k M)(K + \tau M)^{-1}, b)$. This property is known as the shift-invariant property. The algorithm to solve multiple shifted systems using a single shift-and-invert preconditioner is as follows: a basis for the Krylov subspace $\mathcal{K}_m(MK_\tau^{-1}, b)$ is obtained by running m steps of the Arnoldi algorithm on the matrix MK_τ^{-1} with the starting vector b . At the end of m steps, the following Arnoldi relationship holds,

$$(K + z_k M) K_\tau^{-1} V_m = V_{m+1} \underbrace{\left(\begin{bmatrix} I \\ 0 \end{bmatrix} + (z_k - \tau) \bar{H}_m \right)}_{\stackrel{\text{def}}{=} \bar{H}_m(z_k; \tau)} = V_{m+1} \bar{H}_m(z_k; \tau) \quad (13)$$

where $V_{m+1} = [v_1, \dots, v_{m+1}]$ is the Krylov basis and $\bar{H}_m$ is a $(m+1) \times m$ matrix formed by the Arnoldi process. We use the notation H_m to denote the corresponding upper Hessenberg matrix, i.e. $H_m = [I, 0] \bar{H}_m$. Having constructed the Arnoldi basis, the solution for each shift is obtained by searching for solutions $x_m(z_k) \in \text{span}\{K_\tau^{-1} V_m\}$, written as $x(z_k) = K_\tau^{-1} V_m y_m(z_k)$ where the vectors $y_m(z_k)$ can be obtained by either by imposing a Petrov-Galerkin condition, i.e. $r_m(z_k) \perp V_m$ which leads to the Full Orthogonalization Method for shifted systems (FOM-Sh) or by minimizing the residual norm $\|r_m(z_k)\|_2$ over all possible vectors in the span of $K_\tau^{-1} V_m$ which leads to the GMRES method for shifted systems (GMRES-Sh). The residual is defined as

$$r_m(z_k) \stackrel{\text{def}}{=} b - (K + z_k M) x_m = V_{m+1} (\beta e_1 - \bar{H}_m(z_k; \tau) y_m(z_k))$$

Numerical evidence in [22] showed that a single preconditioner might not be sufficient to effectively precondition all systems, especially if the values of z_k are spread out relative to the spectrum of the matrices. This is because a single preconditioner only adequately preconditions systems with z_k close to τ . In order to improve the convergence rate across all the shifts, we proposed a Flexible GMRES/FOM solver [22] for shifted systems that allows us the flexibility to choose different preconditioners (of the shift-and-invert type) at each iteration. This requires storing an additional set of vectors $U_m = [u_1, \dots, u_m]$. Consider the use of a preconditioner at each iteration of the form $K + \tau_j M$ where $j = 1, \dots, m$ and j is the iteration number. Systems with shifts z_k close to τ_j converge faster. The procedure to solve the shifted systems of equations follows similarly to the single preconditioner case with the following modified Arnoldi relationship [22],

$$(K + z_k M) U_m = V_{m+1} \underbrace{\left(\begin{bmatrix} I \\ 0 \end{bmatrix} + \bar{H}_m(z_k I_m - T_m) \right)}_{\stackrel{\text{def}}{=} \bar{H}(z_k; T_m)} = V_{m+1} \bar{H}_m(z_k; T_m) \quad (14)$$

where $T_m = \text{diag}\{\tau_1, \dots, \tau_m\}$. The columns of the matrix U_m now no longer span a Krylov subspace but rather a rational Krylov subspace. As in the case of a single preconditioner K_τ^{-1} , the approximate solution is constructed $x_m(z_k) = U_m y_m(z_k)$, so that $x_m(z_k) \in \text{span}\{U_m\}$. Again, the coefficients $y_m(z_k)$ are either computed by using a Petrov-Galerkin condition, leading to the flexible FOM for shifted systems (FFOM-Sh) or by minimizing the residual, leading to flexible GMRES for shifted systems (FGMRES-Sh). The expensive step of computing the matrix U_m is shared across all the systems and the cost of solving smaller subproblem (either by a Petrov-Galerkin projection or a residual minimization) for each shift is negligible compared to the cost of generating U_m . The algorithm is presented in detail in Algorithm 1.

In Algorithm 1, the use of a different preconditioner shift at each iteration would be too expensive since a different preconditioner $(K + \tau_j M)$ has to be factorized at each iteration. Numerical evidence shows that this is unnecessary since we need only pick a few systems with shifts τ such that they effectively precondition the systems over the entire range of shifts. The rule of thumb is that systems with shifts z_k close to τ converge first. In our previous work [22], we considered the solution of shifted systems with 200 shifts that were on the pure imaginary axis. We observed that the systems with shifts closer to the origin converged more slowly. Therefore, to speed convergence, we chose 5 preconditioners with shifts τ that were evenly spaced on a log scale. The choice of shifts and number of preconditioners is entirely application dependent. For the case of the shifts determined by the modified Talbot contours as shown in Figure 1, numerical experiments in Section 4 show the use of two preconditioners produces more favorable

Algorithm 1 Flexible FOM/GMRES for shifted systems

```
1: Given  $K, M, b, \{\tau_1, \dots, \tau_m\}, \{z_1, \dots, z_{N_z/2}\}$  and a tolerance  $tol$ ,
2:  $v_1 = b/\beta$ ,  $k = 1$  and  $\beta \stackrel{\text{def}}{=} \|b\|_2$ 
3: for all  $j = 1, \dots, \maxit$  do
4:   Solve  $(K + \tau_j M)u_j = v_j$ 
5:    $s_j := Mu_j$ 
6:   for all  $i = 1, \dots, j$  do
7:      $h_{ij} := s_j^* v_i$ 
8:     Compute  $s_j := s_j - h_{ij}v_i$ 
9:   end for
10:   $h_{j+1,j} := \|s_j\|_2$ . If  $h_{j+1,j} = 0$  stop
11:   $v_{j+1} = s_j/h_{j+1,j}$ 
12:  Define  $U_j = [u_1, \dots, u_j]$ ,  $V_{j+1} = [v_1, \dots, v_{j+1}]$  and  $\bar{H}_j = \{h_{i,l}\}_{1 \leq i \leq j+1, 1 \leq l \leq j}$ 
13:  Construct  $T_j = \text{diag}\{\tau_1, \dots, \tau_j\}$ 
14:  for all  $k = 1, \dots, N_z/2$  do
15:    If, not converged
16:    Construct  $\bar{H}_j(z_k; T_j) = \begin{bmatrix} I \\ 0 \end{bmatrix} + \bar{H}_j(z_k I_j - T_j)$ 
17:    FOM:
18:       $H_j(z_k; T_j)y_j^{fom} = \beta e_1$ , where  $H_j(z_k; T_j) = [I, 0]\bar{H}_j(z_k; T_j)$ 
19:    GMRES:
20:       $y_j^{gm} \stackrel{\text{def}}{=} \min_{y_j \in \mathbb{C}^j} \|\beta e_1 - \bar{H}_j(z_k; T_j)y_j\|_2$ 
21:    Construct the approximate solution as  $x_j(z_k) = U_j y_j(z_k)$ 
22:  end for
23:  If all systems have converged, exit
24: end for
```

results compared to only a single preconditioner. The shifts are chosen as follows: the first shift corresponds to z_k with the smallest real part called $\bar{\tau}_1$, and the second shift corresponds to the shift with the largest real part called $\bar{\tau}_2$. The sequence τ_j for $j = 1, \dots, m$ is constructed as follows. We choose $\tau_j = \bar{\tau}_1$ for $j = 1, 2, 3$ and $\tau_j = \bar{\tau}_2$ for $j = 4, 5$. This procedure is repeated until all the systems converge to the desired user-defined tolerance. This choice of preconditioners results in fast convergence for the range of parameters we have explored. Future work will focus on the optimal choice of preconditioners.

3.1 Properties of approximation space

In this section, we analyse the convergence of the flexible Krylov subspace approximation to the discretized representation of the contour integral. The error of the Krylov subspace approach in Equation (14) is defined by [19, 17] as

$$e_m \stackrel{\text{def}}{=} \sum_{k=1}^{N_z} w_k ((K + z_k M)^{-1} b - x_m) = \sum_{k=1}^{N_z} w_k ((K + z_k M)^{-1} b - U_m H_m^{-1}(z_k; T_m) \beta e_1) \quad (15)$$

where $x_m = U_m H_m^{-1}(z_k; T_m) \beta e_1$ is the approximate projected solution. Similarly, we can define the residual vector

$$r_m \stackrel{\text{def}}{=} \sum_{k=1}^{N_z} w_k (b - (K + z_k M)x_m) = \sum_{k=1}^{N_z} w_k (b - (K + z_k M)U_m H_m^{-1}(z_k; T_m) \beta e_1) \quad (16)$$

It can be readily verified that the approximate solution $x_m \stackrel{\text{def}}{=} \sum_k w_k U_m H_m^{-1}(z_k; T_m) \beta e_1$ lies in the subspace $\text{span}\{U_m\}$ where

$$\text{span}\{U_m\} = \text{span}\{v_1, T_1 v_1, \dots, (T_{m-1} \cdots T_1) v_1\}, \text{ with } T_k = (K + \tau_k M)^{-1} \quad (17)$$

To simplify the analysis, let us make the following change of variables: $b \leftarrow M^{-1/2}b$, $A \leftarrow M^{-1/2}KM^{-1/2}$ and $x \leftarrow M^{1/2}x$. This variable change is possible since M is a mass matrix and therefore, positive definite.

The projective space U_m is independent of the shift z_k and is not a typical Krylov subspace but rather a rational Krylov subspace. Krylov subspaces generate sequences of vectors that are of polynomial form, i.e., $v_{m+1} = p_m(A)v_1$ where $p(\cdot)$ is a polynomial. On the other hand, rational Krylov subspace generates iterates that can be expressed as rational functions of the matrix, i.e. $v_{m+1} = y_{m-1}(A)z_m(A)^{-1}v_1$, where y_{m-1} is a polynomial of degree at most $m-1$ and z_m is a fixed polynomial of degree m . The rational Krylov method was originally proposed by Ruhe [20] in the context of approximating interior eigenvalues, with which appropriately chosen shifts could potentially accelerate convergence of the eigenvalues in the desired region of the spectrum. They have been quite successful in the context of model reduction, in order to approximate the transfer function for a dynamical system [3]. The rational approach requires either knowledge of the shifts τ_k a priori, or this needs to be computed on-the-fly. In this work, our choice of shifts is done a priori and is based on heuristics that perform well in the applications, as we will demonstrate. In the context of $\mathcal{H}_2$ -optimality reduction, an automated choice of shifts has been proposed in [3]. The following result characterizes the properties of the error due to the flexible Krylov approach and derives an expression for the rational function that defines the rational Krylov subspace.

Proposition 1. *The error e_m satisfies*

$$e_m \in \text{span} \{ (A + z_1 I)^{-1} v_{m+1}, \dots, (A + z_{N_z} I)^{-1} v_{m+1} \} \quad (18)$$

and can be expressed in terms of a rational function as

$$e_m = r(A)v_1 \quad r(\lambda) \stackrel{\text{def}}{=} \sum_{k=1}^{N_z} \frac{w_k}{\lambda + z_k} \frac{\det(G_m - \lambda H_m)}{\prod_{j=1}^m (\lambda + \tau_j)} \quad (19)$$

where $G_m = I - H_m T_m$. Further, the error can be bounded as

$$\|e_m\|_2 \leq \sum_{k=1}^{N_z/2} |w_k| \| (A + z_k I)^{-1} v_{m+1} \|_2 |\eta_k| \quad (20)$$

where $\eta_k \stackrel{\text{def}}{=} h_{m+1,m}(z_k - \tau_m) e_m^ H_m^{-1}(z_k; T_m) \beta e_1$.*

Proof. Let us begin by writing

$$e_m = \sum_{k=1}^{N_z} w_k (A + z_k I)^{-1} (b - (A + z_k I) U_m H_m^{-1}(z_k; T_m) \beta e_1)$$

Using the modified Arnoldi relation in equation (14)

$$b - (A + z_k I) U_m H_m^{-1}(z_k; T_m) \beta e_1 = b - V_{m+1} \bar{H}_m(z_k; T_m) H_m(z_k; T_m)^{-1} \beta e_1 = -v_{m+1} \eta_k \quad (21)$$

since $b = V_m \beta e_1$ and η_k as defined above. Therefore, we have

$$e_m = - \sum_{k=1}^{N_z} w_k \eta_k (A + z_k I)^{-1} v_{m+1}$$

This proves the part of the proposition in Equation (18) the Arnoldi relations $U_m = V_{m+1} \bar{H}_m$ and $AU_m + U_m T_m = V_m$. Eliminating U_m we have $AV_{m+1} \bar{H}_m + V_{m+1} \bar{H}_m T_m = V_m$. The general recurrence for the vector v_m can be written as

$$(A + \tau_m I) v_{m+1} h_{m+1,m} = v_j - \sum_{j=1}^m (A v_j h_{jm} + v_j h_{jm} \tau_m)$$

The individual vectors $(A + z_k I)^{-1} v_{m+1}$ can be expressed in terms of rational functions and convergence this yields $h_{m+1,m} = 0$ following the arguments of [20] we get that

$$v_{m+1} = r(A)v_1 \quad \text{where} \quad r(\lambda) \stackrel{\text{def}}{=} \frac{I - H_m T_m - \lambda H_m}{(\lambda + \tau_1) \cdots (\lambda + \tau_m)}$$

Plugging in this expression into Equation (19) we get the desired result. The inequality for the error follows from the properties of vector norms. $\square$

We also have that the residual satisfies $r_m = -\sum_{k=1}^{N_z} w_k \eta_k v_{m+1}$ and therefore, $\|r_m\|_2 \leq \sum_k |w_k| |\eta_k|$. The quantity $|\eta_k| = |h_{m+1,m}(z_k - \tau_m) e_m^* H_m^{-1}(z_k; T_m) \beta e_1|$ is established as an a posteriori estimate of error of the flexible Krylov approach [22]. The result in Equation (19) suggests that an efficient choice of shifts τ_k for the preconditioners may result in a smaller error e_m after m iterations. In [22] analysis is provided for when the a posteriori measure is small and showed that it is related to the convergence of the eigenvalues of the matrix KM^{-1} (or alternatively, $M^{-1/2}KM^{-1/2}$). Although we do not have an automated procedure to determine the shifts, empirical evidence supports that our heuristic choice of preconditioner shifts τ_k leads to fast convergence for a wide range of parameters of the underlying PDE. This is discussed in the following section.

4 Numerical Experiments

4.1 Problem set-up

In this section we demonstrate using numerical examples the computational gains of the flexible Krylov solver and its robustness to changing various parameters in the forward problem. We consider a 2D depth-averaged aquifer with horizontal confining layers in a square domain satisfying the differential equation (2) with zero Dirichlet boundary conditions on all boundaries. The equations are discretized using standard linear finite elements implemented using FEniCS [14, 15, 16] using Python as the user-interface. The domain size is $[0, L]^2$ where $L = 100$ [m].

For our numerical experiments, we consider two transmissivity fields : the first is a random field generated using an exponential covariance kernel, i.e.

$$\kappa(x, y) = \theta \exp(-r) \quad (22)$$

with $r = \|x - y\|_2/L$. For the second transmissivity field we use a scaled version of Franke's function [4] which is a smooth function obtained by the weighted sum of four exponentials. The natural log of the fields are displayed in Figure 2. The exponential field is rougher than the Franke field and as a result we expect the iterative solver to work hard to converge to the desired tolerance. In both cases, the mean transmissivity was chosen to be $\mu_K = 10^{-4}$ [m²/s] and in these examples storativity was chosen to be constant with $S_s = 10^{-5}$ [-]. We consider one pumping source located at (50, 50) pumping at a constant rate of 0.85 L/s.

We compare the results of two different solvers - 'Single' corresponds to the Krylov solver with a single preconditioner $(K + \tau M)$ where, $\tau = \arg \min \text{Real}(\sigma_k)$ for $k = 1, \dots, N_z$, i.e. the shift corresponding to the smallest real part and 'Flexible' using the flexible Krylov approach described in section 3. The procedure to choose the shifts for the preconditioners is also described there.

4.2 Results

To demonstrate the robustness of our solver, we test it with respect to different parameters - variance of the field, number of times at which the solution needs to be computed and the number of grid points in the domain. First, we vary the variance of the log transmissivity field and note that higher variances correspond to more ill-conditioned systems. We report the number of iterations taken and the CPU time required for the solvers to converge to a relative tolerance of 10^{-10} . The number of shifts N_z was chosen to be 40 and therefore the number of systems required to be solved is $N_z/2 = 20$. The results are displayed in Table 1. The last column in the table is an estimate of the maximum condition number of $K + \sigma M$ across all shifts. This is a lower bound to the condition number in the 1-norm [7, 8]. We see that the fields corresponding to larger variances have more ill-conditioned systems. Since Franke field is smoother than the random field, it is expected that the resulting linear systems are less ill-conditioned and therefore, it would take fewer number of iterations to converge to the desired tolerance. Furthermore, the number of iterations required by both 'Single' and 'Flexible' preconditioning increases with increasing variance, however the number of iterations using 'Flexible' preconditioning does not

Figure 2: Example log transmissivity fields generated by (left) sample drawn from a random field with exponential covariance kernel and (right) scaled version of Franke’s function. The variance for both fields in the figure is 1.6.

Figure 3: Drawdown curve measured at location (70, 70) with the source at (50, 50) pumping at a constant rate of 0.85 L/s. The squares indicate sampling times 1, 5 and 20 min that were used in the results in Table 2.

increase significantly. By contrast, the use of ‘Single’ preconditioning is insufficient because it does not adequately precondition all the systems and the cost can grow significantly.

Franke Field					
Variance	Single		Flexible		κ_{pseudo}
	Iter.	CPU Time [s]	Iter.	CPU Time [s]	
0.8	73	8.6	40	4.6	1.2×10^6
1.6	95	13.0	48	5.8	1.5×10^7
3.5	125	21.3	54	7.1	1.3×10^9
Random Field					
	Iter.	CPU Time [s]	Iter.	CPU Time [s]	
0.8	89	11.6	45	5.5	1.1×10^6
1.6	126	20.1	53	6.8	9.9×10^6
3.5	139	28.9	54	7.3	9.8×10^8

Table 1: Iteration count and CPU time for solving the system of equations (11). The solution was computed at measurement time = 5 min. Two different log transmissivity fields were considered. An estimate for condition number shows that the fields with higher variance corresponded to larger condition number. The proposed ‘Flexible’ solver performs better than ‘Single’, both in terms of iterations and CPU time.

Next, we consider the performance of the solvers in computing the solution at three different measurement times. The number of iterations and the time required by the solver for solving for the hydraulic head at various times is computed and displayed in Table 2. It is observed that the number of iterations required for the systems corresponding to the early times to converge are larger, even though the condition number decreases. The reason for the increase in iterations can be explained by examining the contours obtained from different measurement times (Figure 1). It is the presence of the scaling factor $1/t$ multiplying the shifts z_k that cluster the shifts for larger times, see Equation (7). The clustering of the shifts allows for faster convergence because it allows for better preconditioning (both ‘Single’ and ‘Flexible’) of all the systems since the systems with shifts closer to τ will converge more rapidly than those further away. This is consistent with the analysis done in [22].

Franke Field					
Meas. Time [min]	Single		Flexible		κ_{pseudo}
	Iter.	CPU Time [s]	Iter.	CPU Time [s]	
1	134	22.6	55	7.8	4.3×10^6
5	95	13.0	48	5.8	1.5×10^7
20	58	5.8	33	3.4	5.6×10^7
Random Field					
	Iter.	CPU Time [s]	Iter.	CPU Time [s]	
1	145	26.2	59	8.3	3.0×10^6
5	126	20.1	53	6.8	9.9×10^6
20	92	12	45	5.4	3.8×10^7

Table 2: Iteration count and time for solving the system of equations (11). Two different log transmissivity fields were considered and the solution was computed at different measurement times. Here, variance of the log transmissivity was fixed to be 1.6. An estimate for condition number shows that the systems at earlier measurement time corresponded to larger condition number. The proposed ‘Flexible’ solver performs better than ‘Single’, both in terms of iterations and CPU time.

Thus far, we have only looked at solving for each time point independently, however our solver allows for simultaneously solving for the solutions at multiple times. We consider 40 uniformly spaced time points between 40 min and 60 min. These measurements correspond to ‘late-time’ measurements (see Figure 3). These systems are solved simultaneously using ‘Flexible’ preconditioning. To pick the preconditioner shifts, we first sort all the shifts arising from the discretized contour from the 40 time instances by their real parts and we choose two preconditioners τ_1 corresponding to the overall shift with the largest real part and τ_2 corresponding to the overall shift with the smallest real part. We observe that the maximum number of iterations was 27 across all shifts and all the 40 time sample points, which is roughly the same cost of solving for one single time, i.e., the computational cost associated with solving the problem for multiple times is not significantly higher than the cost of solving the system of equations for a single time. More investigation is needed to extend the solvers to be able to compute the entire time history simultaneously.

Grid size	Single		Flexible	
	Iter.	Time [s]	Iter.	Time [s]
41^2	96	3.6	49	1.5
101^2	95	13.0	48	5.8
201^2	94	52.8	45	25.1
301^2	92	260.2	44	116.2
401^2	91	611.7	44	268.1

Table 3: Iteration count and time respectively for computing pressure field using increasing number of grid points. The log transmissivity field is Franke field with a variance of 1.6.

All of the above experiments have been conducted on a grid of size 101^2 . To show the

independence of the iteration count with respect to the grid size, we consider grids ranging from 41^2 to 401^2 . In all the experiments, the log transmissivity field is the Franke field with a variance of 1.6. Table 3 lists the iteration count of the solver with increasing grid size. For both the solvers ‘Single’ and ‘Flexible’, the iteration count is independent of the number of grid points.

In summary, we can conclude that our solver is robust to various parameters such as field smoothness and variance, as well as the computation of the solutions at different measured times. Additionally, for ‘late enough’ times, the proposed flexible solver further improves efficiency as it allows for the solution of the resulting systems of equations simultaneously. Therefore, this solver can be used within the framework of inverse problems to achieve significant speedup.

5 Inverse Problem

5.1 Geostatistical approach

We consider the geostatistical approach as described in [12, 13], which is one of the prevalent approaches for solving the inverse problem. The objective is to determine aquifer properties (here log transmissivity) given discrete head measurements. In the geostatistical approach the unknowns are modeled as Gaussian random fields and the Bayesian approach is used to infer the posterior probability density function of the unknowns as the product of two parts: the likelihood of the measured data or the ‘data misfit’ and the prior distribution of the parameters which represents the assumed spatial correlation of the parameters. Denote by $s(x) \in \mathbb{R}^{N_s}$ the vector corresponding to the discretization of the function to be estimated with $s \sim \mathcal{N}(X\beta, Q)$ where X is a $N_s \times p$ known matrix, β are p unknown drift coefficients and Q is a covariance matrix with entries $Q_{ij} = \kappa(x_i, x_j)$ where κ is a generalized covariance kernel. The covariance kernel contains information about the degree of smoothness of the random field and the correlation between two points. Popular choices of covariance kernels include the Matérn covariance family, which includes the exponential and Gaussian covariance kernels [34]. The measurement equation is expressed as,

$$y = h(s) + \epsilon \quad \epsilon \sim \mathcal{N}(0, R) \quad (23)$$

where $y \in \mathbb{R}^{N_y}$ denotes the hydraulic head measurements and ϵ represents the measurement error which includes both the error resulting from data collection and the errors in evaluating $h(s)$. The matrices Q, R and X are structural parameters whose values can be optimized using a restricted maximum likelihood approach. More details on the choice of these parameters can be found in [12]. The best estimate is obtained by computing the maximum a posteriori estimate (MAP). This is equivalent to solving the following optimization problem,

$$\arg \min_{\hat{s}, \hat{\beta}} \frac{1}{2} \|y - h(s)\|_{R^{-1}}^2 + \frac{1}{2} \|s - X\beta\|_{Q^{-1}}^2$$

where, the objective function to be minimized is the negative logarithm of the posterior probability density function $p(s, \beta|y)$. To solve the nonlinear optimization problem, Gauss-Newton algorithm is used. The procedure is described in Algorithm 2.

Every iteration in algorithm 2 requires the computation matrices $J_k Q J_k^T$ and $Q J_k^T$. Storing the dense covariance matrix Q while computing $Q J_k^T$ can be expensive, both in terms of memory costs for storage and computational costs, particularly when the number of grid points is large. For covariance matrices that are translation invariant or stationary, the associated cost for storing Q and computing matrix vector products $Q J_k^T$ can be reduced. A Fast Fourier Transform (FFT) method can be used if the grid is regular. If the grid is not regular, the Hierarchical matrix approach can be used [23]. The resulting computations are $\mathcal{O}(N_y N_s \log N_s)$ instead of $\mathcal{O}(N_y N_s^2)$.

5.2 Sensitivity computation

Another step that can be very expensive in the geostatistical method for inversion is the computation of the Jacobian, or sensitivity matrix. When traditional time stepping algorithms are

Algorithm 2 Quasi-linear Geostatistical approach

while not converged **do**

 Compute the $N_y \times N_s$ Jacobian J as,

$$J_k = \left. \frac{\partial h}{\partial s} \right|_{s=s_k} \quad (24)$$

 Solve the following system of equations,

$$\begin{pmatrix} J_k Q J_k^T + R & J_k X \\ (J_k X)^T & 0 \end{pmatrix} \begin{pmatrix} \xi_{k+1} \\ \beta_{k+1} \end{pmatrix} = \begin{pmatrix} y - h(s_k) + J_k s_k \\ 0 \end{pmatrix} \quad (25)$$

s_{k+1} is computed by,

$$s_{k+1} = X \beta_{k+1} + Q J_k^T \xi_{k+1} \quad (26)$$

end while

used for simulating the time-dependent forward problem, computing the sensitivity of the measurement operator with respect to parameters to be inverted for is accomplished by the adjoint state method. This involves cross-correlation of two fields at the same time, one obtained by forward recursion and the other by backward recursion. Typically, the forward recursion is performed first, however the entire time history up to the desired time must be accessible during the backward recursion. For small scale problems, one typically stores the entire time history. However, for large-scale problems arising from finely discretized problems the storage requirements may be so large that it may not be able to be stored in RAM so one must resort to disk storage, in which case memory access then becomes the limiting factor. A standard approach is to use checkpointing [5, 27, 31] which trades computational complexity (by a factor logarithmic in number of steps) while reducing the memory complexity to a few state buffers (also logarithmic in number of steps).

Using the Laplace transform approach to compute the solution of the forward problem at time t avoids the expensive computation of the time history. In this section, we derive expressions for computing the sensitivity of the measurements with respect to the parameters also using the Laplace transform. In addition to the solution of the forward problem $\phi(x, t)$, the algorithm requires computing an adjoint field $\psi(x, t)$ solving for which also leads to a shifted system of equations, and can be efficiently done using the flexible Krylov subspace method described in section 3.

We now derive expressions for the sensitivity. Let S_s be parametrized by s_S and κ be parametrized by s_K . Since we want the reconstruction of S_s and κ to be positive, it is common to consider a log-transformation $S_s = e^{s_S}$ and $\kappa = e^{s_K}$. Consider the functional that we would like to compute the sensitivity of

$$I(t; s_K, s_S) \stackrel{\text{def}}{=} \int_{\Omega} \delta(x - x_m) \phi(x, t) dx$$

which corresponds to a point measurement at the measurement location. We are interested in computing the sensitivity of the functional I with respect to the parameters s_K and s_S . We start by making the following transformation

$$I(t; s_K, s_S) = \frac{1}{2\pi i} \int_{\Omega} \int_{\Gamma} e^{zt} \delta(x - x_m) \hat{\phi}(x, z) dz dx$$

To compute the variation δI with respect to the functions s_K and s_S , we use the standard adjoint field approach. First, take the Laplace transform of Equations (2) and multiply throughout by test functions $\hat{\psi}(x, z)$, then integrate by parts and apply appropriate boundary conditions to get,

$$\int_{\Omega} e^{s_K} \nabla \hat{\phi} \cdot \nabla \hat{\psi} dx + z \int_{\partial\Omega_w} S_y \hat{\phi} \hat{\psi} dx + z \int_{\Omega} e^{s_S} \hat{\phi} \hat{\psi} dx = \int_{\Omega} \hat{q} \hat{\psi} dx$$

Taking the variation $\delta\hat{\phi}$, we have

$$\begin{aligned} & \int_{\Omega} e^{s_K} \nabla \delta\hat{\phi} \cdot \nabla \hat{\psi} dx + z \int_{\partial\Omega_w} S_y \delta\hat{\phi} \hat{\psi} dx + z \int_{\Omega} e^{s_S} \delta\hat{\phi} \hat{\psi} dx \\ &= - \int_{\Omega} \left(e^{s_K} \delta s_K \nabla \hat{\phi} \cdot \nabla \hat{\psi} + z e^{s_S} \delta s_S \hat{\phi} \cdot \nabla \hat{\psi} \right) dx \end{aligned} \quad (27)$$

The adjoint field $\hat{\psi}$ is chosen to satisfy the following set of differential equations, which is similar to the Laplace transformed version of Equations (2) with the forcing term corresponding to a measurement operator at the measurement location x_m .

$$\begin{aligned} -\nabla \cdot \left(e^{s_K} \nabla \hat{\psi}(x, z) \right) + z e^{s_S} \hat{\psi}(x, z) &= -\delta(x - x_m) & x \in \Omega \\ \hat{\psi}(x, z) &= 0, & x \in \partial\Omega_D \\ \nabla \hat{\psi}(x, z) \cdot n &= 0, & x \in \partial\Omega_N \\ e^{s_K} \nabla \hat{\psi}(x, z) \cdot n &= -z S_y \hat{\psi}, & x \in \partial\Omega_w \end{aligned} \quad (28)$$

Similarly, we multiply equation (28) by $\delta\hat{\phi}$ and integrate by parts to obtain,

$$\int_{\Omega} e^{s_K} \nabla \delta\hat{\phi} \cdot \nabla \hat{\psi} dx + z \int_{\partial\Omega_w} S_y \delta\hat{\phi} \hat{\psi} dx + z \int_{\Omega} e^{s_S} \delta\hat{\phi} \hat{\psi} dx = \int_{\Omega} \delta(x - x_m) \delta\hat{\phi}(x, z) dx$$

Equating the right hand sides from the Equation (27) and the equation above, we get

$$\delta I = \frac{1}{2\pi i} \int_{\Gamma} \int_{\Omega} \delta(x - x_m) \delta\hat{\phi}(x, z) dz dx \quad (29)$$

$$= \frac{1}{2\pi i} \int_{\Gamma} \int_{\Omega} e^{zt} \left(e^{s_K} \delta s_K \nabla \hat{\phi} \cdot \nabla \hat{\psi} + z e^{s_S} \delta s_S \hat{\phi} \cdot \hat{\psi} \right) dz dx \quad (30)$$

Finally, the integral w.r.t z can be discretized using the same contour integral approach described in Section 2.2

$$\delta I_{N_z} = \sum_{k=1}^{N_z} w_k \int_{\Omega} \left(e^{s_K} \delta s_K \nabla \hat{\phi}_k \cdot \nabla \hat{\psi}_k + e^{s_S} \delta s_S \hat{\phi}_k \cdot \hat{\psi}_k \right) dx \quad (31)$$

As an example of this calculation, we plot the sensitivity at $t = 5$ [min] (made dimensionless) with respect to log transmissivity. The field here was chosen to be a constant field, hence the symmetry observed between the measurement location (40, 50) and the point source (50, 50). The results are shown for time $t = 5$ [min]. This is displayed in Figure 4.

5.3 Application: Transient Hydraulic Tomography

The objective is to reconstruct log transmissivity of an aquifer given measurements of the hydraulic head at several observation locations. The governing equations are the groundwater equations for a 2D confined aquifer, as described in Section 4. The ‘true’ field is chosen to be that generated by the exponential kernel as shown in Figure 2 also provided in Equation (22). The parameters R and X are chosen to be $R = 10^{-7}I$ and $X = [1, \dots, 1]^T$ respectively. We do not investigate the optimality of these parameters, i.e. those that yield the best reconstruction, since the goal was to demonstrate the performance of our proposed solver in solving inverse problems. We additionally added model error by generating measurements using Crank-Nicolson to solve the governing equations using the ‘true’ log transmissivity field. We have three measurement points: at $t = 8, 10$ and 20 [min] and 36 measurement locations spaced evenly in the square $[20, 80] \times [20, 80]$. We introduce a 2% error in the measurements. The inverse problem is solved using these measurements. Storativity was held constant at $S_s = 10^{-5}$ [-] and the domain was a square domain of length $L = 100$ [m]. The single pumping source, located at (50, 50), has a pumping rate of 0.85 L/s. We use the flexible Krylov solver with $N_z = 20$ for solving the forward and adjoint problems on a problem with 101×101 grid points. The results are shown in Figure 5. The error in the relative L_2 norm is 0.16 within the area of measurements, i.e. the $[20, 80] \times [20, 80]$ m² box, and the total error in the relative L_2 norm is 0.36 for the whole aquifer.

Figure 4: Sensitivity (made dimensionless) with respect to log transmissivity. The field here was chosen to be a constant field, hence the symmetry observed between the measurement location (40, 50) and the point source (50, 50).

Figure 5: True transmissivity field (left) and the estimates transmissivity field (right). The number of grid points is 101, 124. The error in the relative L_2 norm is 0.11 within the area of measurements, i.e. the $[20, 80] \times [20, 80]$ m² box, and the total error in the relative L_2 norm is 0.35 for the whole aquifer.

Figure 6 shows the time required for the various solvers. ‘Direct’ refers to solving each of the shifted systems using a direct solver, and FGMRES-Sh is the solver using 2 preconditioners per time. N_z was chosen to be 20. Note that the slope of the graph corresponding to Crank-Nicolson is dependent on the particular application and on the time step used. The efficiency of Crank-Nicolson relative to FGMRES-Sh depends on the time at which the solution is to be evaluated as well as the size of the time steps taken. For our application, the speedup using FGMRES-Sh as opposed to Crank-Nicolson to solve the forward problem was significant, as demonstrated in the figure. The computational gain is more dramatic in the calculation of the Jacobian since this solution is required for multiple sources and receivers.

Let us assume that we have N_s sources, N_m receivers and collect N_T time measurements. The total number of systems that need to be solved are $N_s N_m N_T$. Assuming the application of the preconditioner can be represented by $\mu(N)$ and the number of iterations using the Krylov solvers are N_{iter} . The total cost of computing the Jacobian is $N_m N_s N_{iter} N_T \mu(N)$. Using a time-stepping scheme with N_t time steps (and therefore, the solution of N_t system of equations) similarly costs $N_m N_s N_t \mu(N)$. The costs are comparable when $N_t \sim N_{iter} N_T$. In Section 2.3 we saw that for increasing accuracy, the number of systems to be solved increase only logarithmically for Laplace transform-based methods whereas it has a square-root growth for time-stepping schemes. For a few time measurement points N_T or when the number of time steps required by the time-stepping schemes is large, we expect that the Laplace transform-based methods are

Figure 6: Comparison between Crank-Nicolson and the Laplace Transform method using direct solver and GMRES for shifted systems. Note that the slope of the line corresponding to Crank-Nicolson depends on the time step that is used.

much more efficient. This is also confirmed by Figure 6.

6 Conclusion and future work

We have proposed a fast method for solving large-scale nonlinear inverse problems based on parabolic PDE. While we have focused on THT, it can be applied to a general class of time-dependent PDEs, for which the Laplace transform can be applied. The resulting system of equations are solved efficiently using a flexible Krylov approach, previously used to accelerate oscillatory hydraulic tomography. For small number of measurement times, our solver is computationally more efficient than standard time-stepping schemes, especially when very small time steps required for stability. We have applied the solver to synthetic problems arising in THT. Since the computation of the Jacobian is often the bottleneck in solving large-scale inverse problems, our approach for computing the Jacobian based on the Laplace transform greatly improves the storage and computational cost.

There are several avenues for future work. In our work, the measurement times are considered known a priori and are few in number. However, the number of measurement times and the sampling rate can be chosen to better improve the accuracy of the reconstruction. For example, choosing the sampling points as the roots of Laguerre polynomials may be used to improve the accuracy of the integral corresponding to the data mismatch. Regarding the flexible Krylov solver, theoretic insight into the properties of the rational Krylov subspace generated U_m can better guide the user to pick preconditioners to efficiently precondition the shifted systems, including those corresponding to multiple times. It remains to be demonstrated, if based on our approach, we can develop a solver that can solve for the entire time history. Another computational bottleneck we would like to consider is the solution of the shifted system in Equation (11) with multiple right hand sides, corresponding to different source and measurement locations. One may either try block approaches, or recycling strategies for systems with multiple shifts and multiple right hand sides.

7 Acknowledgments

This work was supported by NSF award 1215742, “Collaborative Research: Fundamental Research on Oscillatory Flow in Hydrogeology.” We would also like to thank James Lambers and Anil Damle for their careful reading of the manuscript and useful suggestions.

References

- [1] D.A. Boas, D.H. Brooks, E.L. Miller, C.A. DiMarzio, M.E. Kilmer, R.J. Gaudette, and Q. Zhang. Imaging the body with diffuse optical tomography. *Signal Processing Magazine, IEEE*, 18(6):57–75, 2001.
- [2] M. Cardiff, W. Barrash, and P.K. Kitanidis. Hydraulic conductivity imaging from 3-D transient hydraulic tomography at several pumping/observation densities. *Water Resources Research*, 49(11):7311–7326, 2013.
- [3] Vladimir Druskin and Valeria Simoncini. Adaptive rational krylov subspaces for large-scale dynamical systems. *Systems & Control Letters*, 60(8):546–560, 2011.
- [4] R. Franke. A critical comparison of some methods for interpolation of scattered data. Technical report, DTIC Document, 1979.
- [5] A. Griewank and A. Walther. Algorithm 799: REVOLVE: an implementation of checkpointing for the reverse or adjoint mode of computational differentiation. *ACM Transactions on Mathematical Software (TOMS)*, 26(1):19–45, 2000.
- [6] G. Gu, X. Zhou, and L. Lin. A flexible preconditioned Arnoldi method for shifted linear systems. *Journal of Computational Mathematics - International Edition*, 25(5):522, 2007.
- [7] W.W. Hager. Condition estimates. *SIAM Journal on Scientific and Statistical Computing*, 5(2):311–316, 1984.
- [8] N.J. Higham and F. Tisseur. A block algorithm for matrix 1-norm estimation, with an application to 1-norm pseudospectra. *SIAM Journal on Matrix Analysis and Applications*, 21(4):1185–1201, 2000.
- [9] K.J. In’t Hout and J.A.C. Weideman. A contour integral method for the Black-Scholes and Heston equations. *SIAM Journal on Scientific Computing*, 33(2):763–785, 2011.
- [10] Jari Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*, volume 160. Springer Science & Business Media, 2006.
- [11] M.E. Kilmer and E. de Sturler. Recycling subspace information for diffuse optical tomography. *SIAM Journal on Scientific Computing*, 27(6):2140–2166, 2006.
- [12] P. K. Kitanidis. Quasilinear geostatistical theory for inversing. *Water Resour. Res.*, 31(10):2411–2419, 1995.
- [13] P. K. Kitanidis. *Bayesian and Geostatistical Approaches to Inverse Problems*, pages 71–85. John Wiley & Sons, Ltd, 2010.
- [14] A. Logg, K.A.M., G.N. Wells, et al. *Automated Solution of Differential Equations by the Finite Element Method*. Springer, 2012.
- [15] A. Logg and G.N. Wells. Dolfin: Automated finite element computing. *ACM Transactions on Mathematical Software*, 37(2), 2010.
- [16] A. Logg, G.N. Wells, and J. Hake. *DOLFIN: a C++/Python Finite Element Library*, chapter 10. Springer, 2012.
- [17] Luciano Lopez and Valeria Simoncini. Analysis of projection methods for rational function approximation to the matrix exponential. *SIAM Journal on Numerical Analysis*, 44(2):613–635, 2006.
- [18] Shlomo P Neuman. Theory of flow in unconfined aquifers considering delayed response of the water table. *Water Resources Research*, 8(4):1031–1045, 1972.
- [19] Marina Popolizio and V Simoncini. Acceleration techniques for approximating the matrix exponential operator. *SIAM Journal on Matrix Analysis and Applications*, 30(2):657–683, 2008.

- [20] A. Ruhe. Rational Krylov sequence methods for eigenvalue computation. *Linear Algebra and its Applications*, 58:391–405, 1984.
- [21] Yousef Saad. *Iterative methods for sparse linear systems*. Siam, 2003.
- [22] A.K. Saibaba, T. Bakhos, and P.K. Kitanidis. A flexible Krylov solver for shifted systems with application to oscillatory hydraulic tomography. *SIAM Journal on Scientific Computing*, 35(6):A3001–A3023, 2013.
- [23] A.K. Saibaba and P.K. Kitanidis. Efficient methods for large-scale linear inversion using a geostatistical approach. *Water Resources Research*, 48(5), 2012.
- [24] D. Sheen, I.H. Sloan, and V. Thomée. A parallel method for time-discretization of parabolic problems based on contour integral representation and quadrature. *Mathematics of Computation of the American Mathematical Society*, 69(229):177–195, 2000.
- [25] D Sheen, I.H. Sloan, and V. Thomée. A parallel method for time discretization of parabolic equations based on Laplace transformation and quadrature. *IMA Journal of Numerical Analysis*, 23(2):269–299, 2003.
- [26] V. Simoncini and D.B. Szyld. Recent computational developments in Krylov subspace methods for linear systems. *Numerical Linear Algebra with Applications*, 14(1):1–59, 2007.
- [27] W.W. Symes. Reverse time migration with optimal checkpointing. *Geophysics*, 72(5):SM213–SM221, 2007.
- [28] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. siam, 2005.
- [29] V. Thomée. A high order parallel method for time discretization of parabolic type equations based on Laplace transformation and quadrature. *Int. J. Numer. Anal. Model*, 2:121–139, 2005.
- [30] L.N. Trefethen, J.A.C. Weideman, and T. Schmelzer. Talbot quadratures and rational approximations. *BIT Numerical Mathematics*, 46(3):653–670, 2006.
- [31] Q. Wang, P. Moin, and G. Iaccarino. Minimal repetition dynamic checkpointing algorithm for unsteady adjoint calculation. *SIAM Journal on Scientific Computing*, 31(4):2549–2567, 2009.
- [32] J Weideman and L Trefethen. Parabolic and hyperbolic contours for computing the Bromwich integral. *Mathematics of Computation*, 76(259):1341–1356, 2007.
- [33] J.A.C. Weideman. Optimizing Talbot’s contours for the inversion of the Laplace transform. *SIAM Journal on Numerical Analysis*, 44(6):2342–2362, 2006.
- [34] C.K.I. Williams and C.E. Rasmussen. Gaussian processes for machine learning. *the MIT Press*, 2(3):4, 2006.
- [35] M Zaslavsky, V Druskin, A Abubakar, T Habashy, and V Simoncini. Large-scale Gauss-Newton inversion of transient CSEM data using the model reduction framework. *Geophysics*, 2012.