

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

On the use of the saddle formulation in weakly-constrained 4D-VAR data assimilation

Gratton, Serge; Gürol, Selime; Simon, Ehouarn; Toint, Philippe

Publication date:
2017

Document Version
Early version, also known as pre-print

[Link to publication](#)

Citation for published version (HARVARD):

Gratton, S, Gürol, S, Simon, E & Toint, P 2017 'On the use of the saddle formulation in weakly-constrained 4D-VAR data assimilation' Arxiv.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

On the use of the saddle formulation in weakly-constrained 4D-VAR data assimilation

S. Gratton ^{*}, S. Gürol [†], E. Simon [‡] and Ph. L. Toint [§]

19 September 2017

Abstract

This paper discusses the practical use of the saddle variational formulation for the weakly-constrained 4D-VAR method in data assimilation. It is shown that the method, in its original form, may produce erratic results or diverge because of the inherent lack of monotonicity of the produced objective function values. Convergent, variationally coherent variants of the algorithm are then proposed whose practical performance is compared to that of other formulations. This comparison is conducted on two data assimilation instances (Burgers equation and the Quasi-Geostrophic model), using two different assumptions on parallel computing environment. Because these variants essentially retain the parallelization advantages of the original proposal, they often — but not always — perform best, even for moderate numbers of computing processes.

Keywords: data assimilation, variational methods, weakly-constrained 4D-VAR, saddle formulation, parallel computing.

1 Introduction

Data assimilation has long been an integral and important part of weather forecasting, as new (and often incomplete) meteorological observations are integrated in the ongoing process of predicting the weather for the next few days [4]. The question here is that of using the data to determine a “best” current state of the weather system from which elaborate models may then be evolved in time, providing the desired predictions. Among the possible techniques for this task, variational methods have been applied extensively, typically weighting the use of a priori knowledge (often materialized by the specification of a background state x_b) with the quality of the fit to the observations. This is the case, in particular, for the well-known 4D-Var formulation [23, 7]. In recent years, it has also become necessary to take possible model errors into account, thus weighting a priori knowledge, data fitting and model error reduction, an approach which leads to the “weakly-constrained 4D-Var” formulation of the relevant data assimilation problem [36, 31, 29, 30]. In one of the formulations, the total time

^{*}Université de Toulouse, INP, IRIT, Toulouse, France. Email: serge.gratton@enseeiht.fr

[†]CERFACS, Toulouse, France. Email: selime.gurol@cerfacs.fr

[‡]Université de Toulouse, INP, IRIT, Toulouse, France. Email: ehouarn.simon@enseeiht.fr

[§]NAXYS, University of Namur, Namur, Belgium. Email: philippe.toint@unamur.be

horizon (assimilation window) considered is split into a number (N_{sw}) of time sub-windows, and the problem can be written as

$$\min_{x \in \mathbb{R}^s} J(x) \stackrel{\text{def}}{=} \frac{1}{2} \|x^{(0)} - x_b\|_{B^{-1}}^2 + \frac{1}{2} \sum_{j=0}^{N_{sw}} \left\| \mathcal{H}_j(x^{(j)}) - y_j \right\|_{R_j^{-1}}^2 + \frac{1}{2} \sum_{j=1}^{N_{sw}} \|x^{(j)} - \mathcal{M}_j(x^{(j-1)})\|_{Q_j^{-1}}^2 \quad (1.1)$$

where

- $x = (x^{(0)}, x^{(1)}, \dots, x^{(N_{sw})})^T \in \mathbb{R}^s$ ($s = n(N_{sw} + 1)$) is the control variable (with $x^{(j)} = x(t_j)$),
- $x_b \in \mathbb{R}^n$ is the background given at the initial time (t_0).
- $y_j \in \mathbb{R}^{m_j}$ is the observation vector over a given time interval
- \mathcal{H}_j maps the state vector x_j from model space to observation space
- \mathcal{M}_j represents an integration of the numerical model from time t_{j-1} to t_j
- B , R_j and Q_j are the positive-definite covariance matrices for background, observation and model error, respectively.

The incorporation of possible model errors is achieved by the presence of the third term in the objective function.

As it is the case for the standard 4D-Var (consisting of the first two terms in (1.1)), the general unconstrained nonlinear least-squares problem is solved by applying the Gauss-Newton algorithm [8, 18], which iteratively proceeds by linearizing \mathcal{H} and \mathcal{M} at the current iterate and then, often very approximately, minimizing the resulting quadratic function. A practically crucial question is then how to approximately perform this minimization. As this is a main theme of the present paper, we immediately stress here that this aim (approximate minimization) is often very different from approximate gradient/residual reduction (although they coincide if the minimization is exact). Key factors for selecting a subproblem solver include the choice of a quadratic model (re)formulation (known as the variational formulation), the choice of a preconditioner, the parallelization potential of the resulting algorithm.

Three formulations are available (state, forcing and saddle) and are detailed in the next section. The ‘‘saddle formulation’’, discussed in [13, 12, 11], has recently attracted interest of practitioners because of its appealing potential for parallel computing while still allowing a wide choice of preconditioners. However it is fair to say that numerical experience with this approach remains scarce so far, prompting for a more detailed assessment.

The purpose of the present paper is to propose such an assessment. It will be shown that, left to its own devices, the original algorithm for the saddle formulation may produce erratic results or diverge altogether. To circumvent this problem, a more elaborate variant of the same approach will be proposed which, at the same time, guarantees convergence of the overall Gauss-Newton algorithm and essentially retains the excellent parallelization features of the original method. The numerical performance will be illustrated and compared to that of state and forcing formulations on an assimilation example on the nonlinear Burgers equation and on the two-layers Quasi-Geostrophic (QG) atmospheric model⁽¹⁾ provided within OOPS by the

⁽¹⁾Quasi-geostrophic motion means that, in the horizontal direction of the atmospheric flow, the Coriolis force caused by the rotation of the Earth, and the pressure gradient force are in approximate balance.

European Centre for Medium Range Weather Forecasts (ECMWF). This model is widely used in theoretical atmospheric studies, since it is simple enough for numerical experimentations and yet adequately captures the most relevant large-scale dynamics in the atmosphere. For more details on the QG model, see [11, 12]. This comparison will demonstrate that its parallel computing features help to explain why the new variant often outperforms other approaches. Influence of the choice of preconditioner, detailed operator cost and data organization will also be discussed.

The paper is organized as follows. Section 2 provides the necessary background and notations for the three variational formulations mentioned above, including some of the associated preconditioning issues and a discussion of the parallelization bottlenecks. The potentially problematic behaviour of the original saddle method is then discussed and illustrated in Section 3, and the new algorithmic variants described in Section 4. A comparison of these new variants on the Burgers and QG examples is then proposed in Section 5 and 6, respectively. The special case where the inverse of the correlation matrices is not available is briefly considered in Section 7, while conclusions and perspectives are outlined in Section 8.

Notations. The Euclidean product of the vectors x and y is denoted by $y^T x$ and the induced Euclidean norm of x by $\|x\|$.

2 Problem formulations and preconditioning

As indicated above, the formulation of the subproblem at each iteration of the Gauss-Newton algorithm for solving (1.1) is crucial for good computational performance. If the operators M_j are the linearized \mathcal{M}_j and H_j are the linearized \mathcal{H}_j , this subproblem can be expressed in terms of the increment δx as

$$\min_{\delta x \in \mathbb{R}^s} q_{\text{st}}(\delta x) \stackrel{\text{def}}{=} \frac{1}{2} \|L \delta x - b\|_{D^{-1}}^2 + \frac{1}{2} \|H \delta x - d\|_{R^{-1}}^2 \quad (2.1)$$

where

$$L = \begin{pmatrix} I & & & & & \\ -M_1 & I & & & & \\ & -M_2 & I & & & \\ & & \ddots & \ddots & & \\ & & & & -M_{N_{sw}} & I \end{pmatrix}, \quad (2.2)$$

for suitable ‘‘misfit’’ vectors

$$d = (d_0^T, d_1^T, \dots, d_{N_{sw}}^T)^T \quad \text{and} \quad b = (b_0^T, c_1^T, \dots, c_{N_{sw}}^T)^T,$$

and where

$$H = \text{diag}(H_0, H_1, \dots, H_{N_{sw}}), \quad D = \text{diag}(B, Q_1, \dots, Q_{N_{sw}}) \quad (2.3)$$

and

$$R = \text{diag}(R_0, R_1, \dots, R_{N_{sw}}).$$

(Note the incorporation of the background covariance matrix B in D . Also note that we have eschewed correlation across time windows, as is often done in practice.) The approximate minimization of the quadratic subproblem is itself carried out using a Krylov method (often conjugate gradients [22], GMRES [28] or efficient specialized techniques such as RPCG [21]

or RSFOM [19], see also [17]). These iterative methods typically requires preconditioning for achieving reasonable computational efficiency.

Three variants of the above problem can then be defined. In the form presented above, the formulation is called the “state formulation” and its optimality condition is given by the linear system

$$(L^T D^{-1} L + H^T R^{-1} H) \delta x = L^T D^{-1} b + H^T R^{-1} d. \quad (2.4)$$

Another version (called the “forcing formulation”) may be obtained by making the change of variables $\delta p = L \delta x$, then requiring the solution of the minimization problem

$$\min_{\delta p \in \mathbb{R}^s} q_{\text{fo}}(\delta p) \stackrel{\text{def}}{=} \frac{1}{2} \|\delta p - b\|_{D^{-1}}^2 + \frac{1}{2} \|HL^{-1} \delta p - d\|_{R^{-1}}^2 \quad (2.5)$$

whose optimality condition may now be written as

$$(D^{-1} + L^{-T} H^T R^{-1} H L^{-1}) \delta p = D^{-1} b + L^{-T} H^T R^{-1} d \stackrel{\text{def}}{=} b_{\text{fo}}. \quad (2.6)$$

We immediately note that (2.6) may be obtained as a two-sided preconditioning of (2.4) with L^{-T} and L^{-1} . A third version (the “saddle” formulation) is obtained by transforming the terms in (2.1) in a set of equality constraints and writing the Karush-Kuhn-Tucker conditions for the resulting constrained problem, leading to the large “saddle” linear system

$$\begin{pmatrix} D & 0 & L \\ 0 & R & H \\ L^T & H^T & 0 \end{pmatrix} \begin{pmatrix} \delta \lambda \\ \delta \mu \\ \delta x \end{pmatrix} = \begin{pmatrix} b \\ d \\ 0 \end{pmatrix} \quad (2.7)$$

where the control vector $[\delta \lambda^T, \delta \mu^T, \delta x^T]^T$ is a $(2s + m)$ -dimensional vector. For the sake of brevity, we do not cover the details of this latter derivation here (see [12, 11]): it is enough to view it as an algebraic “lifting” of condition (2.4) since this latter condition is recovered by applying Gaussian block elimination to the first two rows and columns. Unfortunately, this reformulation is only that of the optimality condition and it is unclear whether it can be derived from an associated quadratic minimization problem. Unless exact minimization is considered, this will turn out to be problematic, as we will see below.

We immediately observe that matrix-vector products $u = L^{-1}v$ are sequential, because they are defined by the simple recurrence

$$u_0 = v_0, \quad u_i = v_i + M_i u_{i-1} \quad (i = 1, \dots, N_{sw}) \quad (2.8)$$

(a similar recurrence holds for products with L^{-T}), which is a serious drawback in the context of modern computer architectures for high-performance computing. In this respect, using the forcing formulation can be computationally cumbersome and, even if the state and saddle point formulations allow performing matrix-vector products with L in parallel, their suitable (and often necessary) preconditioners involve the operator \tilde{L}^{-1} , where \tilde{L} is a block bi-diagonal approximation of L within which the matrices M_i are replaced by approximations \tilde{M}_i . The choice of such preconditioners is thus restricted to use operators \tilde{L} whose inversion can be parallelized, or whose preconditioning efficiency is such that extremely few sequential products are requested. Given the structure of L , this limits the possible options (see [15]). We have chosen, for the present paper, to follow [12] and to focus on the two simplest choices:

$$\tilde{M}_i = 0, \quad \text{and} \quad \tilde{M}_i = I. \quad (2.9)$$

Once \tilde{L} is determined, it remains to decide on the complete form of the (left⁽²⁾) preconditioner, depending on which of the formulations (2.4)–(2.7) is used. For the state formulation (2.4), we anticipate the background term to dominate and use the approximate inverse Hessian of the first term, given by

$$S^{-1} = \tilde{L}^{-1} D \tilde{L}^{-T}. \quad (2.10)$$

If the forcing formulation (2.6) is considered, an obvious choice is to use D as preconditioner (thus yielding a system which can be viewed as a low-rank modification of the identity). Its efficiency has been considered in [9, 10]. Finally, the choice is more open for the saddle formulation (2.7). For the sequel of this paper, we consider the preconditioners given by the inverse of the matrices

$$P_M = \begin{pmatrix} D & 0 & \tilde{L} \\ 0 & R & 0 \\ \tilde{L}^T & 0 & 0 \end{pmatrix}, \quad P_B = \begin{pmatrix} D & 0 & 0 \\ 0 & R & 0 \\ 0 & 0 & -S \end{pmatrix} \quad \text{and} \quad P_T = \begin{pmatrix} D & 0 & \tilde{L} \\ 0 & R & H \\ 0 & 0 & S \end{pmatrix}, \quad (2.11)$$

P_M being the inexact constraint preconditioner suggested in [2, 3] and used in [12], and P_B and P_T being the triangular and block-diagonal ones inspired by [1] (see also [33]).

3 The original saddle method

Armed with these concepts and notation, we may now consider the original saddle technique as discussed in [12] (and also used in [14]). It is outlined as Algorithm 3.1, where we define

$$r(\delta\lambda, \delta\mu, \delta x) = \begin{pmatrix} D & 0 & L \\ 0 & R & H \\ L^T & H^T & 0 \end{pmatrix} \begin{pmatrix} \delta\lambda \\ \delta\mu \\ \delta x \end{pmatrix} - \begin{pmatrix} b \\ d \\ 0 \end{pmatrix}.$$

Algorithm 3.1: SADDLE-original (SAQ0)

An initial x_0 is given as well as the correlation matrices D and R , a maximum number of inner iterations n_{inner} and a relative residual accuracy threshold $\epsilon_r \in (0, 1)$. Set $k = 0$. While (not converged):

1. Compute $J(x_k)$ and $g_k = \nabla_x J(x_k)$.
2. Apply the preconditioned GMRES algorithm [28] to reduce $\|r(\delta\lambda, \delta\mu, \delta x)\|$ using one of the left preconditioners given by (2.11). Terminate the GMRES iteration at inner iteration j with $(\delta\lambda, \delta\mu, \delta x)$ if

$$\|r(\delta\lambda, \delta\mu, \delta x)\| \leq \epsilon_r (\|b\| + \|d\|) \quad \text{or} \quad j = n_{\text{inner}}, \quad (3.1)$$

yielding a step $\delta x_k = \delta x$.

3. Set $x_{k+1} = x_k + \delta x_k$ and increment k by one.

⁽²⁾Right preconditioning is also possible, but depends more on the detailed nature and discretizations of the dynamical models, which is why it is not considered here.

In practice the operational constraints often impose moderate values of n_{inner} (a few tens) as well as a small number of outer iterations (ten or less), and the improvement obtained by this fast procedure is often sufficient for producing a reasonable forecast. Thus the concept of convergence should be taken with a grain of salt in this context. However, the monotonicity of the values of $J(x_k)$ underlying the convergence idea remains important as a theoretical guarantee that the method is meaningful from the statistical and numerical points of view, preventing the algorithm to produce unreliable results.

In the state formulation, the monotonic decrease of the $J(x_k)$ is promoted by the fact that the method used to minimize (2.1) (conjugate gradient or one of its variants) is itself a monotonic algorithm. As a consequence, any decrease obtained for (2.1) translates into a decrease for (1.1) provided q_{st} reasonably approximates J in the neighbourhood of x_k , as is often the case (or can be enforced by a trust-region scheme [6]). A similar argument applies for the forcing formulation, where a monotonic algorithm is also used to minimize $q_{\text{fo}}(\delta_p)$ (remember that this formulation can be derived from the state one by suitable two-sided preconditioning). However, the GMRES method used in Algorithm 3.1 merely reduces the residual of the system (2.7) without any variational interpretation. If n_{inner} is large enough for the residual to become “sufficiently” small (producing a “sufficiently” accurate solution of (2.7)), then the equivalence between the optimality conditions (2.4) and (2.7) implies that the decrease in q_{st} at the computed step is comparable to that which would be obtained by minimizing this quadratic model exactly, thereby ensuring a suitable decrease in J . The difficulty is to quantify what is meant by “sufficiently”.

To illustrate this point, let us consider an assimilation problem for the one-dimensional nonlinear Burgers equation involving 100 discretization states over 3000 time steps in 50 time sub-windows and 20 observations per subwindows (the complete description of this problem is given in Appendix A1). We apply Algorithm 3.1 to this problem with the number of inner iterations n_{inner} fixed to 50 and using 10 major Gauss-Newton iterations, $\tilde{M}_i = 0$ and the preconditioners defined by (2.11). We use the abbreviation⁽³⁾ SAQ0-P-M to denote the corresponding algorithmic variants, where P is the choice of preconditioner type in (2.11) and M is the particular choice of the model approximations \tilde{M}_i , which can be either I or 0. Figure 3.1 shows the resulting evolutions of the values of q_{st} (dashed curve) and J (continuous curve) for SAQ0-M-0 and SAQ0-T-0 over all inner iterations⁽⁴⁾, major iterations being indicated by vertical dotted lines and the true minimum value of J by an horizontal thick black line. Several important conclusions follow from the examining this figure.

1. None of the two methods achieves a significant reduction of the gap between $J(x_0) \approx 1.5 \times 10^5$ and the optimal value (≈ 63.11), the version using P_M even diverging slowly.
2. The curves for q_{st} and J differ so little for iterations beyond the first that they are mostly undistinguishable, indicating a good fit between q_{st} and J for moderately small steps. The observed stagnation/divergence may therefore not be blamed on the problem’s nonlinearity.
3. The non-monotonic evolution of both q_{st} and J along inner iterations is very obvious. This is true for both SAQ0-M-0 and SAQ0-T-0 at the first iteration and for SAQ0-M-0 at all subsequent ones. We observe in particular that the value of q_{st} (and that of

⁽³⁾The naming convention will become clearer in the sequel of the paper.

⁽⁴⁾The values of J at inner iterations have been computed for illustration purposes only: they are not needed by the algorithm.

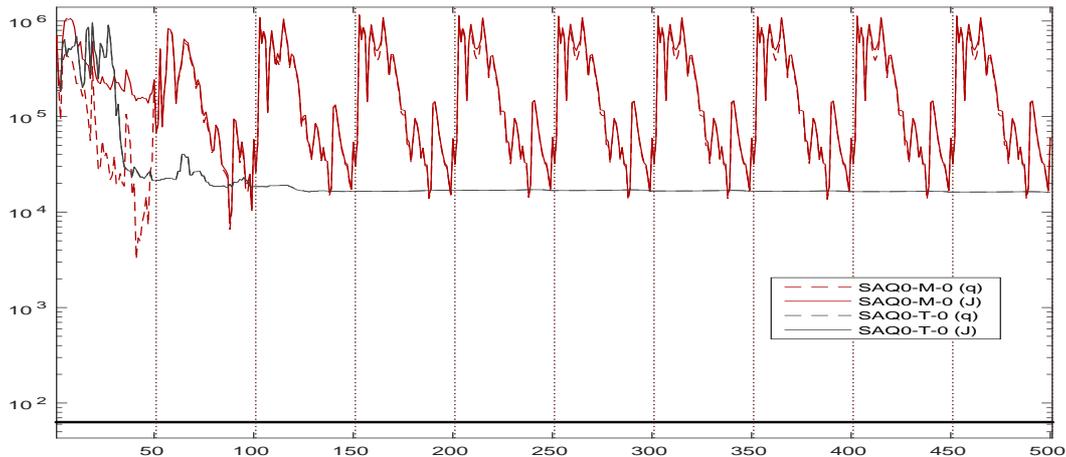


Figure 3.1: Evolution of q (dashed) and J (continuous) as a function of the total number of inner iterations for SAQ0-M-0 and SAQ0-T-0

J) often starts by increasing at the first inner iterations. A stopping rule based on a maximum number of such iterations or on shortening the step therefore essentially relies on luck to produce a decrease in either objectives.

4. A qualitatively similar picture is obtained when using $\tilde{M}_i = I$, the only significant difference being that SAQ0-T-I now (slowly) diverges, despite the fact that the model approximations incorporates more information than for the (marginally more efficient) SAQ0-T-0.

The numerical behaviour of for SAQ0-B-0 (the variant using the block diagonal preconditioner P_B from (2.11)) is not shown in Figure 3.1 because it would be barely visible. As it turns out, it exemplifies to an extreme the fundamental difference between reducing the residual of the system (2.7) and obtaining a decrease in q_{st} or J . In this particular case, the preconditioner at iteration two is good enough to ensure, with a single step of GMRES, a relative reduction of the residual norm of the order of 10^7 , thereby triggering a successful exit from the inner iteration loop. However, since the first step of GMRES is colinear with the initial residual $(b^T, d^T, 0)^T$ (the right-and side of (2.7)), this seemingly excellent step does not alter the values of the state variables at all. Thus $\delta x_k = 0$ for all $k \geq 2$, causing the algorithm to stagnate (this would only show as short horizontal line in the figure). The same undesirable behaviour is also observed for SAQ0-B-I.

Extensive numerical experience with the Burgers assimilation problem indicate that the conclusions drawn from this example are typical of many other problems settings differing by conditioning of the involved correlation matrices, number of time sub-windows or number of observations. Although they obviously remain problem-dependent (as will be shown in Section 5), they show that, in general, *the original saddle method* described in Algorithm 3.1 *is potentially very inefficient or divergent* and that its efficiency might decrease even if some algorithm's ingredients (such as model approximations) are improved⁽⁵⁾.

⁽⁵⁾It is quite remarkable that some practical implementations of the original saddle method actually skip Step 1 of Algorithm 3.1. Since nor the objective function values J nor those of its gradient are ever computed, no convergence guarantee can possibly be given for this technique whose link with an optimization method

4 The globalized SADDLE algorithm

Is it possible to fix this problematic behaviour of the original saddle algorithm? The answer is fortunately positive. Since we know that the desired reduction in q_{st} is obtained if we allow GMRES to fully solve (2.7), the main idea is to adapt the GMRES termination rule so that termination cannot occur before a minimal decrease in q_{st} is obtained. We now detail this strategy as Algorithm 4.1 where we use a generic non-negative sequence $\{\theta_j\} \rightarrow 0$.

Algorithm 4.1: Globalized SADDLE (SAQ1)

An initial x_0 is given as well as the correlation matrices D and R , a target number of inner iterations n_{inner} and a relative residual accuracy threshold $\epsilon_r \in (0, 1)$. A model check frequency $\ell \in \mathbb{N}$ and a model decrease threshold $\epsilon_q \in (0, 1)$ are also given. Set $k = 0$.

While (not converged):

1. Compute $J(x_k)$ and $g_k = \nabla_x J(x_k)$.
2. Apply the preconditioned GMRES algorithm to (2.7), using one of the left preconditioners given by (2.11). At inner iteration j , terminate with $(\delta\lambda, \delta\mu, \delta x)$ if

$$q_{\text{st}}(0) - q_{\text{st}}(\delta x) \geq \max \left[\epsilon_q \min [1, \|g_k\|^2], \theta_j \right] \quad (4.1)$$

or if (2.7) is solved to full accuracy, yielding a step $\delta x_k = \delta x$.

3. Perform a backtracking linesearch [26, p. 37] on J along the direction δx_k , yielding $x_{k+1} = x_k + \alpha \delta x_k$ for some stepsize $\alpha > 0$. Increment k by one.

It is clear that verifying (4.1) requires the (periodic) evaluation of the quadratic model $q_{\text{st}}(\delta x)$, which is an additional computational cost: one needs to apply the L, D^{-1}, H and R^{-1} operators to obtain $q_{\text{st}}(\delta x)$. The GMRES algorithm may also need more than n_{inner} iterations to terminate, potentially increasing its cost further. Note also that the linesearch procedure guarantees that $J(x_{k+1}) \leq J(x_k)$ for all k .

Observe now that, because the covariance matrices are positive-definite, the level set $\{x \in \mathbb{R}^n \mid J(x) \leq J(x_0)\}$ is compact, and thus, using the monotonicity of the algorithm, that there exists a constant $\kappa_g \geq 1$ such that $\|g_k\| \leq \kappa_g$, for all k . On termination of GMRES, we therefore obtain that

$$\epsilon_q \kappa_g^{-2} \|g_k\|^2 \leq q_{\text{st}}(0) - q_{\text{st}}(\delta x_k) = -g_k^T \delta x_k - \frac{1}{2} (L \delta x_k)^T D^{-1} (L \delta x_k) - \frac{1}{2} (H \delta x_k)^T R^{-1} (H \delta x_k).$$

Using the positive-definite character of $\nabla^2 q_{\text{st}}$, the Hessian of q_{st} , we deduce that

$$g_k^T \delta x_k \leq -\epsilon_q \kappa_g^{-2} \|g_k\|^2. \quad (4.2)$$

In addition, the strict convexity of q_{st} ensures that

$$\|\delta x_k\| \leq \frac{2}{\nu_{\min}} \|g_k\|, \quad (4.3)$$

becomes somewhat tenuous.

where $\nu_{\min} > 0$ is the smallest eigenvalue of $\nabla^2 q_{\text{st}}$. Thus (4.2) and (4.3) together guarantee that δx_k is “gradient related” in the sense that

$$g_k^T \delta x_k \leq -\kappa_1 \|g_k\|^2 \quad \text{and} \quad \|\delta x_k\| \leq \kappa_2 \|g_k\| \quad (4.4)$$

for some positive constants $0 < \kappa_1 \leq \kappa_2$. In conjunction with the use of a linesearch, this well-known property of minimization directions is then sufficient to ensure the monotonic decrease of the sequence $\{J(x_k)\}$ and, assuming uniformly bounded condition numbers for $\nabla^2 q_{\text{st}}(0)$ at all major iterations, the theoretical convergence of the outer Gauss-Newton iteration (see [26, Section 3.2], or [6, Section 10.1], for instance). For the test (4.1) to allow for early GMRES termination, it is necessary that ϵ_q is chosen not too small. Yet it should be small enough for (4.1) to be attainable. If it is chosen too large, it may (in the worst case) force GMRES to solve the system (2.7) to full accuracy, in which case (4.1) is guaranteed with $\epsilon_q \kappa_g^{-2} = \|[\nabla^2 q_{\text{st}}]^{-1}\|$ [6, Section 10.1].

While obtaining convergence is often practically out of reach or much too slow in practice, the theoretical guarantee provides a strong reassurance against potentially erratic results. Clearly, this argument also holds for the choice $\theta_j = 0$ for all j . The introduction of that sequence is therefore unnecessary for ensuring mere convergence, but other choices may be instrumental in speeding up decrease. For our Burgers example, the choices

$$\theta_j = \left(\frac{1}{2} q_{\text{st}}(0)\right)^{\max[1, \frac{n_{\text{inner}}}{j}]} - 1 \quad \text{and} \quad \epsilon_g = \frac{1}{100} \quad (4.5)$$

appear to give reasonable results.

We illustrate the behaviour of SAQ1 on the example used in the previous section to highlight the difficulties of SAQ0. Its performance for the three saddle preconditioners of (2.11) is shown in Figure 4.2.

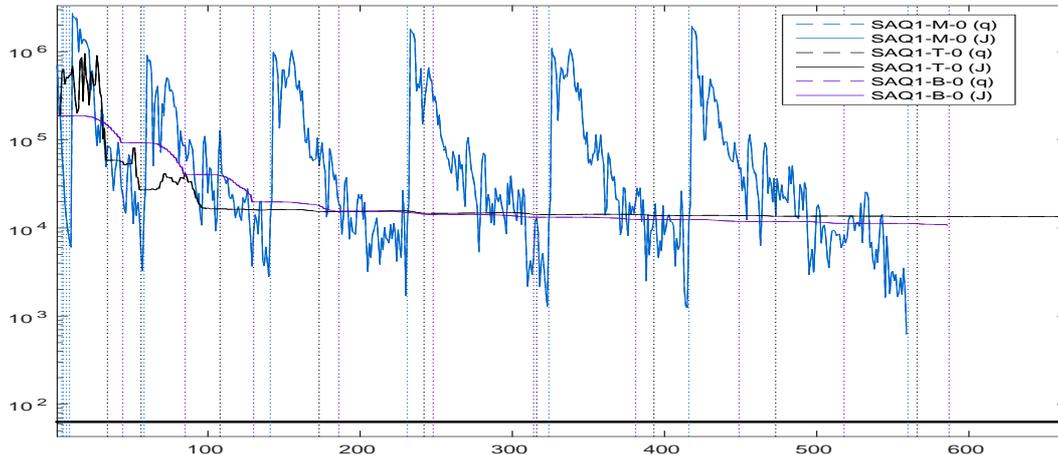


Figure 4.2: Evolution of q (dashed) and J (continuous) as a function of the total number of inner iterations for SAQ1 (Burgers example)

Comparing with the performance of the SAQ0 variants (Fig. 3.1), we may verify that SAQ1 achieves a significant reduction in J in ten Gauss-Newton iterations, albeit at the price of more inner iterations and the cost of one valuation of q_{st} per inner iteration. SAQ1-M-0 oscillates

most, but gets the best decrease, while the performance of SAQ1-T-0 remains disappointing despite the theoretical guarantees and the introduction of the $\{\theta_j\}$.

If the cost of additional evaluations of q_{st} is high, one may try to space them out by checking (4.1) only every $\ell > 1$ inner iteration. The inner iteration termination rule (4.1) then becomes

$$\text{mod}(j, \ell) = 0 \quad \text{and} \quad q_{\text{st}}(0) - q_{\text{st}}(\delta x) \geq \max \left[\epsilon_q \min [1, \|g_k\|^2], \theta_j \right], \quad (4.6)$$

and the resulting algorithms will be denoted by the abbreviation SAQ ℓ in what follows.

The performance of SAQ25-P-0 (that is SAQ25 with preconditioner P in (2.11) and $\tilde{M}_i = 0$) is reported in Figure 4.3 for different preconditioners. As it turns out, the relative performance of the methods including the preconditioners P_T and P_B is again poor. This is unfortunately a constant in our experience and we therefore focus on the use of the more successful P_M only from now on⁽⁶⁾.

We postpone the assessment of the sequential/parallel computational cost of SAQ ℓ as a function of ℓ and the number of computing processes p to Section 5, but we immediately notice that larger values of ℓ may cause SAQ ℓ to require more inner iterations (as termination is checked less often), in turn leading to larger memory and orthogonalisation costs. Thus a value of $\ell \leq n_{\text{inner}}$ seems most reasonable. However, more inner iterations may also result in a better decrease of the quadratic model, and, if the problem is not too nonlinear, of the overall objective function.

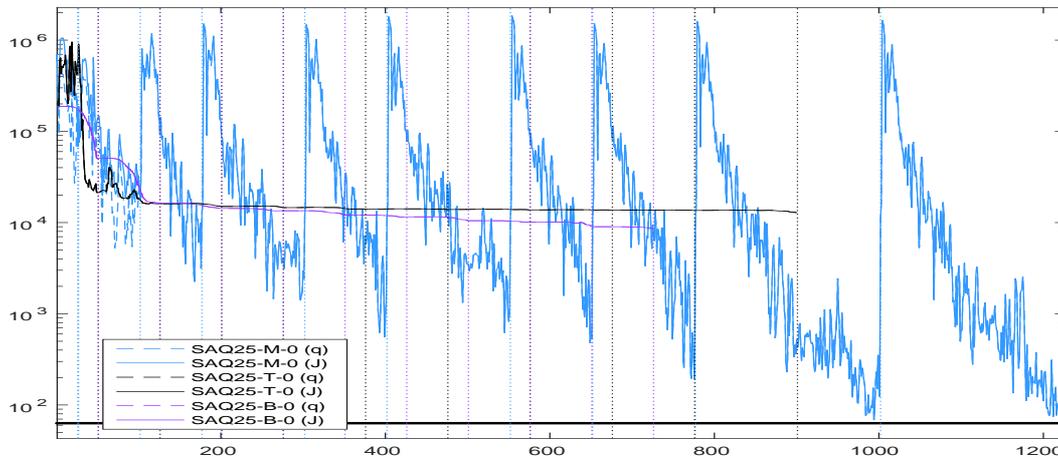


Figure 4.3: Evolution of q (dashed) and J (continuous) as a function of the total number of inner iterations for SAQ25 (Burgers example)

If evaluating q_{st} is so costly that it must be avoided altogether, it is still possible to cure the defects of the original formulation by relying directly on checking the gradient-related property of δx stated by (4.4). However, the resulting algorithm then suffers more directly from the need to estimate κ_1 and κ_2 a priori, and moreover appears to be significantly slower

⁽⁶⁾We remark that this is not in contradiction with a comment in [33] that, *for the symmetric case*, extended preconditioner formulations are unlikely to be efficient because of an alternating property in MINRES. Indeed, the square root of P_M^{-1} is not well defined because it is indefinite, and the left-preconditioned system matrix is no longer symmetric (which is why GMRES is used).

than the SAQ ℓ versions, mostly because it cannot incorporate the forcing sequence $\{\theta_j\}$. It will therefore not be discussed here.

5 Numerical comparisons on the Burgers example

We now turn to a comparative evaluation of the computational costs associated with the SAQ ℓ algorithms, as well as state-of-the-art algorithms for the state and forcing variational formulation. Our evaluation attempts to provide conclusions in a context where parallel computing is available.

The comparison will involve several implementations of the state and forcing formulations. The first set (which, for now, we call the ST algorithms) uses a standard Gauss-Newton algorithm where the quadratic model q_{st} is minimized at each outer iteration using the left-preconditioned Full Orthogonalization Method (FOM) [27] expressed in the inner product defined by the inverse preconditioner⁽⁷⁾, which is in general preferable [19] to the conjugate-gradient algorithm with reorthogonalization (the practitioner's most common choice so far). The preconditioner used is given by $S^{-1} = \tilde{L}^{-1}D\tilde{L}^{-T}$, the approximate inverse of the Hessian of the first term in (2.1). The implementations of the forcing formulation (which we call the FO algorithms) use the same left-preconditioned FOM method to minimize q_{fo} given by (2.5) as a function of δp using D as a preconditioner and deduce δx from

$$\delta x = L^{-1}\delta p. \quad (5.1)$$

All these methods use a standard trust-region scheme [6] for ensuring theoretical convergence. In these algorithms, the value of q_{st} is readily available at the price of a single inner product at the end of each inner iteration (remember that $q_{\text{fo}}(\delta p) = q_{\text{st}}(L^{-1}\delta p)$). It is therefore most coherent to terminate the inner iterations when (4.6)-(4.5) holds or, in the worst case, if the relevant system has been solved to full accuracy (i.e. residual norm below 10^{-12} in our tests).

This lead us to relatively large set of algorithms, which differ by four possible choices: the variational formulation (SA, ST or FO), the frequency ℓ of the quadratic model check for terminating inner iterations in (4.6), the type of preconditioner used (inexact constraint M, triangular T or block-diagonal B for SA, the Schur complement S for ST, and the block-diagonal D for FO), and finally the choice of model approximation \tilde{M}_i used in defining \tilde{L} (0 or I, for SA and ST only). Using a naming convention coherent with that already introduced for the SAQ ℓ algorithms, we will, in the sequel, consider the algorithmic variants whose names are of the form AAQ ℓ -P-M, where AA denotes the variational formulation, Q ℓ the frequency of the check for quadratic decrease in (4.6), P the preconditioner type and M the choice of \tilde{M}_i , as summarized in Table 5.1.

Thus algorithm STQ15-S-0 uses the state formulation, checks for sufficient quadratic decrease every 15-th inner iteration, uses the Schur complement preconditioner in which \tilde{L} is defined using $\tilde{M}_i = 0$. In order to limit the number of variants, we have chosen $\ell \in \{1, 15, 25, 50\}$. We also introduced the 'n' preconditioner type, which stands for not using preconditioning at all. Altogether, we therefore obtain a set of 36 different algorithms⁽⁸⁾.

⁽⁷⁾In order to handle the unsymmetric matrix resulting from left-preconditioning, see Algorithm A.1 in Appendix A3.

⁽⁸⁾SAQ1-n, SAQ1-M-0, SAQ1-M-I, SAQ1-M-M, SAQ15-n, SAQ15-M-0, SAQ15-M-I, SAQ15-M-M, SAQ25-n, SAQ25-M-0, SAQ25-M-I, SAQ25-M-M, SAQ50-n, SAQ50-M-0, SAQ50-M-I, SAQ50-M-M, STQ1-n, STQ1-S-0, STQ1-S-I, STQ1-S-M, STQ15-n, STQ15-S-0, STQ15-S-I, STQ15-S-M, STQ25-n, STQ25-S-0, STQ25-S-I, STQ25-S-M, STQ50-n, STQ50-S-0, STQ50-S-I, STQ50-S-M, FOQ1-D, FOQ15-D, FOQ25-D, FOQ50-D.

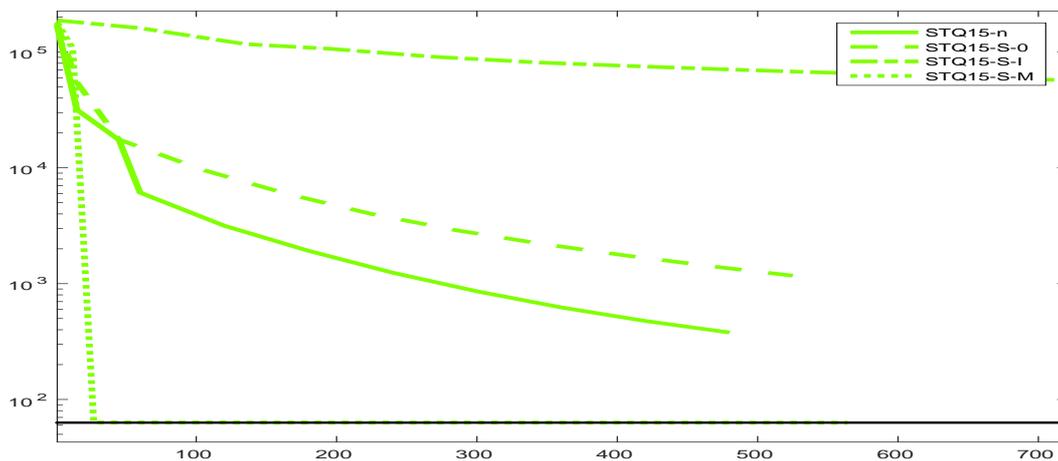
Var. Form. (AA)	quad. check freq. (Qℓ)	P type (P)	\tilde{M}_i (M)
SA	Q1, Q15, Q25, Q50	M, T, B, n	0, I, M
ST	Q1, Q15, Q25, Q50	S, n	0, I, M
FO	Q1, Q15, Q25, Q50	D, n	

Table 5.1: Naming conventions for the considered algorithmic variants

Before we embark on the comparison of computational costs, some further comments on the various methods are in order. The first is that the use of preconditioning with the state formulation is not without risks. Indeed, [16] provides an analysis of the inherent difficulty of preconditioning weighted least-squares problems, caused by the interplay between the eigenstructures of \tilde{L} and D . To illustrate how problematic this can be, we borrow the following illustrative example from this latter paper: let $\alpha > 1$ be a parameter and

$$L = \begin{pmatrix} 1 & 0 \\ \alpha & 1 \end{pmatrix}, \quad \tilde{L} = \begin{pmatrix} 1 & 0 \\ 2 + \alpha & 1 \end{pmatrix} \quad \text{and} \quad D = \begin{pmatrix} \alpha & 0 \\ 0 & 1 \end{pmatrix}.$$

Then it can be verified that $\tilde{L}^{-1}\tilde{L}^{-T}$ is a good preconditioner of $L^T L$ in the sense that the condition number of $\tilde{L}^{-1}\tilde{L}^{-T}L^T L$ is finite for all α (its value is equal to 33.97 for all $\alpha > 1$) while that $\tilde{L}^{-1}D\tilde{L}^{-T}L^T D^{-1}L$ tends to infinity when α grows. This discussion suggests that a comparison between the various choices of \tilde{M}_i within the state formulation can be useful for each specific problem. Figure 5.4 and other tests not reported here indicate that using the choice $\tilde{M}_i = I$ (STQ15-S-I) is, in the Burgers example, very inefficient compared with the three other choices. By contrast, the choice of $\tilde{M}_i = M_i$ (STQ15-S-M) expectedly yields maximal accuracy in a very small number of inner iterations.


 Figure 5.4: Evolution of J as a function of the number of inner iterations for STQ15-n, STQ15-S-0, STQ15-S-I, and STQ15-S-M (Burgers example)

Other significant differences exist between STQℓ and FOQℓ algorithms on one hand, and SAQℓ algorithms on the other hand.

1. At each inner iteration, the STQ ℓ algorithm compute a matrix-vector product with the Hessian of q_{st} , which involve using the D^{-1} operator. The situation is better for the FOQ ℓ -D variants, since (formally) the D^{-1} term in the Hessian of (2.5) is in this case premultiplied by D , which is readily simplified not to involve D^{-1} at all.
2. The FOQ ℓ -D methods also require products with L^{-1} and L^{-T} at each inner iteration. Unfortunately, these products are inherently sequential (see [15]) and therefore potentially very costly. However, the quantity $L^{-1}\delta p$ can be recurred within the FOM algorithm itself at marginal cost, making the backsolve (5.1) unnecessary. A description of the resulting FOM algorithm (stripped from its trust-region enforcing features) is given as Algorithm A.2 in Appendix A3.
3. The use of the preconditioners $S^{-1} = \tilde{L}^{-1}D\tilde{L}^{-T}$ for the STQ ℓ methods is fully parallelizable, given our simple choices for \tilde{M}_i .

Note that all methods use the D^{-1} operator for jointly computing the values of $J(x_k)$ and g_k (once per major iteration) as well as for the periodic evaluations of q_{st} every ℓ -th inner iteration. These remarks suggest that the two main parameters influencing the parallel computing costs of the considered methods are the (parallelizable) cost of computing D^{-1} and that the purely sequential ones of computing L^{-1} and L^{-T} . Due to the form of L in (2.2), the two latter costs are bounded by a small multiple of that of integrating the full nonlinear model over the complete assimilation window. Our objective is therefore to assess the efficiency of the various algorithms in a parametric study varying the cost of applying D^{-1} .

The model used for the execution of parallel tasks is fairly simple, but it hoped that it can nevertheless be sufficient for the broad type of analysis presented. Let us denote by c_{op} the cost of evaluating (possibly in parallel) the operator op . Then the cost of evaluating the tasks of costs c_1, \dots, c_k in parallel on p parallel computing processes is approximated by

$$\pi_p(c_1, \dots, c_k) = \max \left[\left\lceil \frac{k}{p} \right\rceil \frac{1}{k} \sum_{j=1}^k c_j, \max_{j=1, \dots, k} c_j \right].$$

Taken alone, this approximation is not enough to provide the description of a parallel computing environment, as it is crucial to consider the impact of communications, which is beyond the scope of the present paper. In what follows, we consider two complementary cases, and discuss their associated parallel computing costs successively.

5.1 A fully MPI approach

A first, if somewhat restrictive, setting is to assume that the computation is performed in p MPI processes allowing the parallelizable operators to be executed simultaneously for different time windows. In particular, this implies that parallel products with L_i , L_i^T , L_i^{-1} , L_i^{-T} , D_i , D_i^{-1} and H_i are excluded because each of them already uses the full available parallelism. By the same argument, parallel products with R_i , R_i^{-1} and H_i^T are also banned.

In this framework, the cost of evaluating q_{st} is given by

$$c_q = c_L + c_{D^{-1}} + c_H + c_{R^{-1}}, \quad (5.2)$$

while that of evaluating J and its gradient is

$$c_J = c_{\mathcal{M}} + c_{\mathcal{H}} + c_{L^T} + c_{D^{-1}} + c_{H^T} + c_{R^{-1}}. \quad (5.3)$$

Note that the quantities $D^{-1}b$ and $R^{-1}d$ are both available once J and its gradient have been evaluated. We first investigate the computational costs of the components of GMRES, for which it can be verified that the cost of a Krylov iteration for (2.7) is

$$c_{K,sa} = c_L + c_D + c_{L^T} + c_H + c_{H^T} + c_R \quad (5.4)$$

while that of applying the saddle preconditioner P_M of (2.11) involves

$$c_{S^{-1}} = c_{\tilde{L}^{-T}} + c_D + c_{\tilde{L}^{-1}}, \quad (5.5)$$

(the cost of applying $S^{-1} = \tilde{L}^{-1}D\tilde{L}^{-T}$) and is

$$c_{P_M} = c_{S^{-1}} + c_{R^{-1}}. \quad (5.6)$$

The cost of applying n_o outer iterations of SAQ ℓ -M for a total of n_i inner iterations may then be approximated by

$$c_{SAQ\ell-M} \approx n_o(c_J + c_{P_M}) + n_i(c_{K,sa} + c_{P_M}) + \frac{n_i}{\ell}c_q, \quad (5.7)$$

the second term in the first bracket of the right-hand side accounting for the preconditioning, at each outer iteration, of the initial inner-iteration residual, and the last term accounting for the periodic evaluations of q_{st} within the termination criterion. Similarly, it can be verified that, for the state formulation,

$$c_{K,st} = c_L + c_{D^{-1}} + c_{L^T} + c_H + c_{H^T} + c_{R^{-1}} \quad (5.8)$$

and thus, using (5.5) and

$$c_{rhs,st} = c_{L^T} + c_{H^T} \quad (5.9)$$

the cost of computing the right-hand side of (2.4), that

$$c_{STQ\ell-S} \approx n_o(c_J + c_{rhs,st} + c_{S^{-1}}) + n_i(c_{K,st} + c_{S^{-1}}). \quad (5.10)$$

Finally, for the forcing formulation,

$$c_{K,fo} = c_D + c_{L^{-1}} + c_H + c_{R^{-1}} + c_{H^T} + c_{L^{-T}}, \quad c_{rhs,fo} = c_{L^{-T}} + c_{H^T} \quad (5.11)$$

and

$$c_{FOQ\ell} \approx n_o(c_J + c_{rhs,fo} + c_D) + n_i c_{K,fo}. \quad (5.12)$$

We next assign approximate costs for all building blocks other than D^{-1} , where one unit of cost is given by the integration of the model \mathcal{M} on the complete time window⁽⁹⁾. Assuming p computing processes are available and defining $e_{N_{sw}}$ to be the vector of all ones and length N_{sw} , let

$$c_{\mathcal{M}} = 1, \quad c_{\mathcal{H}} = \frac{1}{20 N_{sw}} \pi_p(e_{N_{sw}}), \quad c_D = \frac{1}{2 N_{sw}} \pi_p(e_{N_{sw}}), \quad c_R = \frac{1}{100 N_{sw}} \pi_p(e_{N_{sw}}), \quad (5.13)$$

$$c_{R^{-1}} = \frac{1}{100 N_{sw}} \pi_p(e_{N_{sw}}), \quad c_H = \frac{1}{10 N_{sw}} \pi_p(e_{N_{sw}}), \quad c_{H^T} = \frac{1}{10 N_{sw}} \pi_p(e_{N_{sw}}), \quad (5.14)$$

⁽⁹⁾The cost values are based on discussions with practitioners.

$$c_L = \frac{2}{N_{sw}} \pi_p(e_{N_{sw}}), \quad c_{L^T} = \frac{4}{N_{sw}} \pi_p(e_{N_{sw}}), \quad c_{L^{-1}} = 2, \quad \text{and} \quad c_{L^{-T}} = 4, \quad (5.15)$$

where we have used the block-diagonal structure of D , R , R^{-1} , H , H^T , L and L^T to allow their costs to decrease with p . Using these admittedly fairly rough approximations, we may re-analyze the behaviour of the SAQ ℓ algorithms this time as a function of computational effort.

For the specific choice $c_{D^{-1}} = 0.5$, Figure 5.5 corresponds to Figure 4.2 where the horizontal axis now indicates sequential computational costs, respectively (instead of inner-iteration counts) and where the evolution of the nonlinear cost J is only shown (in thicker lines) between major iterations. The linesearch is active at the first major iteration for all reported variants except SAQ1-M-0, as is shown by the nearly vertical lines in the top left of the graph.

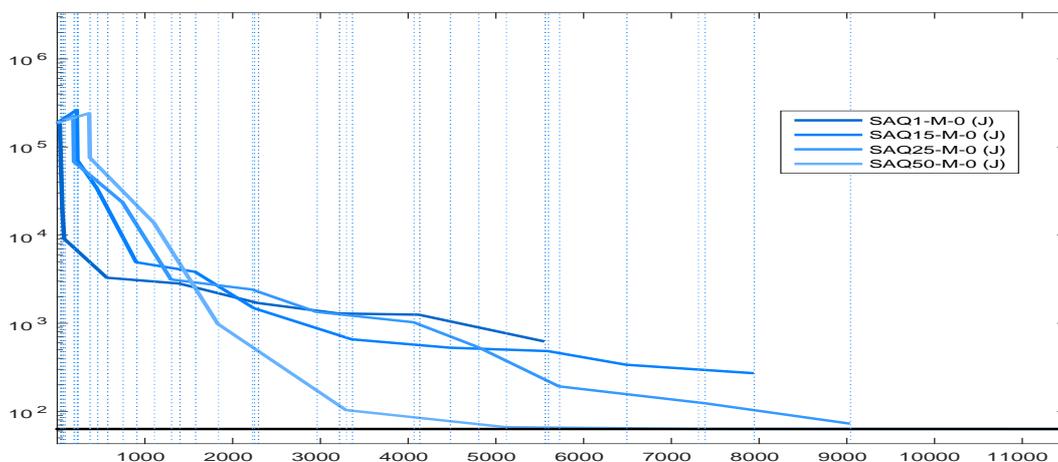


Figure 5.5: Evolution of J as a function of the sequential computational cost, for SAQ ℓ -M-0, ($\ell = 1, 15, 25, 50$) using $c_{D^{-1}} = \frac{1}{2}$ (Burgers example, fully MPI model)

If we are now interested in plotting the parallel computational cost for the same algorithms, the picture looks entirely similar (as all SAQ ℓ -M-0 scale in the same way), but the maximal total cost appearing on the horizontal axis shrinks from 11475 to 542 for 50 computing processes, a 21-fold speedup. This illustrates the excellent parallelization potential of the SAQ ℓ -M-I methods, despite the chaotic (but controlled) evolution of the quadratic model's values. Similar plots and number could be presented for the SAQ ℓ -M-I and STQ ℓ -S algorithms, all showing reasonable parallelization potential. As expected due to the use of the inherently sequential operators L^{-1} and L^{-T} , no such gains can be obtained with the FOQ ℓ -D variants, whose costs only vary marginally with p .

We are now left with the question of choosing a solution algorithm among our 36 SAQ ℓ , STQ ℓ and FOQ ℓ variants ($\ell = 1, 15, 25, 50$), depending on the relative costs of applying D^{-1} . To answer this question, we first applied each of the 36 methods to the Burgers example for $c_{D^{-1}} \in [\frac{1}{2}, 10]$. We then discarded all methods for which the total decrease in J differed by a factor more than $\rho \in (0, 1)$ of the optimal decrease, that is

$$J(x_0) - J(x_f) > \rho(J(x_0) - J(x_*))$$

where $J(x_*)$ was obtained by running STQ1-S-M to full accuracy and x_f denotes the final value of x resulting from the application of the algorithm. We finally selected the method for

which the computational cost was least for $p = 1, 10, 25, 50$. The maps indicating the winning method for each pair (c_{D-1}, ρ) with $\rho \in [10^{-1}, 10^{-3}]$ and each p are given in Figures 5.6. Each such map is accompanied with a picture of the surface of the minimum computational costs over all (c_{D-1}, ρ) pairs. The legend provides a correspondance between colors on the maps (the left graph in each box) and the algorithmic variants⁽¹⁰⁾.

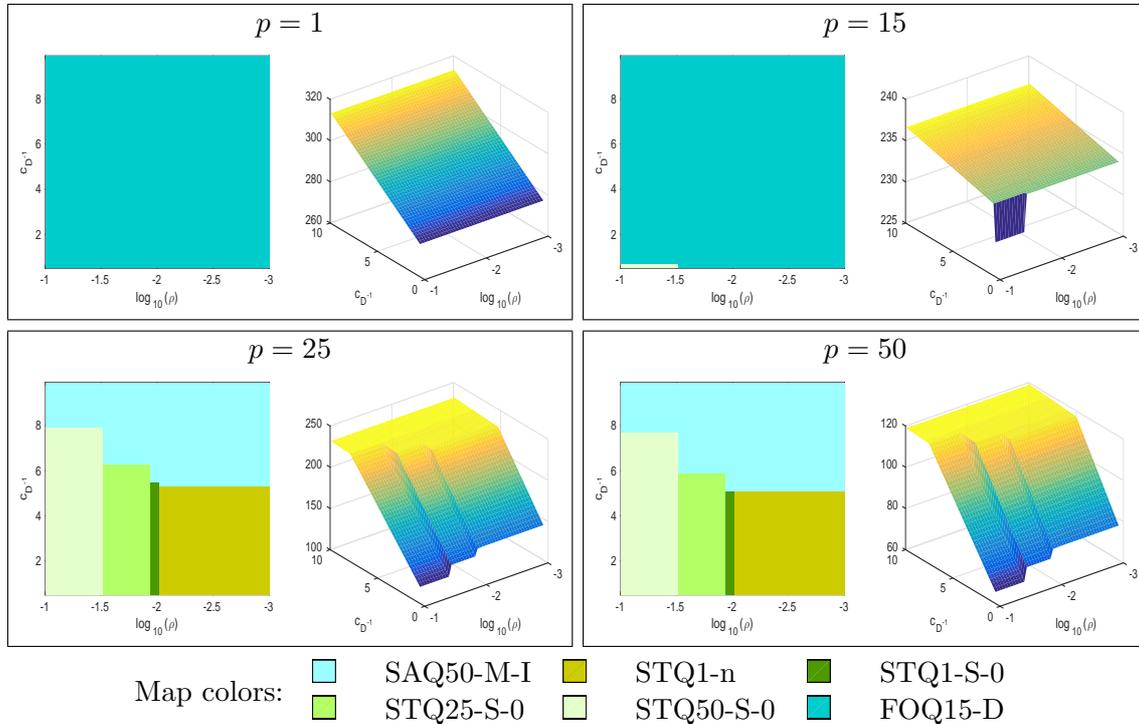


Figure 5.6: Best algorithmic variants as a function of c_{D-1} the number of computing processes p and the reliability factor ρ (Burgers example, fully MPI model)

Several conclusions *for the Burgers example* may be drawn from these maps and cost surfaces.

- Using the forcing formulation dominates all other variants when computations are sequential, for a computing cost with maximal values close of 313.5. We note the proportionally faster increase of the total cost with c_{D-1} compared to requiring higher accuracy. It is interesting that the variant FOQ15-D is best, indicating that terminating the inner iterations as soon as possible (as is the case with FOQ1-D) may be sub-optimal when the preconditioner is excellent.

The usefulness of the forcing formulation clearly decreases when the number of computing processes increases, as expected.

- Using a saddle-based algorithm clearly supposes a high value of c_{D-1} and the availability of several computing processes. The frequency $\ell = 50 = n_{\text{inner}}$ appears to provide the

⁽¹⁰⁾The colors used by MATLAB© for the minimum cost surfaces at the right of each box are meaningless here.

best compromise between good decrease on the quadratic model and excessive number of inner iterations.

3. When $p = 25$ or $p = 50$ and c_{D-1} is moderate, the algorithms using the state formulation dominate with a frequency ℓ diminishing for increasing accuracy, the unpreconditioned version being suitable for maximum accuracy. This is coherent with our comments in the beginning of this section and Figure 5.4, as is the fact that none of the STQ ℓ -S-I variants ever appears on the podium of best methods.
4. Compared to the sequential case, the most parallel of our scenarii ($p = 50$) provides a reduction of computational costs from 315 to 119, which corresponds to a speed-up of approximately 2.6.
5. The additional computing cost necessary for obtaining improved accuracy is negligible, irrespective of the number of computing processes.

5.2 An hybrid MPI/OpenMP approach

Let us now assume that a more elaborate parallel computing environment is considered, such as an hybrid MPI/OpenMP system, where we assume that p processes are available, each of which with two computing cores. This means that we may now apply two time-parallel operators simultaneously. Then the computing costs (5.2)-(5.10) may be rewritten as

$$c_q = \pi_2(c_L + c_{D-1}, c_H + c_{R-1}), \quad (5.16)$$

$$c_J = c_M + c_H + \pi_2(c_{LT} + c_{D-1}, c_{HT} + c_{R-1}), \quad (5.17)$$

$$c_{K,sa} = \pi_2(c_L + c_D + c_H, c_{LT} + c_R + c_{HT}), \quad (5.18)$$

$$c_{S-1} = c_{\tilde{L}-T} + c_D + c_{\tilde{L}-1}, \quad c_{P_M} = \pi_2(c_{S-1}, c_{R-1}). \quad (5.19)$$

$$c_{SAQ\ell-M} \approx n_o(c_J + c_{P_M}) + n_i(c_{K,sa} + c_{P_M}) + \frac{n_i}{\ell}c_q, \quad (5.20)$$

$$c_{K,st} = \pi_2(c_L + c_{D-1} + c_{LT}, c_H + c_{R-1} + c_{HT}), \quad c_{rhs,st} = \pi_2(c_{LT}, c_{HT}), \quad (5.21)$$

$$c_{STQ\ell-S} \approx n_o(c_J + c_{rhs,st} + c_{S-1}) + n_i(c_{K,st} + c_{S-1}), \quad (5.22)$$

$c_{K,fo}$, $c_{rhs,fo}$ and $c_{FOQ\ell-D}$ being unmodified.

We may then repeat our experiments in this new setting, which allows further gains compared to the fully MPI case, as shown in Figure 5.7 (note the doubly logarithmic axis). Although the improvement by switching from the fully MPI model to the hybrid MPI/OpenMP model is far from negligible, we nevertheless note that most of the advantage obtained by parallel processing is due to the parallelization in time, with a very strong correlation with N_{sw} , the number of subwindows (see (5.13)(5.15)). This already apparent at the very beginning of the computation, as evaluating $J(x_0)$ already makes a significant difference (both parallel runs of the algorithm are completed before $J(x_0)$ is evaluated in the sequential mode). Figure 5.8 then illustrates how these gains in computational costs are translated in the new 'best method' maps/minimum cost surfaces.

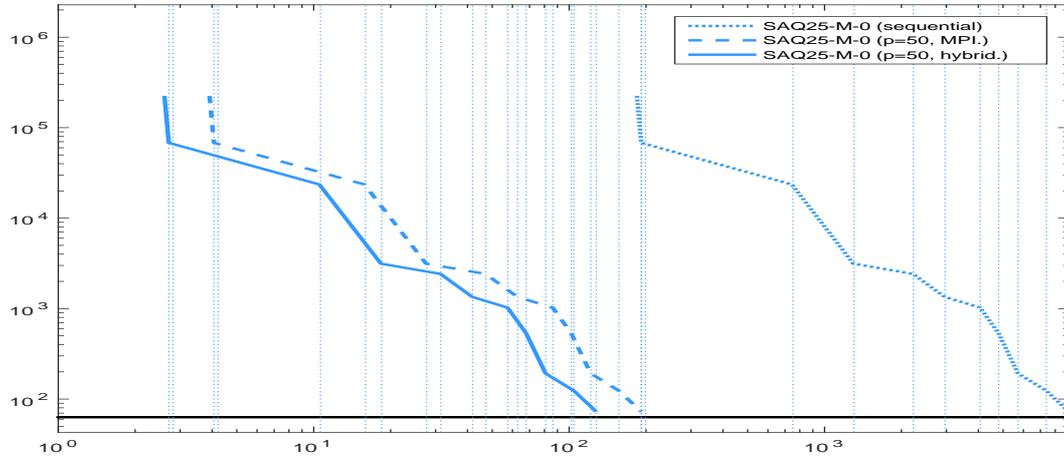


Figure 5.7: Evolution of J as a function of computational cost for SAQ25-M-0 for $c_{D-1} = \frac{1}{2}$ in the sequential, fully MPI and hybrid MPI/OpenMP settings (Burgers example)

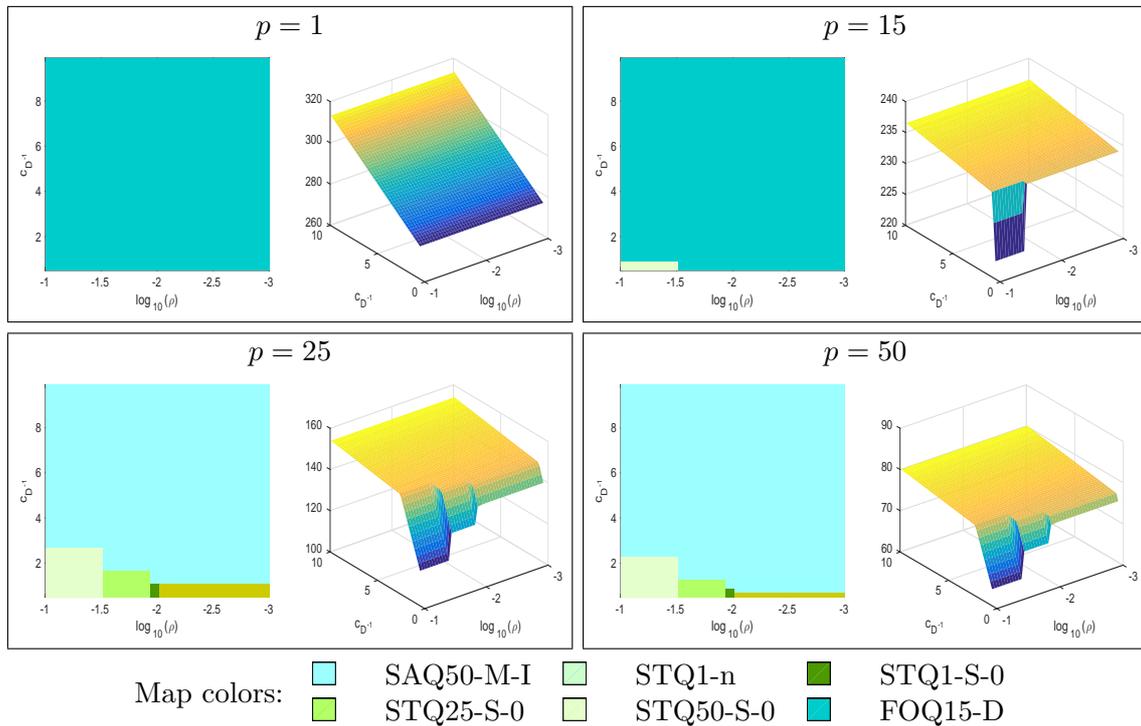


Figure 5.8: Best algorithmic variants as a function of c_{D-1} the reliability factor ρ and the number of computing processes p (Burgers example, hybrid MPI/OpenMP model)

This figure broadly confirms and amplifies the trends already present in Figure 5.6, with the SAQ50-M-I method becoming more important especially when p grows. The reduction of total computational cost is also noticeable, the smallest cost now being approximately 80 (which corresponds to a speed-up close to 4).

6 Numerical comparisons on the QG example

We now turn to the results obtained for the two-layers ECMWF QG example, also using $n_{\text{inner}} = 50$ and at most 5 Gauss-Newton iterations (see Appendix A2 for a more complete description of the problem). We first observe in Figure 6.9 that the fields of interest do evolve (relatively) slowly over the complete assimilation time, and thus even more so within each of the 48 subwindows considered. Hence we may expect $\tilde{M}_i = I$ to be a reasonable approximation of M_i , at variance with the Burgers case (see Figure A.14). In particular, the caveat on using STQ ℓ -S-I may no longer apply.

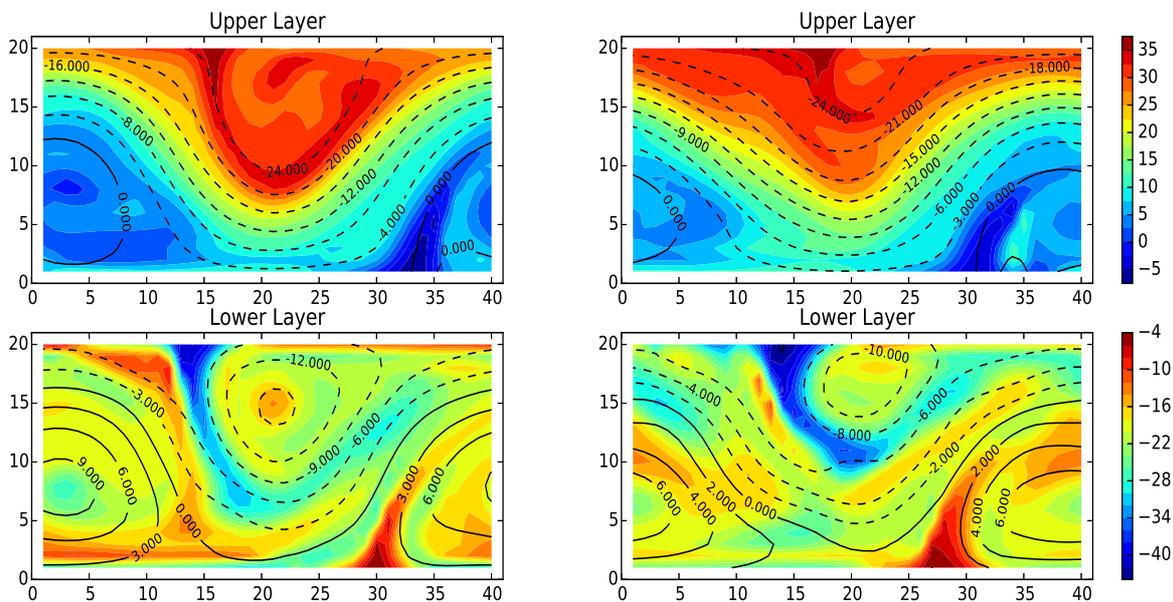


Figure 6.9: Initial (left) and final (right) forecasts using the QG model

Keeping this in mind, we next examine the performance of the original saddle methods SAQ0-M and SAQ0-M and compare them, first in terms of number of inner iterations, for the parallelizable preconditioners (2.9), to competing variants such as SAQ15 or STQ15. The outcome is presented in Figure 6.10 for the choice $\tilde{M}_i = 0$ and Figure 6.11 for the choice $\tilde{M}_i = I$.

We see in the first of these figures that the simply preconditioned SAQ0-M-0 performs relatively well, albeit a bit slowly compared to the globalized saddle SAQ15-M-0 and the state algorithm STQ15-S-0. The dominance of this latter formulation with STQ15-S-I is even more obvious in the second figure, SAQ15-M-I and SAQ0-M-I following the same initial curve, the latter then levelling off and diverging slowly.

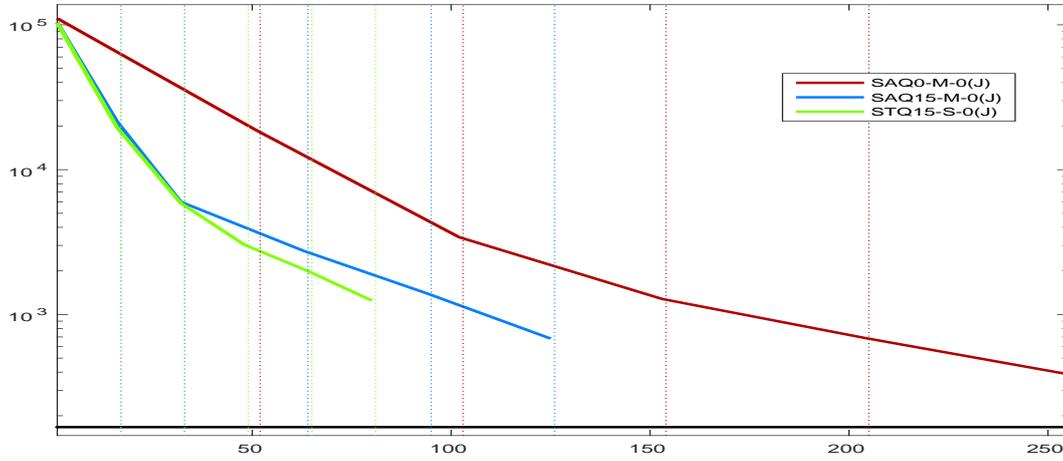


Figure 6.10: Evolution of q (dashed) and J (continuous) as a function of the number of inner iterations for SAQ0-M-0, SAQ15-M-0 and STQ15-S-0 (QG example)

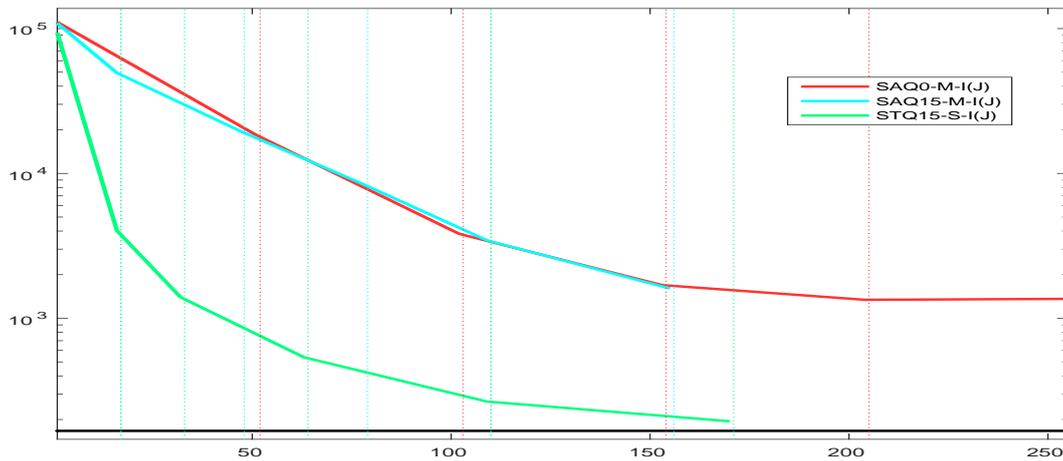


Figure 6.11: Evolution of q (dashed) and J (continuous) as a function of the number of inner iterations, for SAQ0-M-I, SAQ15-M-I and STQ15-S-I (QG example)

We may now again compute the same ‘best method’ maps and minimum cost surfaces, mimicking our analysis for the Burgers case, yielding Figures 6.12 for the fully MPI model, and Figure 6.13 for the hybrid MPI/OPenMP one.

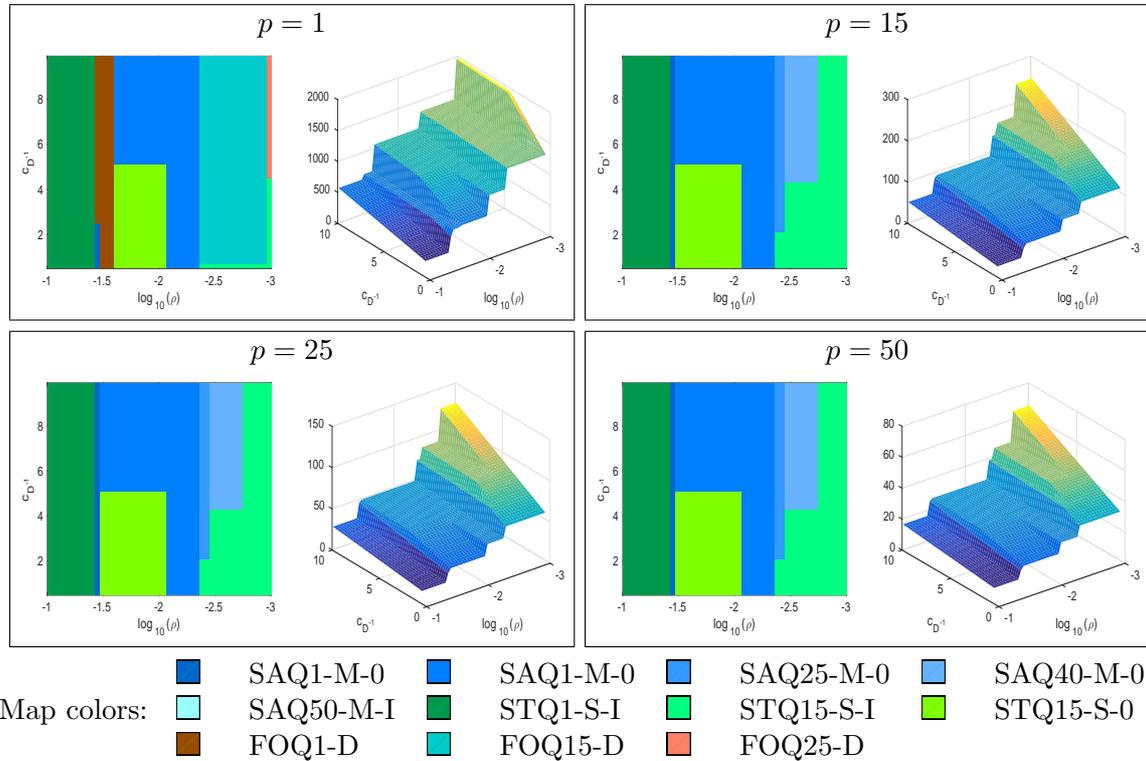


Figure 6.12: Best algorithmic variants as a function of $c_{D^{-1}}$ the number of computing processes p and the reliability factor ρ (QG example, fully MPI model)

The situation is much more intricate here than for the Burgers example. In particular, the competition between state and forcing formulations is very tight in the sequential case. As above, the latter steps being advantageous when the number of computing processes increases. The maximal computational costs now range from approximately 1977 ($p = 1$) to 65 ($p = 50$), giving an excellent speed-up of 30. Looking at the ‘‘best-method’’ maps in more detail, we see that state-based algorithms seems to obtain a better decrease quickly, then being outperformed by saddle-based methods, for finally nevertheless taking over (using $\tilde{M}_i = I$) for the more stringent accuracy requirements. We also note by looking at the minimum-cost surfaces that such requirements come at a significant computational cost (irrespective of p), in contrast with what was observed for the Burgers example. The original saddle algorithm, which we kept in the comparison here, is never the best method.

7 Avoiding the use of D^{-1}

As already noted, the cost of the D^{-1} operator may vary considerably from application to application. In some oceanographic models, D is computed using a diffusion operator which is integrated using an implicit scheme [34, 25, 20, 35]. This makes the cost of D^{-1} very comparable to that of D . In some other applications, such as atmospheric modelling

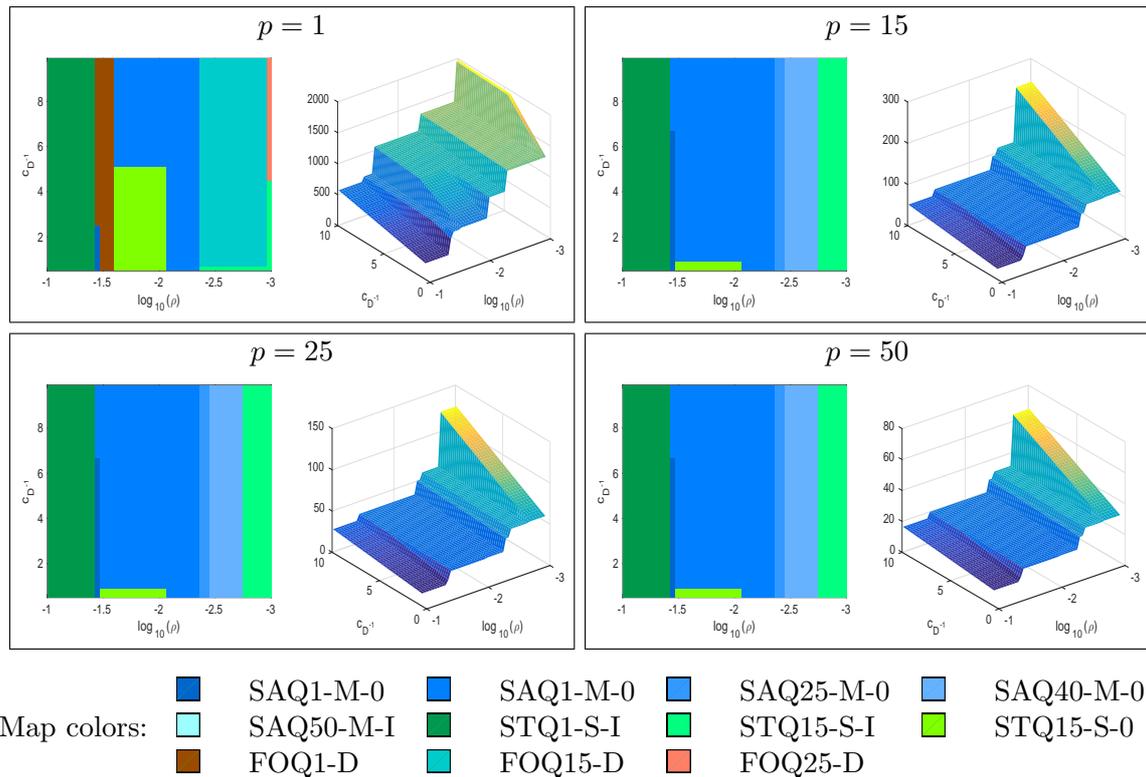


Figure 6.13: Best algorithmic variants as a function of $c_{D^{-1}}$ the number of computing processes p and the reliability factor ρ (QG example, hybrid MPI/OpenMP model) and weather forecasting, the D operator may involve more complicated elements, such as localization schemes [32, 5, 24], which makes applying D^{-1} potentially more costly. The authors are aware of the strong reluctance of practitioners in these areas to even provide the D^{-1} operator at all. If that is the case, and if one nevertheless desires to enjoy the security of a solid global convergence theory at an acceptable cost, it is possible to use an approximate D^{-1} operator, for instance by computing $y = D^{-1}x$ by approximately solving the linear system $Dy = x$. Various iterative methods can be considered for this task, including conjugate-gradients or FOM, which should already be available in the data assimilation system. In addition, as in the Burgers and QG examples, parallelism can often be exploited to make this computation efficient. We have however met two difficulties when experimenting with the idea. The first is that a very inaccurate solution of the linear system may result in an unsymmetric D^{-1} operator, which then empties the very state formulation of its meaning (in addition to causing numerical havoc). The second is that too inexact solutions may also slow down the convergence of $STQ\ell$ significantly, although the saddle-based algorithms seem more robust. We have however found that a relatively modest number of conjugate gradient iterations is very often sufficient to reach this accuracy level. In preliminary tests on the Burgers example with 25 unpreconditioned CG iterations, the SAQ50-M-I algorithm turned out to be the best choice for both parallel computing models, and the parallel computing cost ($p = 50$) increased by less than a factor three compared to using the exact D^{-1} (with $c_{D^{-1}} = c_D \in [\frac{1}{2}, 10]$). Very similar conclusions can be reached when applying the same strategy (with 20 unpreconditioned CG iterations) to the QG example. As number of outer and inner iterations differ only very marginally from that observed when using the exact D^{-1} and since the cost of 20 products with D brings $c_{D^{-1}} = 10$ to the top of the range considered

for exact D^{-1} , the results are essentially identical to those obtained for $c_{D^{-1}} = 10$ in Section 6.

There is little doubt that preconditioning and problem specific tuning would reduce the cost of the approximate D^{-1} even further. Our experiments therefore show that, even if the D^{-1} operator is unavailable, considering the SAQ ℓ or STQ ℓ algorithms may be the best option, irrespective of the number of computing processes and of the parallel computing model.

We finally note that using the forcing formulation makes it possible to avoid using B^{-1} (not D^{-1}) altogether by starting the Gauss-Newton algorithm with $x_0 = x_b$ and recurring $B^{-1}(x_k - x_b)$ over successive major iterations from by-products of the FOM or CG algorithms. This technique however suffers from the same parallelization problems as the standard FOM and does not avoid using the (possibly approximate) Q_i^{-1} operators (see (2.3)). This is why we haven't considered this variant in detail in our parallel computing assessment.

8 Conclusion and perspectives

In this paper, we have exposed the problematic behaviour of the original saddle formulation as a general method for solving the weakly constrained 4D-Var problem. Its undesirable features are caused by the very poor correlation, for approximate solutions, between quadratic model decrease (the objective) and reduction of the residual of the associated optimality conditions (the mean). This mismatch in turn causes the values of the cost function(s) to behave chaotically and makes terminating the inner iteration too much dependent on chance, potentially resulting in divergence of the whole process. We have nevertheless proposed a strategy (and a corresponding class of algorithms) which cures the problem and for which strong global convergence results can be proved.

We have then experimented with this new class of saddle-based algorithms and compared their performance with that of methods associated with alternative variational formulations of the problem. This comparison was conducted on two different and complementary examples of data assimilation, taking into account not only performance in terms of number of iterations, but also considering two more elaborate approximations of computational costs in a two different parallel computing models. A parametric study of the sequential and parallel computing cost as a function of the costs of applying the D^{-1} operator and the accuracy obtained has been conducted, showing the relative merits of the new saddle algorithms and the more classical CG/FOM solvers for the state formulation. Both appear to have their place in the data assimilation toolbox. We have also provided a preliminary discussion of the application of both classes of algorithms in the case where the D^{-1} operator is unavailable, indicating that similar conclusions hold if it is approximated.

Several issues remain to be explored further, one of which is the use of approximate operators: we only briefly touched the question in our discussion of the use of the approximate D^{-1} , and further elaboration including preconditioning and the possible use of inexact products [17] might be of interest. The second and most important one is the translation of our conclusions, drawn in a relatively controlled context, to the more complex environments of truly parallel operational systems.

References

- [1] M. Benzi and A. Wathen. Some preconditioning techniques for saddle point problems. In *Model Order Reduction*, number 13 in Mathematics in Industry, pages 195–211, Heidelberg, Berlin, New York, 2008.

Springer Verlag.

- [2] L. Bergamaschi, J. Gondzio, M. Venturin, and G. Zilli. Inexact constraint preconditioners for linear systems arising in interior point methods. *Computational Optimization and Applications*, 36(2-3):136–147, 2007.
- [3] L. Bergamaschi, J. Gondzio, M. Venturin, and G. Zilli. Erratum to: Inexact constraint preconditioners for linear systems arising in interior point methods. *Computational Optimization and Applications*, 49(2):401–406, 2011.
- [4] F. Bouttier and P. Courtier. Data assimilation concepts and methods. Technical report, ECMWF, Reading, England, 1999. ECMWF Meteorological Training Course Lecture Series.
- [5] A. M. Clayton, A. C. Lorenc, and D. M. Barker. Operational implementation of a hybrid ensemble/4D-Var global data assimilation system at the Met Office. *Quarterly Journal of the Royal Meteorological Society*, 139:1445–1461, 2013.
- [6] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *Trust-Region Methods*. MPS-SIAM Series on Optimization. SIAM, Philadelphia, USA, 2000.
- [7] Ph. Courtier. Dual formulation of four-dimensional variational assimilation. *Quarterly Journal of the Royal Meteorological Society*, 123:2449–2461, 1997.
- [8] Ph. Courtier, J.-N. Thépaut, and A. Hollingsworth. A strategy for operational implementation of 4D-Var using an incremental approach. *Quarterly Journal of the Royal Meteorological Society*, 120:1367–1388, 1994.
- [9] A. El-Said. *Variational Data Assimilation Problem for Numerical Weather Prediction*. PhD thesis, University of Reading, Reading, UK, 2015.
- [10] A. El-Said, N. K. Nichols, and A. S. Lawless. Conditioning of the weak-constraint 4DVAR problem. Technical report, University of Reading, Reading, UK, 2017.
- [11] M. Fisher, S. Gratton, S. Gürol, Y. Trémolet, and X. Vasseur. Low rank updates in preconditioning the saddle point systems arising from data assimilation problems. *Optimization Methods and Software*, (to appear), 2017.
- [12] M. Fisher and S. Gürol. Parallelisation in the time dimension of four-dimensional variational data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 143(703):1136–1147, 2017.
- [13] M. Fisher, Y. Trémolet, H. Auvinen, D. Tan, and P. Poli. Weak-constrained and long window 4D-Var. Technical Report 655, ECMWF, 2011.
- [14] M. A. Freitag and D. L. H. Green. A low-rank approach to the solution of weak constraint variational data assimilation problems. arXiv:1702.07278v1, 2017.
- [15] S. Gratton, S. Gürol, E. Simon, and Ph. L. Toint. Issues in making the weakly-constrained 4DVar formulation computationally efficient. Oberwolfach Reports 47, 2017.
- [16] S. Gratton, S. Gürol, E. Simon, and Ph. L. Toint. Preconditioning weighted linear least-squares with an application to weakly constrained variational data assimilation. (in preparation), 2017.
- [17] S. Gratton, S. Gürol, Ph. L. Toint, J. Tshimanga, and A. Weaver. Krylov methods in the observation space for data assimilation. Oberwolfach Reports, 2012.
- [18] S. Gratton, A. Lawless, and N. K. Nichols. Approximate Gauss-Newton methods for nonlinear least-squares problems. *SIAM Journal on Optimization*, 18:106–132, 2007.
- [19] S. Gratton, Ph. L. Toint, and J. Tshimanga. Range-space variants and inexact matrix-vector products in Krylov solvers for linear systems arising from inverse problems. *SIAM Journal on Matrix Analysis*, 32(3):969–986, 2011.
- [20] S. Gratton, Ph. L. Toint, and J. Tshimanga. Conjugate-gradients versus multigrid solvers for diffusion-based correlation models in data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 139:1481–1487, 2013.
- [21] S. Gratton and J. Tshimanga. An observation-space formulation of variational assimilation using a restricted preconditioned conjugate-gradient algorithm. *Quarterly Journal of the Royal Meteorological Society*, 135:1573–1585, 2009.
- [22] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of the National Bureau of Standards*, 49:409–436, 1952.

- [23] F.-X. Le Dimet and O. Talagrand. Variational assimilation of meteorological observations: Theoretical aspect. *Tellus*, 38A:97–110, 1986.
- [24] A. C. Lorenc. Improving ensemble covariances in hybrid variational data assimilation without increasing ensemble size. *Quarterly Journal of the Royal Meteorological Society*, 143:1062–1072, 2017.
- [25] I. Mirouze and A. T. Weaver. Representation of correlation functions in variational assimilation using an implicit diffusion operator. *Quarterly Journal of the Royal Meteorological Society*, 136:1421–1443, 2010.
- [26] J. Nocedal and S. J. Wright. *Numerical Optimization*. Series in Operations Research. Springer Verlag, Heidelberg, Berlin, New York, 1999.
- [27] Y. Saad. *Iterative Methods for Sparse Linear Systems*. PWS Publishing Company, Boston, USA, 1996.
- [28] Y. Saad and M. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7:856–869, 1986.
- [29] Y. Trémolet. Accounting for an imperfect model in 4D-Var. *Quarterly Journal of the Royal Meteorological Society*, 132(621):2483–2504, 2006.
- [30] Y. Trémolet. Model error estimation in 4D-Var. *Quarterly Journal of the Royal Meteorological Society*, 133(626):1267–1280, 2006.
- [31] P. A. Vidard, A. Piacentini, and F.-X. Le Dimet. Variational data analysis with control of the forecast bias. *Tellus*, 56A:177–188, 2004.
- [32] X. Wang, T. M. Hamill, J. S. Whitaker, and C. H. Bishop. A comparison of hybrid ensemble transform kalman filter - optimum interpolation and ensemble square root filter analysis scheme. *Monthly Weather Review*, 135:1055–1076, 2007.
- [33] A. Wathen. Preconditioning. *Acta Numerica*, 24:329–376, 2015.
- [34] A. T. Weaver and Ph. Courtier. Correlation modelling on the sphere using a generalized diffusion equation. *Quarterly Journal of the Royal Meteorological Society*, 127:1815–1846, 2001.
- [35] A. T. Weaver, J. Tshimanga, and A. Piacentini. Correlation operators based on implicitly formulated diffusion equation solved with the Chebyshev iteration. *Quarterly Journal of the Royal Meteorological Society*, 142:455–471, 2016.
- [36] D. Zupanski. A general weak constraint applicable to operational 4DVAR data assimilation systems. *Monthly Weather Review*, 125, 1997.

A1. The Burgers assimilation problem

We consider the one-dimensional Burgers equation on the spatio-temporal domain $\Omega = [0, T] \times [0, 1]$ whose governing equation is

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} - \nu \frac{\partial^2 u}{\partial x^2} = g \text{ on } \Omega \quad (\text{A.1})$$

with Dirichlet boundary conditions

$$u(0, t) = u(1, t) = 0 \quad \forall t \geq 0. \quad (\text{A.2})$$

The forcing term is then given by

$$\begin{aligned} g(x, t) = & \pi k [x + k(t+1) \sin \pi(1-x)(t+1)] \cos \pi x(t+1) \sin \pi(1-x)(t+1) \\ & + \pi k [1-x - k(t+1) \sin \pi x(t+1)] \sin \pi x(t+1) \cos \pi(1-x)(t+1) \\ & + 2\nu k^2 \pi^2 (t+1)^2 [\sin \pi x(t+1) \sin \pi(1-x)(t+1) \\ & \quad + \cos \pi x(t+1) \cos \pi(1-x)(t+1)] \end{aligned} \quad (\text{A.3})$$

The discretization uses a first-order upwind scheme in time and a second-order centered scheme in space. With a space step Δx and a time step Δt , this gives

$$\frac{1}{\Delta t} (u_i^{n+1} - u_i^n) + \frac{u_i^n}{2\Delta x} (u_{i+1}^n - u_{i-1}^n) - \frac{\nu}{(\Delta x)^2} (u_{i+1}^n - 2u_i^n + u_{i-1}^n) = g(i\Delta x, n\Delta t) \quad (\text{A.4})$$

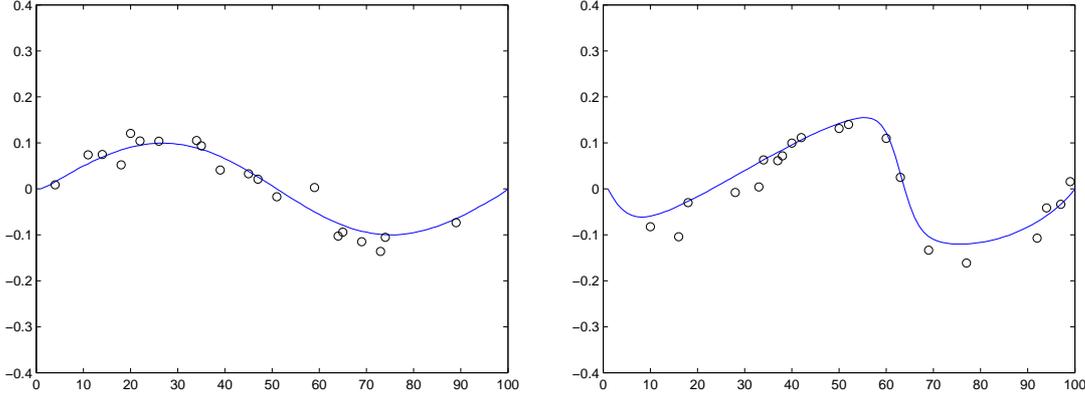


Figure A.14: Reference trajectory and observations at the end of the first (left) and last (right) subwindow

with $u_i^n = u(i\Delta x, n\Delta t)$. We choose $\Delta x = 0.01$ leading to a state vector of dimension $n = 100$, $\Delta t = 10^{-5}$ and a diffusion coefficient $\nu = 0.25$. The length of the assimilation window T is equal to 0.03. It is divided into $N_{sw} = 50$ subwindows of equal length.

The reference solution is built by running the model with the initial condition

$$u_{\text{true}}(x, 0) = k \sin(2\pi x), \quad \forall x \in [0, 1] \quad (\text{A.5})$$

with $k = 0.1$ as for the forcing function g , and by adding a Gaussian random variable at the end of each subwindow:

$$x_j^t = \mathcal{M}_j(x_{j-1}^t) + \epsilon_j^m, \quad \epsilon_j^m \sim \mathcal{N}(0, \sigma_m^2 I_n) \quad (j = 1, \dots, N_{sw})$$

with $x_j^t = [u_1^j, \dots, u_n^j]^T$ the state vector at time t_j , and \mathcal{M}_j the integration of the numerical model from time t_{j-1} to t_j . We choose $\sigma_m^2 = 1.10^{-4}T/N_{sw}$.

At the end of each subwindow (time t_j), $m_j = 20$ observations are built by randomly selecting components of the reference solution and then adding a Gaussian random variable:

$$\forall j = 1 : N_{sw} \quad y_j = \mathcal{H}_j(x_j^t) + \epsilon_j^o, \quad \epsilon_j^o \sim \mathcal{N}(0, \sigma_o^2 I_{m_j})$$

with \mathcal{H}_j the observation operator at time t_j (basically the random selection of m_j components of the reference solution). We choose $\sigma_o^2 = 10^{-3}$. This strategy results in a total of 1000 assimilated observations for the whole assimilation window. The reference solution and the observations at the end of the first and last subwindows are shown in Figure 8.

The observation error covariance matrices R_j , with $j = 1 : N_{sw}$, introduced in the definition of the 4D-Var cost function are diagonal and their diagonal entries are chosen such that they are positive, the largest one is equal to one and the condition number of R_j is equal to 10^3 .

The background solution x_b corresponds to the sum of the reference solution at initial time (A.5) and a random variable $\epsilon^b \sim \mathcal{N}(0, \sigma_b^2 I_n)$, with $\sigma_b^2 = 10^{-2}$. A model error is introduced during the numerical integration of the model by adding a Gaussian random variable at the

end of each subwindow as for the reference solution. The background error covariance matrix corresponds to the weighted sum of the squared exponential covariance and the identity matrix given by

$$B = \sigma_b^2(\alpha I_n + (1 - \alpha)\tilde{B}), \quad \text{with } \tilde{B}_{i,j} = e^{-\frac{d(i,j)^2}{L^2}} \quad (\text{A.6})$$

with $d(i, j)$ the distance between the spatial grid points i and j , $L = 0.25$ a specified length scale, and $\alpha \in [0, 1]$ the weight associated with both matrices. This last parameter is a simple way to allow variation in the condition number of the matrix B . We choose $\alpha = 0.001$ which results in a condition number of 1.10^5 . The model error covariance matrices Q_j , with $j = 1 : N_{sw}$, are built using the same strategy except that $L = 0.05$ and $\alpha = 0.01$. The condition number of these matrices is close to $1.65 \cdot 10^3$.

A2. The ECMWF QG problem

In the quasi Geostrophic (QG) problem, a reference stream function is generated from a model with layer depths of $H_1 = 6000 \text{ m}$ and $H_2 = 4000 \text{ m}$, and the time step is set to 600 s , whereas the assimilating model has layer depths of $H_1 = 5500 \text{ m}$ and $H_2 = 4500 \text{ m}$, and the time step is set to 3600 s . These differences in the layer depths and the time steps provide a source of model error.

Observations of the non-dimensional stream function, vector wind and wind speed were taken from the reference of the model at 100 points randomly distributed over both levels for each hour. Observation errors were assumed to be independent from each others and uncorrelated in time, the standard deviations were chosen to be 0.4 for the stream function observation error, 0.6 for the vector wind and 1.2 for the wind speed. The observation operator is the bi-linear interpolation of the model fields to horizontal observation locations.

The background error covariance matrix (B matrix) and the model error covariances (matrices Q_i) correspond to vertical and horizontal correlations. The vertical and horizontal structures are assumed to be separable. In the horizontal plane, covariance matrices correspond to isotropic, homogeneous correlations of stream function with Gaussian spatial structure. For the background error covariance matrix B , the standard deviation and the horizontal correlation length scale are set to 0.8 and 10^6 m respectively. For the model error covariance matrices Q_i , the standard deviation and the horizontal correlation length scale are set to 0.6 and $2 \times 10^5 \text{ m}$ respectively. The vertical correlation is assumed to be constant over the horizontal grid and the correlation coefficient value between the two layers was taken as 0.2 for B and 0.5 for Q_i .

The length of the assimilation window is set to 48 hours, divided into 48 equal sub-windows of 1 hour each.

A3. The left-preconditioned FOM algorithm and its application to the forcing formulation

We first state, as Algorithm A.1 on the following page the left-preconditioned FOM algorithm for general symmetric positive definite system $Ax = b$ with symmetric positive definite preconditioning matrix M . This algorithm uses the inner product induced by M^{-1} . In the description, we use the notation $1:k$ as a short-hand for $\{1, \dots, k\}$.

Algorithm A.1: Left-preconditioned FOM algorithm for solving $MAx = Mb$

1. Initialization. Symmetric positive definite matrices $A, M \in \mathbb{R}^{n \times n}$ are given, as well as a right-hand side $b \in \mathbb{R}^n$.

- 1.1 $w \leftarrow Mb$
- 1.2 $\beta \leftarrow \sqrt{b^T w}$
- 1.3 $U_{1:n,1} \leftarrow w/\beta$
- 1.4 $Q_{1:n,1} \leftarrow b/\beta$
- 1.5 $z_1 \leftarrow U_{1:n,1}^T b$

2. Main loop. For $i = 1, \dots, \text{maxit}$,

- 2.1 $w \leftarrow AU_{1:n,k}$
- 2.2 $v \leftarrow Mw$
- 2.3 for $j = 1, \dots, k$
 - 2.3.1 $H_{j,k} \leftarrow Q_{1:n,j}^T v$
 - 2.3.2 $v \leftarrow v - U_{1:n,j} H_{j,k}$
 - 2.3.3 $w \leftarrow w - Q_{1:n,j} H_{j,k}$
- 2.4 $H_{k+1,k} \leftarrow \sqrt{w^T v}$
- 2.5 $y \leftarrow \beta H_{1:k,1:k}^{-1} e_1$
- 2.6 $\gamma_k \leftarrow |H_{k+1,k} y_k|$
- 2.7 $q_k \leftarrow -\frac{1}{2} z^T y$
- 2.8 In view of q_k and γ , terminate with $x = Uy$?
- 2.9 $U_{1:n,k+1} \leftarrow v/H_{k+1,k}$
- 2.10 $Q_{1:n,k+1} \leftarrow w/H_{k+1,k}$
- 2.11 $z_{k+1} \leftarrow U_{1:n,k+1}^T b$

Note that Steps 1.5, 2.7 and 2.11 are only necessary if the value of the model quadratic $q(x) = \frac{1}{2}x^T Ax - b^T x$ must be tracked in the course of the inner iterations. If this is the case q_k is the value of $q(x) = q(Uy)$ at iteration k , while $\gamma_k = \|\nabla_x q(x)\|_M$ (the preconditioned norm of the system's residual). These values may then be used to decide on termination in Step 2.8 (for instance according to (4.6)). Also note that the operator M^{-1} is only used implicitly and never appears in the algorithm.

We next consider applying this algorithm (with $M = D$) to find $\delta x = L^{-1}\delta p$, where δp (approximately) solves (2.6). This gives Algorithm A.2 on the next page. Storing the matrix P in Step 2.1 of this algorithm allows avoiding the backsolve (5.1) by returning $\delta x = Py$ instead of $\delta p = Uy$.

Algorithm A.2: Specialized FOM algorithm for (2.6)-(5.1)

1. Initialization. The matrices $L, D, H, R \in \mathbb{R}^{n \times n}$ are given, as well as a right-hand side $r = D^{-1}b + L^{-T}H^T R^{-1}d \in \mathbb{R}^n$.

$$1.1 \quad w \leftarrow Dr$$

$$1.2 \quad \beta \leftarrow \sqrt{w^T r}$$

$$1.3 \quad U_{1:n,1} \leftarrow w/\beta$$

$$1.4 \quad Q_{1:n,1} \leftarrow r/\beta$$

$$1.5 \quad z_1 \leftarrow U_{1:n,1}^T r$$

2. Main loop. For $i = 1, \dots, \text{maxit}$,

$$2.1 \quad P_{1:n,k} \leftarrow L^{-1}U_{1:n,k}$$

$$2.2 \quad v \leftarrow L^{-T}H^T R^{-1}HP_{1:n,k}$$

$$2.3 \quad w \leftarrow Q_{1:n,k} + v$$

$$2.4 \quad v \leftarrow U_{1:n,k} + Dv$$

2.5 for $j = 1, \dots, k$

$$2.6.1 \quad T_{j,k} \leftarrow Q_{1:n,j}^T v$$

$$2.6.2 \quad v \leftarrow v - U_{1:n,j}T_{j,k}$$

$$2.6.3 \quad w \leftarrow w - Q_{1:n,j}T_{j,k}$$

$$2.7 \quad T_{k+1,k} \leftarrow \sqrt{w^T v}$$

$$2.8 \quad y \leftarrow \beta T_{1:k,1:k}^{-1} e_1$$

$$2.9 \quad \gamma_k \leftarrow |T_{k+1,k} y_k|$$

$$2.10 \quad q_k \leftarrow -\frac{1}{2} z^T y$$

2.11 In view of q_k and γ , terminate with $\delta x = Py$?

$$2.12 \quad U_{1:n,k+1} \leftarrow v/T_{k+1,k}$$

$$2.13 \quad Q_{1:n,k+1} \leftarrow w/T_{k+1,k}$$

$$2.14 \quad z_{k+1} \leftarrow U_{1:n,k+1}^T r$$