A Time-parallel Approach to Strong-constraint Four-dimensional Variational Data Assimilation

Vishwas Rao^a, Adrian Sandu^a

^a*Computational Science Laboratory, Department of Computer Science, Virginia Tech. Blacksburg, Virginia 24060*

Abstract

A parallel-in-time algorithm based on an augmented Lagrangian approach is proposed to solve four-dimensional variational (4D-Var) data assimilation problems. The assimilation window is divided into multiple sub-intervals that allows parallelization of cost function and gradient computations. The solution to the continuity equations across interval boundaries are added as constraints. The augmented Lagrangian approach leads to a different formulation of the variational data assimilation problem than the weakly constrained 4D-Var. A combination of serial and parallel 4D-Vars to increase performance is also explored. The methodology is illustrated on data assimilation problems involving the Lorenz-96 and the shallow water models.

1. Introduction

Predicting the behavior of complex dynamical systems, such as the atmosphere, requires using information from observations to decrease the uncertainty in the forecast. Data assimilation combines information from a numerical model, prior knowledge, and observations (all with associated errors) in order to obtain an improved estimate of the true state of the system. Data assimilation is an important application of data-driven application systems (DDDAS [4], or InfoSymbiotic systems) where measurements of the physical system are used to constrain simulation results. Two approaches to data assimilation have gained widespread popularity: variational and ensemble-based methods. The ensemble-based methods are rooted in statistical theory, whereas the variational approach is derived from optimal control theory. The variational approach formulates data assimilation as a nonlinear optimization problem constrained by a numerical model. The initial conditions (as well as boundary conditions, forcing, or model parameters) are adjusted to minimize the discrepancy between the model trajectory and a set of time-distributed observations. In real-time operational settings the data assimilation process is performed in cycles: observations within an assimilation window are used to obtain an optimal trajectory, which provides the initial condition for the next time window, and the process is repeated in the subsequent cycles. The variational methodology is widely adopted by national and international numerical weather forecast centers to provide the initial state for their forecast models [5].

The process of solving the nonlinear optimization to obtain the initial conditions of a 4D-Var is highly sequential since the minimization procedure consists of sequential iterations, each iteration requires at least one forward model integration followed by an integration of the adjoint model, and each integration performs a sequence of timesteps over the entire analysis window [18]. Parallel 4D-Var simulations currently rely on spatial decomposition. Trémolet and Le Dimet have shown how variational data assimilation can be used to couple models and to perform it in parallel in the spatial direction [37]. In [27], the author considers different data distribution strategies to perform parallel variational data assimilation in the spatial direction. An implementation of 4D-Var by exploiting spatial parallelism for chemical data assimilation has been described in [17]. A scalable approach for variational data assimilation is presented in [15]. The authors

Email addresses: visrao@vt.edu (Vishwas Rao), sandu@cs.vt.edu (Adrian Sandu)

achieve parallelism by dividing the 3D-Var into multiple local 3D-Var sub-problems. Multiple copies of the modified 3D-Var problem, which ensures feasibility at the boundaries of the sub-domains, are solved across processors and the global 3D-Var minimum is obtained by collecting the local minima.

The current generation of computers are becoming more parallel, while the individual processor speed is increasing slowly. As our understanding of the physics improves, the computer models become more complex. Space parallelism alone is likely to be insufficient to run 4D-Var efficiently on the new computer architectures. Unless advanced and scalable parallel algorithms to solve the 4D-Var are developed, determining initial state to forecast using 4D-Var might become infeasible in the foreseeable future.

Fisher [18] achieves time parallelism in the context of incremental weak-constraint 4D-Var. Weak-constraint 4D-Var removes the model constraints are replaced by penalties in the cost function. The incremental approach solves a linearized optimization problem in the inner loop. This linearized problem is equivalent to solving a saddlepoint system, and this solution can be efficiently time-parallelized [18]. Ulbrich [38] proposes an SQP approach where the time dimension is parallelized using a “parareal” approach. A different approach is to formulate the discrete optimality system and to solve it using parallel linear algebra: Biros and Ghattas use Krylov space schemes [8], while Borzi uses space-time multigrid schemes [10]. Navon et. al. have used the augmented Lagrangian approach to perform data assimilation in [20, 26]. However, to the best of our knowledge this is the first work that uses the augmented Lagrangian framework to propose a time-parallel implementation of strong-constraint 4D-Var. A conference version of the current paper was presented at ICCS 2015 [31].

The focus of this work is on developing a scalable time-parallel algorithm to solve the strong-constraint 4D-Var problem. The algorithm evaluates the important components of 4D-Var namely, cost function and gradient computations in parallel. This work addresses an important challenge associated with computing the solutions to 4D-Var problem: scalability in the temporal direction. The solution presented here is based on the augmented Lagrangian framework. The paper is organized as follows: Section 2 gives a brief description about data assimilation and 4D-Var. In Section 3 we show how the 4D-Var can be reformulated using augmented Lagrangian technique so that the gradient and cost function evaluations can be parallelized. The most important contribution of this work is that it achieves parallelism in temporal direction. This work can be easily extended to achieve parallelism in spatio-temporal directions. Section 4 gives a detailed description of the parallel algorithm. Section 5 shows the numerical experiments with a small but chaotic Lorenz-96 model and a relatively large shallow water model on a sphere and Section 6 gives concluding remarks and future directions.

2. Four-dimensional variational data assimilation

Data assimilation (DA) is the fusion of information from priors, imperfect model predictions, and noisy data, to obtain a consistent description of the true state $\mathbf{x}^{\text{true}}$ of a physical system [9, 14, 21, 33, 34]. The estimate that fuses all these sources of information is called the analysis $\mathbf{x}^{\text{a}}$, which under certain statistical assumptions of the model and errors is provably the best estimate of the true state.

The prior information encapsulates our current knowledge of the system. Usually the prior information is contained in a background estimate of the state $\mathbf{x}^{\text{b}}$ and the corresponding background error covariance matrix $\mathbf{B}$.

The model captures our knowledge about the physical laws that govern the evolution of the system. The model evolves an initial state $\mathbf{x}_0 \in \mathbb{R}^n$ at the initial time t_0 to future states $\mathbf{x}_k \in \mathbb{R}^n$ at future times t_k . A general model equation is represented as follows:

$$\mathbf{x}_k = \mathcal{M}_{t_0 \rightarrow t_k}(\mathbf{x}_0) . \quad (1)$$

Observations are noisy snapshots of reality available at discrete time instances. Specifically, measurements $\mathbf{y}_k \in \mathbb{R}^m$ of the physical state $\mathbf{x}^{\text{true}}(t_k)$ are taken at times t_k , $k = 1, \dots, N$. The model state is related to observations by the following relation:

$$\begin{aligned} \mathbf{y}_k &= \mathcal{H}(\mathbf{x}_k) - \varepsilon_k^{\text{obs}}, \quad k = 1, \dots, N, \\ \varepsilon_k^{\text{obs}} &= \varepsilon_k^{\text{representativeness}} + \varepsilon_k^{\text{measurement}} . \end{aligned} \quad (2)$$

The observation operator $\mathcal{H}$ maps the model state space onto the observation space. The observation error term ($\varepsilon_k^{\text{obs}}$) accounts for both measurement and representativeness errors. Measurement errors are due to imperfect sensors. The representativeness errors are due to the inaccuracies of the mathematical and numerical approximations inherent to the model.

Variational methods solve the data assimilation problem in an optimal control framework, where one finds the control variable which minimizes the mismatch between the model forecasts and the observations. Strong-constraint 4D-Var assumes that the model (1) is perfect [33, 34]. The control parameters are the initial conditions $\mathbf{x}_0$, which uniquely determine the state of the system at all future times via the model equation (1). The background state is the prior best estimate of the initial conditions $\mathbf{x}_0^b$, and has an associated initial background error covariance matrix $\mathbf{B}_0$. Observations $\mathbf{y}_k$ at t_k have the corresponding observation error covariance matrices $\mathbf{R}_k$, $k = 1, \dots, N$. The 4D-Var problem provides the estimate $\mathbf{x}_0^a$ of the true initial conditions as the solution of the following optimization problem

$$\mathbf{x}_0^a = \arg \min_{\mathbf{x}_0} \mathcal{J}(\mathbf{x}_0) \quad \text{subject to (1),} \quad (3)$$

with the following cost function:

$$\mathcal{J}(\mathbf{x}_0) = \frac{1}{2} (\mathbf{x}_0 - \mathbf{x}_0^b)^T \mathbf{B}_0^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b) + \frac{1}{2} \sum_{k=1}^N (\mathcal{H}_k(\mathbf{x}_k) - \mathbf{y}_k)^T \mathbf{R}_k^{-1} (\mathcal{H}_k(\mathbf{x}_k) - \mathbf{y}_k). \quad (4)$$

The first term of the sum (4) quantifies the departure of the solution $\mathbf{x}_0$ from the background state $\mathbf{x}_0^b$ at the initial time t_0 . The second term measures the mismatch between the forecast trajectory (model solutions $\mathbf{x}_k$) and observations $\mathbf{y}_k$ at all times t_k in the assimilation window. The weighting matrices $\mathbf{B}_0$ and $\mathbf{R}_k$ need to be predefined, and their quality influences the accuracy of the resulting analysis.

Weak constraint 4D-Var [34] removes the perfect model assumption by allowing a model error $\eta_{k+1} = \mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)$. Under the assumption that the model errors are normally distributed, $\eta_k \in \mathcal{N}(0, \mathbf{Q}_k)$, the weak constraint 4D-Var solution is the unconstrained minimizer of the cost function

$$\mathcal{J}^{\text{weak}}(\mathbf{x}_0, \dots, \mathbf{x}_N) = \mathcal{J}(\mathbf{x}_0) + \frac{1}{2} \sum_{k=0}^{N-1} (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k))^T \mathbf{Q}_{k+1}^{-1} (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)). \quad (5)$$

The control variables are the states of the system at all times in the assimilation window.

In this paper we focus on the strong constraint formulation (3). The minimizer of (3) is computed iteratively using gradient-based numerical optimization methods. First-order adjoint models provide the gradient of the cost function [11], while second-order adjoint models provide the Hessian-vector product (e.g., for Newton-type methods). The methodology for building and using various adjoint models for optimization, sensitivity analysis, and uncertainty quantification is discussed in [12, 33]. Various strategies to improve the 4D-Var data assimilation system are described in [13]. The procedure to estimate the impact of observation and model errors is developed in [28, 30]. A framework to perform derivative free variational data assimilation using the trust-region framework is given in [32].

The iterative solution of (3) is highly sequential: first, one iteration follows the other; next, the forward and adjoint models are run sequentially forward and backward in time, respectively. In order to reveal additional parallelism the solution to 4D-Var problem is approached using the augmented Lagrangian framework. In this framework, the assimilation window is divided into multiple sub-intervals, and the model constraints are explicitly imposed at the boundaries. This approach bears similarities with the Parareal approach that exploits time parallelism in the solution of ordinary differential equations [29].

3. 4D-Var solution by the augmented Lagrangian approach

The 4D-Var cost function (4) is minimized subject to the generic model constraints (1). The model equations can also be written as

$$\mathbf{x}_{k+1} = \mathcal{M}_{k,k+1}(\mathbf{x}_k), \quad k = 0, 1, \dots, N-1, \quad (6)$$

where $\mathcal{M}_{k,k+1}$ represents the model solution operator that propagates the state $\mathbf{x}_k$ at t_k to the state $\mathbf{x}_{k+1}$ at t_{k+1} . The minimizer of (4) under the constraints (6) is the unconstrained minimizer of the Lagrangian

$$\mathbf{L}(\mathbf{x}_0, \dots, \mathbf{x}_N; \lambda_0, \dots, \lambda_N) = \mathcal{J}(\mathbf{x}_0) - \sum_{k=0}^{N-1} \lambda_{k+1}^T \cdot (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)). \quad (7)$$

To expose time parallelism the assimilation window is divided into N sub-intervals, namely,

$$[t_0, t_N] = [t_0, t_1] \cup \dots \cup [t_{N-1}, t_N]. \quad (8)$$

The forward model and adjoint model states at the interval boundaries are denoted by

$$\mathbf{x} = [\mathbf{x}_0, \dots, \mathbf{x}_N], \quad \boldsymbol{\lambda} = [\lambda_0, \dots, \lambda_N], \quad (9)$$

respectively. We denote by $\mathbf{x}^a$ the optimal solution and by $\boldsymbol{\lambda}^a$ the optimal value of the Lagrange multiplier in (7).

The augmented Lagrangian [39, Section 17.3] associated with (4) and (6) reads:

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \mu) = & \frac{1}{2} (\mathbf{x}_0 - \mathbf{x}_0^b)^T \mathbf{B}_0^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b) + \frac{1}{2} \sum_{k=1}^N (\mathcal{H}(\mathbf{x}_k) - \mathbf{y}_k)^T \mathbf{R}_k^{-1} (\mathcal{H}(\mathbf{x}_k) - \mathbf{y}_k) \\ & - \sum_{k=0}^{N-1} \lambda_{k+1}^T (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)) \\ & + \frac{\mu}{2} \sum_{k=0}^{N-1} (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k))^T \mathbf{P}_{k+1}^{-1} (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)), \end{aligned} \quad (10)$$

where $\mathbf{P}_k$'s are error scaling matrices. This is the (regular) Lagrangian for the problem that minimizes the cost function

$$\mathcal{J}^{\text{augmented}}(\mathbf{x}_0, \dots, \mathbf{x}_N) = \mathcal{J}(\mathbf{x}_0) + \frac{\mu}{2} \sum_{k=0}^{N-1} (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k))^T \mathbf{P}_{k+1}^{-1} (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)) \quad (11)$$

subject to the model constraint (6). The constrained minimization of (11) is equivalent to (3) since the additional term is zero along the constraints. Note that (11) is a constrained minimization problem, unlike (5), which is unconstrained.

The original 4D-Var problem in (3) is solved in the augmented Lagrangian framework by performing a sequence of unconstrained minimizations

$$\mathbf{x}^{\{\ell\}} = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \tilde{\boldsymbol{\lambda}}^{\{\ell\}}, \mu^{\{\ell\}}), \quad \ell = 0, 1, \dots \quad (12)$$

If $\tilde{\boldsymbol{\lambda}}^{\{\ell\}} \approx \boldsymbol{\lambda}^a$ then $\tilde{\mathbf{x}}^{\{\ell\}} \approx \mathbf{x}^a$, and the solution error decreases with increasing μ [39, Section 17.3].

The optimization proceeds in cycles of inner and outer iterations. Inner iterations solve the optimization problem (12) for particular values of $\mu^{\{\ell\}}$ and $\tilde{\boldsymbol{\lambda}}^{\{\ell\}}$. After each solution of (12) the outer iteration ℓ is completed by updating the Lagrange multiplier approximation and the penalty parameter, as follows:

$$\mu^{\{\ell+1\}} = \rho \mu^{\{\ell\}}, \quad (13a)$$

$$\lambda_k^{\{\ell+1\}} = \lambda_k^{\{\ell\}} - \mu^{\{\ell\}} (\mathbf{x}_k^{\{\ell\}} - \mathcal{M}_{k-1,k}(\mathbf{x}_{k-1}^{\{\ell\}})), \quad k = 0, \dots, N. \quad (13b)$$

The penalty parameter is progressively increased by a constant $\rho > 1$ in order to impose the model constraints (6). Other strategies to update μ and $\boldsymbol{\lambda}$ can be used [6].

3.1. Augmented Lagrangian optimization

We illustrate the convergence process of the augmented Lagrangian 4D-Var in the context of the Lorenz 96 problem in Figure 1 below. There are a few key general observations we make here. The initial value of the control vector contains the background states $\mathbf{x}_k^b$ at the beginning of each of the sub-interval, and is therefore a continuous curve. The outer iterations start with a small value of μ that only imposes the constraints (6) loosely. Consequently, the solutions during the first outer iterations show large discontinuities at the interval boundaries. Subsequently, μ is increased and as a result the constraints are satisfied accurately. This results in a smooth solution curve resembling closely the serial 4D-Var solution.

4. The parallel algorithm

Most of the computational time required by the 4D-Var solution is claimed by the multiple cost function and gradient evaluations. In the augmented Lagrangian 4D-Var formulation (12)–(13) the forward and adjoint models can be run in parallel over sub-intervals, as explained next.

4.1. Parallel-in-time runs of the forward model

The value of the augmented Lagrangian cost function (10) can be computed in parallel. Specifically, on each sub-interval $[t_k, t_{k+1}]$ in (8) a forward solution $\mathcal{M}_{k,k+1}(\mathbf{x}_k)$ is computed starting from the initial value $\mathbf{x}_k$ (9). The forward model runs on each sub-interval can be carried out concurrently. The computational steps are detailed in Algorithm 1.

Algorithm 1 Parallel_CostFunction

```

1: procedure PARALLEL_COSTFUNCTION
2:   Input:  $[\mathbf{x}_0, \dots, \mathbf{x}_N; \lambda_1, \dots, \lambda_N]$ 
3:   Output:  $\mathcal{L}$ 
4:    $\mathcal{L} \leftarrow 0$ 
5:   For all  $0 \leq k \leq N - 1$  do in parallel
6:      $\Delta \mathbf{x}_{k+1} \leftarrow \mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)$  ▷ Solution mismatch at sub-interval boundaries
7:      $\Delta \mathbf{y}_{k+1} \leftarrow \mathcal{H}(\mathbf{x}_{k+1}) - \mathbf{y}_{k+1}$  ▷ Solution mismatch with observations
8:   end For all
9:   for  $0 \leq k \leq N - 1$  do ▷ Evaluate the cost function using (10)
10:     $\mathcal{L} \leftarrow \mathcal{L} + \frac{\mu}{2} \Delta \mathbf{x}_{k+1}^T \mathbf{P}_{k+1}^{-1} \Delta \mathbf{x}_{k+1} - \lambda_{k+1}^T \Delta \mathbf{x}_{k+1} + \frac{1}{2} \Delta \mathbf{y}_{k+1}^T \mathbf{R}_{k+1}^{-1} \Delta \mathbf{y}_{k+1}$ 
11:   end for
12:    $\mathcal{L} \leftarrow \mathcal{L} + \frac{1}{2} (\mathbf{x}_0 - \mathbf{x}_0^b)^T \mathbf{B}_0^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b)$ 
13: end procedure

```

4.2. Parallel-in-time runs of the adjoint model

The gradient of the augmented Lagrangian (10) with respect to the forward model state is given by:

$$\nabla_{\mathbf{x}_0} \mathcal{L} = \mathbf{B}_0^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b) + \mathbf{M}_{0,1}^T \lambda_1 - \mu \mathbf{M}_{0,1}^T \mathbf{P}_0^{-1} (\mathbf{x}_1 - \mathcal{M}_{0,1}(\mathbf{x}_0)), \quad (14a)$$

$$\begin{aligned} \nabla_{\mathbf{x}_k} \mathcal{L} = & \mathbf{H}_k^T \mathbf{R}_k^{-1} (\mathcal{H}(\mathbf{x}_k) - \mathbf{y}_k) - \lambda_k^T + \mathbf{M}_{k,k+1}^T \lambda_{k+1} \\ & + \mu \mathbf{P}_{k-1}^{-1} (\mathbf{x}_k - \mathcal{M}_{k-1,k}(\mathbf{x}_{k-1})) \\ & - \mu \mathbf{M}_{k,k+1}^T \mathbf{P}_k^{-1} (\mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)), \quad k = 1, \dots, N-1, \end{aligned} \quad (14b)$$

$$\nabla_{\mathbf{x}_N} \mathcal{L} = \mathbf{H}_N^T \mathbf{R}_N^{-1} (\mathcal{H}(\mathbf{x}_N) - \mathbf{y}_N) - \lambda_N^T + \mu \mathbf{P}_{N-1}^{-1} (\mathbf{x}_N - \mathcal{M}_{N-1,N}(\mathbf{x}_{N-1})), \quad (14c)$$

where the tangent linear operators of the observation and model solution operators are

$$\mathbf{H}_k = \left. \frac{\partial \mathcal{H}_k}{\partial \mathbf{x}_k} \right|_{\mathbf{x}_k}, \quad \mathbf{M}_{k,k+1} = \left. \frac{\partial \mathcal{M}_{k,k+1}}{\partial \mathbf{x}_k} \right|_{\mathbf{x}_k},$$

respectively, and their transposes are the corresponding adjoint operators. Using the notation

$$\begin{aligned} \Delta \mathbf{x}_k &:= \mathbf{x}_k - \mathcal{M}_{k-1,k}(\mathbf{x}_{k-1}) \\ \Delta \mathbf{y}_k &:= \mathcal{H}(\mathbf{x}_k) - \mathbf{y}_k \end{aligned}$$

the gradient (14) can be written as

$$\begin{aligned} \nabla_{\mathbf{x}_0} \mathcal{L} &= \mathbf{B}_0^{-1}(\mathbf{x}_0 - \mathbf{x}_0^b) + \mathbf{M}_{0,1}^T \lambda_1 - \mu \mathbf{M}_{0,1}^T \mathbf{P}_1^{-1} \Delta \mathbf{x}_1, \\ \nabla_{\mathbf{x}_k} \mathcal{L} &= \mathbf{H}_k^T \mathbf{R}_k^{-1} \Delta \mathbf{y}_k - \lambda_k^T + \mathbf{M}_{k,k+1}^T (\lambda_{k+1} - \mu \mathbf{P}_{k+1}^{-1} \Delta \mathbf{x}_{k+1}) \\ &\quad + \mu \mathbf{P}_k^{-1} \Delta \mathbf{x}_k, \quad k = 1, \dots, N-1, \\ \nabla_{\mathbf{x}_N} \mathcal{L} &= \mathbf{H}_N^T \mathbf{R}_N^{-1} \Delta \mathbf{y}_N - \lambda_N^T + \mu \mathbf{P}_{N-1}^{-1} \Delta \mathbf{x}_N. \end{aligned}$$

The augmented Lagrangian gradient (14) can be evaluated in parallel. On each sub-interval $[t_k, t_{k+1}]$ in (8) an adjoint solution $\mathbf{M}_{k,k+1}^T \lambda_{k+1}$ is computed starting from the terminal value λ_{k+1} (9). The adjoint model runs on each sub-interval can be carried out concurrently. The computational steps are detailed in Algorithm 2.

Algorithm 2 Parallel_Gradient

```

1: procedure PARALLEL_GRADIENT
2:   Input:  $[\mathbf{x}_0, \dots, \mathbf{x}_N; \lambda_1, \dots, \lambda_N]$ 
3:   Output:  $[\nabla_{\mathbf{x}_0} \mathcal{L}, \dots, \nabla_{\mathbf{x}_N} \mathcal{L}]$ 
4:   For all  $0 \leq k \leq N-1$  do in parallel
5:      $\Delta \mathbf{x}_{k+1} \leftarrow \mathbf{x}_{k+1} - \mathcal{M}_{k,k+1}(\mathbf{x}_k)$  ▷ Solution mismatch at sub-interval boundaries
6:      $\Delta \mathbf{y}_{k+1} \leftarrow \mathcal{H}(\mathbf{x}_{k+1}) - \mathbf{y}_{k+1}$  ▷ Innovation vector
7:   end For all
8:   For all  $0 \leq k \leq N-1$  do in parallel
9:      $\mathbf{b}_{k+1} \leftarrow \mu \mathbf{P}_{k+1}^{-1} \Delta \mathbf{x}_{k+1} - \lambda_{k+1}$  ▷ Translate and scale solution mismatch
10:     $\mathbf{d}_{k+1} \leftarrow \mathbf{H}_N^T \mathbf{R}_N^{-1} \Delta \mathbf{y}_{k+1}$  ▷ Scale innovation vector
11:  end For all
12:  For all  $0 \leq k \leq N-1$  do in parallel ▷ Perform the adjoint integrations in parallel
13:     $\mathbf{a}_k \leftarrow \mathbf{M}_{k,k+1}^T \mathbf{b}_{k+1}$ 
14:  end For all
15:  For all  $1 \leq k \leq N-1$  do in parallel ▷ Compute gradient using (14)
16:     $\nabla_{\mathbf{x}_k} \mathcal{L} \leftarrow \mathbf{b}_k + \mathbf{d}_k - \mathbf{a}_k$ 
17:  end For all
18:   $\nabla_{\mathbf{x}_0} \mathcal{L} \leftarrow \mathbf{B}_0^{-1}(\mathbf{x}_0 - \mathbf{x}_0^b) - \mathbf{a}_0$ 
19:   $\nabla_{\mathbf{x}_N} \mathcal{L} \leftarrow \mathbf{d}_N + \mathbf{b}_N$ 
20: end procedure

```

4.3. Initial solution guess

The optimization needs to start with some initial guess for $\mathbf{x}^{\{0\}}$ and $\lambda^{\{0\}}$. The initial guess for the state is obtained by performing a serial forward integration using the background value for initial conditions, $\mathbf{x}_k^{\{0\}} = \mathbf{x}_k^b$. The initial value for the adjoint variable could be obtained by solving the adjoint model once, serially, along the background trajectory. In our experiments we choose the simpler, and less expensive, initialization $\lambda_k^{\{0\}} = 0$.

4.4. Updating Lagrange multipliers

In order to accelerate the convergence of the optimization process we replace the standard updates (13) with the strategy proposed in [19]. The new update process uses information from the previous two iterations instead of just one iteration. The process takes the following steps:

1. Choose $\boldsymbol{\lambda}^0$ and set $\alpha^{\{1\}} = 1$.
2. Let $\mathbf{x}^{\{\ell\}}$ be the solution obtained by solving the optimization problem (12) for particular values of $\mu^{\{\ell\}}$ and $\boldsymbol{\lambda}^{\{\ell\}}$. Apply the classical update (13) to obtain $\mu^{\{\ell+1\}}$ and $\tilde{\boldsymbol{\lambda}}^{\{\ell+1\}}$.
3. The updated Lagrange multiplier is obtained as follows:

$$\begin{aligned}\alpha^{\{\ell+1\}} &= \frac{1}{2} \left(1 + \sqrt{1 + 4(\alpha^{\{\ell\}})^2} \right), \\ \boldsymbol{\lambda}^{\{\ell+1\}} &= \tilde{\boldsymbol{\lambda}}^{\{\ell+1\}} + \left(\frac{\alpha^{\{\ell\}} - 1}{\alpha^{\{\ell+1\}}} \right) (\tilde{\boldsymbol{\lambda}}^{\{\ell+1\}} - \tilde{\boldsymbol{\lambda}}^{\{\ell\}}) + \left(\frac{\alpha^{\{\ell\}}}{\alpha^{\{\ell+1\}}} \right) (\tilde{\boldsymbol{\lambda}}^{\{\ell+1\}} - \boldsymbol{\lambda}^{\{\ell\}}).\end{aligned}$$

It is important to note that above procedure requires the values of $\tilde{\boldsymbol{\lambda}}$ from two successive outer iterations, namely ℓ and $\ell + 1$.

5. Numerical experiments

We study the performance of the parallel implementation of augmented Lagrangian 4D-Var algorithm using the Lorenz-96 model with 40 variables [23] and a shallow water on the sphere model with $\sim 8,000$ variables [36].

5.1. Lorenz-96 model

The Lorenz-96 model [23] is given by

$$\frac{d\mathbf{x}_k}{dt} = \mathbf{x}_{k-1} (\mathbf{x}_{k+1} - \mathbf{x}_{k-2}) - \mathbf{x}_k + F, \quad k = 1, \dots, 40, \quad (15)$$

with periodic boundary conditions and the forcing term $F = 8$ [23]. We use synthetic observations generated by perturbing the reference trajectory with normal noise with mean zero and standard deviation of 5% of average magnitude of the reference solution. The background uncertainty is set to 8% of average magnitude of the reference solution. The background and observation error covariance matrices are assumed to be diagonal. A vector of equidistant components ranging from -2 to 2 was integrated forward in time for 200 time steps and the final state is taken as a reference initial condition for the experiments.

5.2. Shallow water model on the sphere

The shallow water equations have been used extensively as a simple model of the atmosphere since they contain the essential wave propagation mechanisms found in general circulation models [36]. The shallow water equations in spherical coordinates are:

$$\frac{\partial u}{\partial t} + \frac{1}{a \cos \theta} \left(u \frac{\partial u}{\partial \lambda} + v \cos \theta \frac{\partial u}{\partial \theta} \right) - \left(f + \frac{u \tan \theta}{a} \right) v + \frac{g}{a \cos \theta} \frac{\partial h}{\partial \lambda} = 0, \quad (16a)$$

$$\frac{\partial v}{\partial t} + \frac{1}{a \cos \theta} \left(u \frac{\partial v}{\partial \lambda} + v \cos \theta \frac{\partial v}{\partial \theta} \right) + \left(f + \frac{u \tan \theta}{a} \right) u + \frac{g}{a} \frac{\partial h}{\partial \theta} = 0, \quad (16b)$$

$$\frac{\partial h}{\partial t} + \frac{1}{a \cos \theta} \left(\frac{\partial (hu)}{\partial \lambda} + \frac{\partial (hv \cos \theta)}{\partial \theta} \right) = 0. \quad (16c)$$

Here f is the Coriolis parameter given by $f = 2\Omega \sin \theta$, where Ω is the angular speed of the rotation of the Earth, h is the height of the homogeneous atmosphere, u and v are the zonal and meridional wind components, respectively, θ and λ are the latitudinal and longitudinal directions, respectively, a is the

radius of the earth and g is the gravitational constant. The space discretization is performed using the un-staggered Turkel-Zwas scheme [24, 25]. The discretized spherical grid has $n_{lon}=36$ nodes in longitudinal direction and $n_{lat}=72$ nodes in the latitudinal direction. The semi-discretization in space leads to a discrete model of the form (6). In (6) the zonal wind, meridional wind and the height variables are combined into the vector $\mathbf{x} \in \mathbb{R}^n$ with $n = 3 \times n_{lat} \times n_{lon}$. We perform the time integration using an adaptive time-stepping algorithm. For a tolerance of 10^{-8} the average time step size of the time-integrator is 180 seconds. A reference initial condition is used to generate a reference trajectory.

Synthetic observation errors at various times t_k are normally distributed with mean zero and a diagonal observation error covariance matrix with entries equal to $(\mathbf{R}_k)_{i,i} = 1$ for u and v components and $(\mathbf{R}_k)_{i,i} = 10^6$ for h components. The $\mathbf{R}_k$ values correspond to a standard deviation of 5% for u and v components, and 2% for h component. We construct a flow dependent background error covariance matrix as described in [1, 2]. The standard deviation of the background errors for the height component is 2% of the average magnitude of the reference height component in the reference initial condition. The standard deviation of the background errors for the wind components is 15% of the average magnitude of the reference wind component in the reference initial condition.

5.3. Experimental setup

All numerical experiments are carried out in MATLAB. The parallel implementations are developed using MATLAB's parallel toolbox. We compare the performance of the proposed parallel implementation with that of the standard 4D-Var. The accuracy of numerical solutions is measured by the root mean square error (RMSE) with respect to a reference solution. The RMSE is given by

$$\mathbf{RMSE} = \sqrt{\frac{1}{N} \sum_{k=1}^N N_{\text{var}}^{-1} \|\mathbf{x}_k^a - \mathbf{x}_k^{\text{reference}}\|^2}, \quad (17)$$

where the reference solution $\mathbf{x}^{\text{reference}}$ and the analysis $\mathbf{x}^a$ are propagated forward in time using the full model, and the difference is measured at all times throughout the assimilation window.

Computing the cost functions and gradients is the most important aspect of the 4D-Var algorithm and hence it is necessary that their computations are scalable. To evaluate the gradient and cost function we carry out numerical integrations of the forward and adjoint models using MATLODE [3]. MATLODE is a MATLAB version of FATODE, which was developed in FORTRAN [40]. The optimization is carried out using the LBFGS solver [22] implemented in the Poblano optimization toolbox developed at Sandia National Laboratory [16].

5.4. Results with the Lorenz-96 model

Figure 1 illustrates the convergence of augmented Lagrangian 4D-Var iterations. The intermediate solutions are discontinuous at sub-interval boundaries. The corresponding errors, with respect to the traditional 4D-Var solution, are shown in Figure 2. They decrease quickly to zero, showing that the augmented Lagrangian 4D-Var solution converges to the solution of traditional 4D-Var.

The errors of the augmented Lagrangian and traditional 4D-Var solutions with respect to the reference solution are shown in Figure 3. The reference solution is obtained by propagating the reference initial condition using the forward model in (15). The sequential 4D-Var requires 230 gradient and 574 cost function evaluations, where as the parallel 4D-Var requires a total of 100 gradient and 650 cost function evaluations for 6 observations. The Weak scalability results are presented in Figure 4. As the length of the assimilation window increases, the number of sub-intervals increases, and so does the number of observations (one per sub-interval). The number of cores on which the parallel algorithm is run increases such as to remain equal to the number of sub-intervals. The parallel algorithm is scalable in a weak sense: the total computational time increases very slowly with an increased problem size. The most time consuming calculations are those of the cost function and gradient evaluations, which require running the forward and the adjoint models, respectively. The results shown in Figure 5 confirm the good weak scalability of the cost function and gradient computations.

Figure 1: Lorenz-96 model (15) results. Convergence of the parallel augmented Lagrangian 4D-Var solution over the first several iterations. The final solution is computed with the traditional 4D-Var. Three different variables are shown as an illustration.

Figure 2: Lorenz-96 model (15) results. Difference between the augmented Lagrangian solutions at different iterations and the solution of traditional 4D-Var. Three different variables are shown as an illustration.

Figure 3: Lorenz-96 model (15) results. RMSE errors of traditional and augmented Lagrangian solutions with respect to the reference analysis. The two implementations give nearly identical results. Traditional 4D-Var approach requires 230 iterations, whereas the augmented Lagrangian requires 100 iterations.

Figure 4: Lorenz-96 model (15) results. The scaling of overall computational times with increasing number of sub-intervals (which leads to an increase in the length of the assimilation window).

Figure 5: Lorenz-96 model (15) results. Weak scalability results for the cost function and gradient evaluations.

5.5. Results with the shallow water model

Figure 6 shows the weak scalability of cost function and gradient evaluations with the shallow water model. It can be seen that in both cases the parallel computational time is nearly constant with increasing problem size (number of sub-intervals) and a proportional increase in the number of cores (the number of cores is equal to the number of sub-intervals). Figure 7 presents the work-precision diagrams, i.e., the evolution of solution accuracy (RMSE) with increasing number of iterations (increasing CPU time). The initial iterates of the parallel 4D-Var solution proceed rapidly, but afterwards the convergence slows down. At this stage the penalty parameter μ is fairly large and the optimization problem becomes more difficult for the LBFGS algorithm. The performance can be improved by replacing LBFGS with algorithms specially tailored to solve optimization problems in the augmented Lagrangian framework [7]. An alternative strategy is to use a hybrid method: employ parallel 4D-Var for several iterations in the beginning, then continue with traditional serial 4D-Var. Here we perform two outer iterations with small values of μ and then use this solution as an initial guess for the serial 4D-Var. This strategy improves the performance considerably as seen in Figure 7(a). Table 1 provides the computational times for parallel and serial 4D-Var algorithms. The serial 4D-Var for 9 observations requires 95 gradient and 200 cost function evaluations, whereas the parallel 4D-Var algorithm requires 350 gradient and 720 cost function evaluations. In the hybrid methodology, we perform 200 iterations of parallel 4D-Var and 30 iterations of serial 4D-Var for 9 observations. The final RMSE over the assimilation window for both the serial 4D-Var and hybrid methods is ~ 125 . We notice a steady increase in speedup as the problem size (number of sub-intervals) is increased. It is possible to further improve the performance of the parallel algorithm by using second derivative information in the form of Hessian-vector products and employ Newton-type methods [12].

No. of sub-intervals	Serial time [sec.]	Parallel/Hybrid time [sec.]	Speedup
5	8,515	10,834	0.786
7	15,745	12,782	1.232
9	22,277	12,971	1.756

Table 1: Shallow water equations (16) results. Performance comparison of augmented Lagrangian/parallel and traditional/serial 4D-Var and the corresponding speedups.

Figure 6: Shallow water equations (16) results. Weak scalability results for the cost function and gradient evaluations.

Figure 7: Shallow water equations (16) results. Work-performance diagrams comparing parallel and hybrid methods with serial 4D-Var for 9 sub-intervals.

6. Conclusions and future work

This work presents an augmented Lagrangian framework to perform strong-constraint 4D-Var data assimilation in parallel. The assimilation window is split in sub-intervals; cost function and gradient evaluations, which are the main components of the algorithm, are performed by running the forward and the adjoint model in parallel across different sub-intervals.

Future work will focus on tuning the optimization procedure to improve performance on large scale problems, e.g., data assimilation with the weather research and forecast model [35]. The size of the control variable in the augmented Lagrangian framework increases with the number of sub-intervals and can become a bottleneck for the optimization. One possible strategy to overcome this is to perform optimization on a coarse grid and use the projected solution as an initial guess for the fine grid. A natural extension to our methodology is to use the augmented Lagrangian framework to expose space-time parallelism in the data assimilation problem such as to create more parallel tasks and improve the overall scalability. Space parallelism in a penalty formulation has been recently discussed in [15]. We will consider the use of optimization algorithms that are specifically tuned to work well in the augmented Lagrangian framework [7]. Hessian information can be used to accelerate the convergence significantly when the iterates are close to minima. In order to implement this it is useful to explore the construction of second order adjoint models that compute the Hessian-vector products in parallel. As the length of the assimilation window increases, the optimization procedure slows down due to the increase in condition number of the associated Hessian. Hence there is a need to build efficient preconditioners to overcome this bottleneck. Finally, we plan to implement the parallel algorithm in a lower level language (C or FORTRAN using OpenMP or MPI) to have better control over task parallelism.

Acknowledgements

This work was supported in part by the awards AFOSR FA9550-12-1-0293-DEF, AFOSR 12-2640-06, NSF DMS-1419003, NSF CCF-1218454, and by the Computational Science Laboratory at Virginia Tech.

References

- [1] A. Attia, V. Rao, and A. Sandu. A hybrid Monte-Carlo sampling smoother for four dimensional data assimilation. *Ocean Dynamics*, Submitted, 2014.
- [2] A. Attia, V. Rao, and A. Sandu. A sampling approach for four dimensional data assimilation. In *Proceedings of the Dynamic Data Driven environmental System Science Conference*, 2014.
- [3] A. Augustine and A. Sandu. MATLODE: A MATLAB ODE solver and sensitivity analysis toolbox. *ACM TOMS*, Submitted, 2015.
- [4] A. Aved, F. Darema, and E. Blasch. Dynamic data driven application systems. www.1dddas.org, 2014.
- [5] P. Bauer, T. Aulign, W. Bell, A. Geer, V. Guidard, S. Heilliette, M. Kazumori, M.J. Kim, E.H.C Liu, A.P. McNally, B. Macpherson, K. Okamoto, R. Renshaw, and L.P. Riishjgaard. Satellite cloud and precipitation assimilation at operational nwp centres. *Quarterly Journal of the Royal Meteorological Society*, 137(661):1934–1951, 2011.
- [6] E. G. Birgin and J. M. Martínez. Improving ultimate convergence of an augmented Lagrangian method. *Optimization Methods and Software*, 23(2):177–195, 2008.
- [7] E. G. Birgin and J. M. Martínez. *Practical augmented Lagrangian methods for constrained optimization*, volume 10. SIAM, 2014.
- [8] G. Biros and O. Ghattas. Parallel Lagrange-Newton-Krylov-Schur methods for PDE-constrained optimization. Part I: the Krylov-Schur solver. *SIAM Journal on Scientific Computing*, 27(2):687–713, 2005.
- [9] E. Blayo, M. Bocquet, E. Cosme, and L.F. Cugliandolo. Advanced data assimilation for geosciences. In *International Summer School-Advanced Data Assimilation for Geosciences*, page 608, 2014.
- [10] A. Borzi. *Space-Time Multigrid Methods for Solving Unsteady Optimal Control Problems*, chapter 5, pages 97–113. SIAM, 2007.
- [11] D. G. Cacuci, M. Ionescu-Bujor, and I. M. Navon. *Sensitivity and uncertainty analysis, volume II: applications to large-scale systems*, volume 2. CRC Press, 2005.
- [12] A. Cioaca, M. Alexe, and A. Sandu. Second-order adjoints for solving PDE-constrained optimization problems. *Optimization methods and software*, 27(4-5):625–653, 2012.
- [13] A. Cioaca and A. Sandu. An optimization framework to improve 4D-Var data assimilation system performance. *Journal of Computational Physics*, 275:377–389, 2014.
- [14] R. Daley. *Atmospheric data analysis*, volume 2. Cambridge University Press, 1993.

- [15] L. D’Amore, R. Arcucci, L. Carracciolo, and A. Murli. A scalable approach for variational data assimilation. *Journal of Scientific Computing*, 61(2):239–257, 2014.
- [16] D. M. Dunlavy, T. G. Kolda, and E. Acar. Poblano v1.0: A MATLAB toolbox for gradient-based optimization. Technical report, Sandia National Laboratories, 2010.
- [17] H. Elbern, H. Schmidt, and A. Ebel. Implementation of a parallel 4d-variational chemistry data-assimilation scheme. *Environmental management and health*, 10(4):236–245, 1999.
- [18] M. Fisher. Parallelization of 4D-Var in the time dimension using a saddlepoint algorithm. Slides., 2013.
- [19] B. He and X. Yuan. On the acceleration of augmented Lagrangian method for linearly constrained optimization. *optimization online*, 2010.
- [20] MJ Hossen, IM Navon, and F Fang. A penalized four-dimensional variational data assimilation method for reducing forecast error related to adaptive observations. *International Journal for Numerical Methods in Fluids*, 70(10):1207–1220, 2012.
- [21] E. Kalnay. *Atmospheric modeling, data assimilation, and predictability*. Cambridge University Press, 2003.
- [22] D. C. Liu and J. Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- [23] E. N. Lorenz. Predictability: A problem partly solved. In *Proceedings of Seminar on Predictability*, pages 40–58, 1996.
- [24] I. M. Navon and R. De Villiers. The application of the Turkel-Zwas explicit large time-step scheme to a hemispheric barotropic model with constraint restoration. *Monthly Weather Review*, 115(5):1036–1052, 1987.
- [25] I. M. Navon and J. Yu. Exshall: A Turkel-Zwas explicit large time-step FORTRAN program for solving the shallow-water equations in spherical coordinates. *Computers and Geosciences*, 17(9):1311 – 1343, 1991.
- [26] IM Navon and R De Villiers. Combined penalty multiplier optimization methods to enforce integral invariants conservation. *Monthly Weather Review*, 111(6):1228–1243, 1983.
- [27] J. Rantakokko. Strategies for parallel variational data assimilation. *Parallel Computing*, 23(13):2017 – 2039, 1997.
- [28] V. Rao and A. Sandu. A-posteriori error estimates for inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*, Submitted, 2014.
- [29] V. Rao and A. Sandu. An adjoint-based scalable algorithm for time-parallel integration. *Journal of Computational Science*, 5(2):76 – 84, 2014. Empowering Science through Computing + BioInspired Computing.
- [30] V. Rao and A. Sandu. A posteriori error estimates for DDDAS inference problems. *Procedia Computer Science*, 29:1256–1265, 2014.
- [31] V. Rao and A. Sandu. Parallel solution of dddas variational inference problems. *Procedia Computer Science*, 51:2474–2482, 2015.
- [32] E. D. N. Ruiz and A. Sandu. A derivative-free trust region framework for variational data assimilation. *Journal of Computational and Applied Mathematics*, 2015.
- [33] A. Sandu, D. N. Daescu, G. R. Carmichael, and T. Chai. Adjoint sensitivity analysis of regional air quality models. *Journal of Computational Physics*, 204(1):222–252, 2005.
- [34] Adrian Sandu and Tianfeng Chai. Chemical data assimilation—an overview. *Atmosphere*, 2(3):426–463, 2011.
- [35] William C Skamarock, Joseph B Klemp, Jimmy Dudhia, David O Gill, Dale M Barker, Wei Wang, and Jordan G Powers. A description of the advanced research wrf version 2. Technical report, DTIC Document, 2005.
- [36] A. St-Cyr, C. Jablonowski, J. M. Dennis, H. M. Tufo, and S. J. Thomas. A comparison of two shallow water models with nonconforming adaptive grids. *Monthly Weather Review*, 136:1898–1922, 2008.
- [37] Y. Trémolet and F.-X. Le Dimet. Parallel algorithms for variational data assimilation and coupling models. *Parallel Computing*, 22(5):657 – 674, 1996.
- [38] S. Ulbrich. Generalized SQP methods with ‘parareal’ time-domain decomposition for time-dependent PDE-constrained optimization. In *Real-time PDE-constrained optimization*, volume 3 of *Comput. Sci. Eng.*, pages 145–168. SIAM, Philadelphia, PA, 2007.
- [39] S. J. Wright and J. Nocedal. *Numerical optimization*, volume 2. Springer New York, 2006.
- [40] H. Zhang and A. Sandu. FATODE: A library for forward, adjoint, and tangent linear integration of ODEs. *SIAM Journal on Scientific Computing*, 36(5):C504–C523, 2014.