Toward variational data assimilation for coupled models: first experiments on a diffusion problem

Rémi Pellerej, Arthur Vidard, Florian Lemarié

To cite this version:

Rémi Pellerej, Arthur Vidard, Florian Lemarié. Toward variational data assimilation for coupled models: first experiments on a diffusion problem. CARI 2016, Oct 2016, Tunis, Tunisia. hal-01337743

HAL Id: hal-01337743
https://hal.science/hal-01337743

Submitted on 27 Jun 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CARI 2016

Data assimilation for coupled models

Toward variational data assimilation for coupled models : first experiments on a diffusion problem

Rémi Pellerej[1], Arthur Vidard[2], Florian Lemarié[3]

Inria, Univ. Grenoble-Alpes, CNRS, LJK, F-38000 Grenoble, France
[1]remi.pellerej@imag.fr
[2]arthur.vidard@imag.fr
[3]florian.lemarie@inria.fr

ABSTRACT. Nowadays, coupled models are increasingly used in a wide variety of fields including weather forecasting. We consider the problem of adapting existing variational data assimilation methods to this type of application while imposing physical constraints at the interface between the models to be coupled. We propose three data assimilation algorithms to address this problem. The proposed algorithms are distinguished by their choice of cost function and control vector as well as their need to reach convergence of the iterative coupling method (the Schwarz domain decomposition method is used here). The performance of the methods in terms of computational cost and accuracy are compared using a linear 1D diffusion problem.

RÉSUMÉ. De nos jours, les modèles couplés sont de plus en plus utilisés dans de nombreux domaines, dont les prévisions météorologiques. Nous essayons ici d'adapter les méthodes courantes d'assimilation de données variationelles à ce type d'applications tout en imposant des contraintes physiques entre les deux modèles couplés. Nous proposons trois méthodes d'assimilation de données pour ce problème. Les différents algorithmes se distinguent par le choix de leur fonction coût, de leur vecteur de contrôle et du nombre d'itérations de couplage nécessaires (nous utilisons les méthodes de Schwarz pour coupler nos modèles). Ces méthodes sont comparées dans le cadre d'un problème linéaire de diffusion 1D en analysant leur coût de calcul et la qualité de leur analyse.

KEYWORDS : Coupled data assimilation, Schwarz methods, Optimal control

MOTS-CLÉS : Assimilation de données couplée, Méthodes de Schwarz, Contrôle optimal

1. Introduction

In the context of operational meteorology and oceanography, forecast skills heavily rely on proper combination of model prediction and available observations via data assimilation techniques. Historically, numerical weather prediction is made separately for the ocean and the atmosphere in an uncoupled way. However, in recent years, fully coupled ocean-atmosphere models are increasingly used in operational centers to improve the reliability of seasonal forecasts and tropical cyclones predictions. For coupled problems, the use of separated data assimilation schemes in each medium is not satisfactory since the result of such assimilation process is generally inconsistent across the interface, thus leading to unacceptable artefacts [4]. Hence, there is a strong need for adapting existing data assimilation techniques to the coupled framework, as initiated in [5]. In this paper, three general data assimilation algorithms, based on variational data assimilation techniques [3], are presented and applied to a simple coupled problem. The dynamical equations of this problem are coupled using an iterative Schwarz domain decomposition method [1]. The aim is to properly take into account the coupling in the assimilation process in order to obtain a coupled solution close to the observations while satisfying the physical conditions across the air-sea interface. The paper is organized as follows. The model problem and coupling strategy are given in Sec. 2. In Sec. 3 we briefly recall some theoretical aspects of variational data assimilation techniques, and we introduce and discuss three algorithms to solve coupled constrained minimization problems. The performance of the proposed schemes are illustrated by numerical experiments in Sec. 4.

2. Model problem and coupling strategy

We consider a problem defined on $\Omega = \mathbb{R}$. We decompose Ω in two nonoverlapping subdomains Ω_1 and Ω_2 with an interface $\Gamma = \{z = 0\}$. A model is defined on each space-time domain $\Omega_d \times [0, T]$ ($d = 1, 2$) thanks to a differential operator $\mathcal{L}_d$ which acts on the variable u_d. The problem is to couple the two models at their interface Γ. To do so, we introduce the operators $\mathcal{F}_d$ and $\mathcal{G}_d$ which define the interface conditions. Those operators must be chosen to satisfy the required consistency on Γ. We propose to use a global-in-time Schwarz algorithm (a.k.a. Schwarz waveform relaxation, see [1] for a review) to solve the corresponding coupling problem. This method consists in solving iteratively each model on their respective space-time subdomain using the interface conditions on Γ computed during the previous iteration. For a given initial condition $u_0 \in H^1(\Omega_1 \cup \Omega_2)$ and *first-guess* $u_1^0(0, t)$, the corresponding coupling algorithm reads

$$\left\{\begin{array}{lr} \mathcal{L}_2 u_2^k = f_2 & \text{on } \Omega_2 \times T_W \\ u_2^k(z,0) = u_0(z) & z \in \Omega_2 \\ \mathcal{G}_2 u_2^k = \mathcal{G}_1 u_1^{k-1} & \text{on } \Gamma \times T_W \end{array}\right. \quad \left\{\begin{array}{lr} \mathcal{L}_1 u_1^k = f_1 & \text{on } \Omega_1 \times T_W \\ u_1^k(z,0) = u_0(z) & z \in \Omega_1 \\ \mathcal{F}_1 u_1^k = \mathcal{F}_2 u_2^k & \text{on } \Gamma \times T_W \end{array}\right. \qquad (1)$$

where k is the iteration number, $T_W = [0, T]$, and $f_d \in L^2(0, T; L^2(\Omega_d))$ is a given right-hand side. At convergence, this algorithm provides a mathematically strongly coupled solution which satisfies $\mathcal{F}_1 u_1 = \mathcal{F}_2 u_2$ and $\mathcal{G}_2 u_2 = \mathcal{G}_1 u_1$ on $\Gamma \times T_W$. The convergence speed of the method greatly depends on the choice for $\mathcal{F}_d$ and $\mathcal{G}_d$ operators, and the choice of the *first-guess*. Note that in this paper we restrict ourselves to linear differential operators for $\mathcal{L}_d$, $\mathcal{G}_d$, and $\mathcal{F}_d$, and to the multiplicative form of the Schwarz method where each model is run sequentially.

3. Data assimilation

Let us now suppose that some discrete estimates $\mathbf{y}$ of the solution to problem (1) are available over an irregular set of points in the interval $\Omega \times T_W$. In this context we are interested in using a data assimilation (DA) procedure to account for this additional source of information. For the present study we use the variational methods of DA, based on optimal control theory. Our aim is to evaluate a set of parameter $\mathbf{x}_0$, including for instance the initial condition u_0 of problem (1), through the minimization of a cost function $J(\mathbf{x}_0)$ ($\mathbf{x}_0$ is the control vector) which quantifies in some sense the misfit between the observations $\mathbf{y}$ and the model prediction. This minimization requires the gradient of $J(\mathbf{x}_0)$, which can be computed using adjoint methods [3].

3.1. Uncoupled variational data assimilation

We first briefly describe the variational DA approach in the uncoupled case to introduce the necessary notations. The control vector is restricted to subdomain Ω_d and is noted $\mathbf{x}_{0,d} = u_0|_{z \in \Omega_d}$. The optimal control problem amounts to find $\mathbf{x}^a_{0,d}$, the *analysed state*, which best fit observations $\mathbf{y}$ and a previous estimate of the initial state $\mathbf{x}^b_d$ called the *background*. Noting H the observation operator that goes from model space to the observations space and $\mathbf{x}_d = u_d$ the state vector, the cost function to minimize reads

$$J(\mathbf{x}_{0,d}) = \overbrace{\left\langle \mathbf{x}_{0,d} - \mathbf{x}^b_d, \mathbf{B}^{-1}(\mathbf{x}_{0,d} - \mathbf{x}^b_d) \right\rangle_{\Omega_d}}^{J^b(\mathbf{x}_{0,d})} + \overbrace{\int_0^T \left\langle \mathbf{y} - H\left(\mathbf{x}_d\right), \mathbf{R}^{-1}(\mathbf{y} - H\left(\mathbf{x}_d\right)) \right\rangle_{\Omega_d} \mathrm{dt}}^{J^o(\mathbf{x}_{0,d})} \tag{2}$$

where $\mathbf{R}$ is the covariance matrix associated to observation errors, $\mathbf{B}$ is the background error covariance matrix, and $\langle \cdot \rangle_\Sigma$ is the usual Euclidian inner product on a spatial domain Σ. Obviously, if the DA process is done separately on each subdomain (with prescribed boundary conditions on the interface Γ), the initial condition $u_0 = (\mathbf{x}^a_{0,1}, \mathbf{x}^a_{0,2})^T$ obtained on Ω does not satisfy the interface conditions, hence $u_0 \notin H^1(\Omega)$ and well-posedness of the coupled problem is no longer guaranteed. In practice this type of imbalance in the initial condition can severely damage the forecast skills of coupled models [4].

3.2. Toward a coupled variational data assimilation

Our objective is now to properly take into account the coupling in the assimilation process. To do this, we introduce in this section three types of variational DA algorithms whose aim is to provide a solution close to the observations while satisfying the interface conditions on Γ; or at least a weak form of it. The key properties of those algorithms are summarized in Tab. 1.

Full Iterative Method (FIM)

A first possibility is to consider a monolithic view of the problem by ignoring the presence of an interface in the assimilation process. In this case the state vector is $\mathbf{x_0} = u_0(z),\ z \in \Omega$ and for each model integration we iterate the models on Ω_1 and Ω_2 till convergence of the Schwarz algorithm. If we note k_{cvg} the number of iterations to satisfy the stopping criterion, the cost function for the FIM is

$$J(\mathbf{x}_0) = J^b(\mathbf{x}_0) + \int_0^T \left\langle \mathbf{y} - H\left(\mathbf{x}^{\mathrm{cvg}}\right), \mathbf{R}^{-1}(\mathbf{y} - H\left(\mathbf{x}^{\mathrm{cvg}}\right))\right\rangle_\Omega \mathrm{dt} \tag{3}$$

where $\mathbf{x}^{\mathrm{cvg}} = (u_1^{k_{\mathrm{cvg}}}, u_2^{k_{\mathrm{cvg}}})^T$. Since the first-guess u_1^0 in (1) is updated after each minimization iteration with the converged solution obtained during the previous model integration, the Schwarz algorithm will converge more rapidly over the minimization iteration. It can readily be seen that cost function (3) is identical to the cost function we would use for an uncoupled problem defined on Ω. The solution provided by this approach is strongly coupled. Note that the FIM requires the adjoint of the strongly coupled model (1) which can be tedious to derive. The main drawback of this method is that it possibly requires a very large number of Schwarz iterations since we systematically iterate till convergence.

Truncated Iterative Method (TIM)

In order to improve the computational cost of the FIM algorithm, we propose to truncate the Schwarz iterations in the direct and adjoint model after $k_{\max}$ iterations, with $k_{\max} < k_{\mathrm{cvg}}$. Because we do not iterate till convergence, the coupled solution strictly satisfies only one of the two interface conditions, for example we would have $\mathcal{F}_1 u_1 = \mathcal{F}_2 u_2$ and $\mathcal{G}_2 u_2 \neq \mathcal{G}_1 u_1$ if iteration $k_{\max}$ is done first on Ω_2 and then on Ω_1. As proposed by [2] in the context of river hydraulics, a convenient way to propagate the information from one subdomain to the other during the minimization iterations is to use an extended cost function which includes the misfit in the interface conditions. The idea behind this approach is to enforce a weak coupling within the minimization iterations. The control vector $\mathbf{x}_0 = (u_0(z), u_1^0(0,t))^T$ now includes the first-guess on the interface and the cost function reads

$$J(\mathbf{x}_0) = J^b(\mathbf{x}_0) + \int_0^T \left\langle \mathbf{y} - H\left(\mathbf{x}^{\mathrm{trunc}}\right), \mathbf{R}^{-1}(\mathbf{y} - H\left(\mathbf{x}^{\mathrm{trunc}}\right))\right\rangle_\Omega \mathrm{dt} + J^s \tag{4}$$

where $J^s = \alpha_{\mathcal{F}} \|\mathcal{F}_1 u_1(0,t) - \mathcal{F}_2 u_2(0,t)\|^2_{[0,T]} + \alpha_{\mathcal{G}} \|\mathcal{G}_1 u_1(0,t) - \mathcal{G}_2 u_2(0,t)\|^2_{[0,T]}$ with $\|a\|^2_\Sigma = \langle a, a \rangle_\Sigma$ and $\mathbf{x}^{\mathrm{trunc}} = (u_1^{k_{\max}}, u_2^{k_{\max}})^T$. As mentioned above, if the model is integrated first on Ω_2 and then on Ω_1 we have $\mathcal{F}_1 u_1 = \mathcal{F}_2 u_2$ and only $\alpha_{\mathcal{G}}$ is a relevant parameter in the penalization of the interface conditions in (4). Note that, unlike FIM, the first-guess is part of the control vector here, but this method still requires the adjoint of the coupling. Since the first-guess u_1^0 is updated at the end of each minimization iteration, we can expect that we will converge toward a good approximation of the strongly coupled solution.

Coupled Assimilation Method with Uncoupled models (CAMU)

The last possibility we propose to investigate is to suppress the coupling iterations and rely only on the minimization iterations to weakly couple the two models. This approach only requires the adjoint of each individual model but not the adjoint of the coupling as for the previous algorithms. The control vector is $\mathbf{x}_0 = (\mathbf{x}_{0,1}, \mathbf{x}_{0,2})^T$ with $\mathbf{x}_{0,d} = (u_0|_{z \in \Omega_d}, u_d^0(0,t))$. The corresponding cost function is

$$J(\mathbf{x}_0) = \left\{ \sum_{d=1}^{2} (J^b(\mathbf{x}_{0,d}) + J^o(\mathbf{x}_{0,d})) \right\} + J^s.$$

It is straightforward to see that this algorithm provides only a weakly coupled solution. We proceed only to one iteration of the models (which can be run in parallel) with boundary conditions on Γ provided by the term $u_d^0(0,t)$ taken from the control vector. Note that both parameters $\alpha_{\mathcal{F}}$ and $\alpha_{\mathcal{G}}$ have an impact on the solution of the minimisation. In the next section the three DA algorithms presented so far are compared in terms of computational cost and accuracy.

Algo	Control vector	# of coupling iterations	extended cost function	Adjoint of the coupling	Coupling
FIM	$(u_0(z))$	k_{cvg}	no	yes	strong
TIM	$(u_0(z), u_1^0)^T$	$k_{\max}$	yes	yes	$\sim$strong
CAMU	$(u_0(z), u_1^0, u_2^0)^T$	0	yes	no	weak

Table 1. *Overview of the properties of the coupled variational DA methods described in Sec. 3.2. Notations are consistent with those introduced in the text.*

4. Application to a 1D diffusion problem

In this section, previous algorithms are applied on a 1D diffusion problem. We, thus, consider $\mathcal{L}_d = \partial_t + \nu_d \partial_z^2$ in (1) with $\nu_1 \neq \nu_2$ the diffusion coefficients in each subdomain. The computational domain is $\Omega =]-L_1, L_2[$ with $L_1, L_2 \in \mathbb{R}^{+*}$. We choose the interface operators on Γ to obtain a Dirichlet-Neumann algorithm, i.e. $\mathcal{F}_d = \nu_d \partial_z$ and $\mathcal{G}_d = \mathrm{Id}$. The right hand side f_d is chosen so that the analytical solution $u_j^\star$ of the coupled problem on each subdomain is

$$u_j^\star(z,t) = \frac{U_0}{4} e^{-\frac{|z|}{\alpha_d}} \left\{ 3 + \cos^2\left(\frac{3\pi t}{\tau}\right) \right\} \qquad \text{on } \Omega_j \times T_W. \tag{5}$$

where $U_0 = 20\ ^\circ\mathrm{C}$ and $\tau = 22\ \mathrm{h}$. Note $\alpha_1 \nu_2 = \alpha_2 \nu_1$ is required to ensure the proper regularity of the coupled solution across the interface Γ. To satisfy this constraint we choose $\alpha_1 = 4\ \mathrm{km}$, $\alpha_2 = 0.4\ \mathrm{km}$, $\nu_1 = 1\ \mathrm{m}^2/\mathrm{s}$, $\nu_2 = 0.1\ \mathrm{m}^2/\mathrm{s}$. The model problem (1) is discretized using a backward Euler scheme in time and a second-order scheme in space. The resolution in each subdomain is $\Delta z = 20\ \mathrm{m}$ with $L_1 = L_2 = 1\ \mathrm{km}$ and the time-step is $\Delta t = 180$ s. The total simulation time is $T = 12$ h and we start the Schwarz iterations with a random first-guess.

For the assimilation experiments, we consider that the true-state $\mathbf{x}^t$ is the solution of the Schwarz algorithm (1) while the background $\mathbf{x}^b$ corresponds to the solution obtained with a biased initial condition. In both cases, the Schwarz algorithm converges in $k_{\mathrm{cvg}} = 50$ iterations with a tolerance $\epsilon = 10^{-6}$. Some observations $\mathbf{y}$ of the true-state are generated such that $\mathbf{y} = H(\mathbf{x}^t)$, with H the observation operator. The observation and background errors covariance matrices are considered diagonal such that $\mathbf{R} = 10\ \mathrm{Id}$ and $\mathbf{B} = 100\ \mathrm{Id}$. For the extended cost function we consider $\alpha_{\mathcal{F}} = \frac{\alpha_1}{\nu_1} \alpha_{\mathcal{G}}$ with different values of $\alpha_{\mathcal{G}}$. All the minimisation are done until convergence of a conjugate gradient algorithm with a stopping criterion $\| \nabla J(\mathbf{x}_0) \|_\infty < 10^{-5}$.

Single column observation experiment

For our experiments, we consider that observations are available in $\Omega \setminus \{\Gamma\}$ only at the end of the time-window (i.e. at $t = T$). In this case, the results obtained for different assimilation schemes are reported in table 2 where the performance of each scheme is presented in terms of the number of minimisation and models runs. Note that the computational cost of a given method is almost entirely dominated by the model integration. To evaluate the strength of the coupling we define an *interface imbalance indicator* which corresponds to the value of J_s at the end of the DA process, with $\alpha_{\mathcal{G}} = 0.01$ and $\alpha_{\mathcal{F}} = 40$. Values of J^s close to zero indicate that the analysed state is strongly coupled. In table 2, a root mean square error (RMSE) defined as $\sqrt{\mathbb{E}\left((\mathbf{x}^a - \mathbf{x}^t)^2\right)}$ on $\Omega \times T_W$ is also used to evaluate how much the analysed state is close to the true-state.

From table 2, we can first note that the FIM algorithm requires few minimisation iterations to obtain a low RMSE value and a strongly coupled analysed state ($J^s \sim 10^{-12}$). A drawback of this approach is a high computational cost (1169 models runs). Since in the TIM approach the coupling iterations are truncated and the first-guess u_1^0 is part of the control vector, we expect a reduced computational cost compared to FIM. It is however the case only if the J^s term is included in the cost function (i.e. $\alpha_{\mathcal{G}} \neq 0$ or

Algo	$\alpha_{\mathcal{G}}$	$\alpha_{\mathcal{F}}$	$k_{\max}$	# of minimisation iterations	# of models runs	Interface imbalance indicator	RMSE in °C
FIM	-	-	k_{cvg}	58	1169	$3.69\ 10^{-12}$	0.220
TIM	0	-	k_{cvg}	48	2016	$5.63\ 10^{-12}$	0.220
TIM	0	-	5	320	1600	$2.95\ 10^{-2}$	0.216
TIM	0	-	2	1521	3042	3.77	0.272
TIM	0.01	-	2	391	782	$9.25\ 10^{-7}$	0.217
TIM	0.01	-	1	350	350	$8.60\ 10^{-7}$	0.215
CAMU	0.01	40	0	1308	1308	$1.40\ 10^{-4}$	0.229
CAMU	0.001	4	0	268	268	$9.38\ 10^{-3}$	0.240
CAMU	0.0001	0.4	0	758	758	$3.30\ 10^{-1}$	0.327
Uncoupled	0	0	0	101	101	29.0	1.717

Table 2. *Results obtained for the three coupled variational DA methods described in Sec. 3.2 with observations available in $\Omega \setminus \{\Gamma\}$ at the end of the time-window.*

$\alpha_{\mathcal{F}} \neq 0$), otherwise the TIM requires a very large number of models runs to reach an analysed state which is of a lesser quality than with FIM. On the one hand decreasing the value of $k_{\max}$ increases the number of minimization iterations. Indeed, going to Schwarz convergence ($k_{\max} = k_{\mathrm{cvg}}$) procures the best model solution, it then needs few minimisation iterations. However, for the next iteration, the background interface is given by the control vector rather than the previous converged estimate; therefore it requires again numerous Schwarz iterations. On the other hand, by reducing the $k_{\max}$ value, the number of Schwarz iterations is reduced and the update of the first-guess more significant, but the quality of the coupling is affected and this leads to a slower minimisation convergence. Here, a good compromise is to choose $k_{\max} = 5$. When taking J^s into account in TIM (i.e. for $\alpha_{\mathcal{G}} \neq 0$), it leads to a better analysed state with significantly less models runs. Smaller values of $k_{\max}$ provide a faster convergence of the algorithm. With $k_{\max} = 1$, which corresponds to a one-way coupling, it requires only 350 models runs to provide a good approximate of the strongly coupled solution ($J^s = 8.6\ 10^{-7}$, RMSE $= 0.215\,°\mathrm{C}$). In this case, the interface condition $\mathcal{F}_1 u_1 = \mathcal{F}_2 u_2$ is imposed in a strong way in the coupling iterations while the other condition $\mathcal{G}_1 u_1 = \mathcal{G}_2 u_2$ is established in a weak way through J^s during the minimisation. For $k_{\max} > 1$ the interface condition $\mathcal{G}_1 u_1 = \mathcal{G}_2 u_2$ is also imposed in a strong way in the coupling iterations, and seems to conflict with the weak constraint from J^s. By considering uncoupled models in the CAMU algorithm, a proper choice for $\alpha_{\mathcal{G}}$ and $\alpha_{\mathcal{F}}$ to balance J^s and J^o in the cost function can lead to an efficient method (268 models runs). Too big values imply a more constrained cost function, which leads to more minimisation iterations. At the opposite, too small values do not constrain enough the interface and therefore produce poor model solutions. The analysed state shows a larger interface imbalance indicator compared to FIM and TIM, which confirms that CAMU provides a weakly coupled solution, but is significantly better than the uncoupled DA in that respect.

5. Conclusion and perspectives

We addressed in this paper the problem of variational data assimilation for coupled models. The aim of was to introduce coupled DA algorithms. In this context, a difficulty is to determine how to combine the two iterative processes at play, namely the Schwarz iterations in the coupling and the minimisation iterations in the DA problem. The proposed algorithms are distinguished by their choice of cost function and control vector as well as their need to reach convergence of the Schwarz coupling method. We showed that adding a physical constraint on the interface conditions in the cost function can have a beneficial effect on the performance of the method and allow to save coupling iterations. Moreover, an approach which only requires the adjoint of each individual model but not the adjoint of the coupling showed promising results. Since the objective is to apply such methods to ocean-atmosphere coupled models, increasingly complex models including physical parameterisations for subgrid scales will be considered in future work.

6. Acknowledgements

The work described in this article was supported by the ERA-CLIM2 project, funded by the European Union's Seventh Framework Programme under grant n°607029.

References

[1] M.J. Gander. Schwarz methods over the course of time. *Electron. Trans. Numer. Anal.*, 31:228–255, 2008. http://etna.math.kent.edu/vol.31.2008/pp228-255.dir.

[2] I.Y. Gejadze and J. Monnier. On a 2d 'zoom' for the 1d shallow water model: Coupling and data assimilation. *Comput. Methods Appl. Mech. and Engrg.*, 196(45–48):4628 – 4643, 2007.

[3] F.X. Le Dimet and O. Talagrand. Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus A*, 38A(2):97–110, 1986.

[4] D.P. Mulholland, P. Laloyaux, K. Haines, and M.A. Balmaseda. Origin and impact of initialization shocks in coupled atmosphere-ocean forecasts. *Mon. Wea. Rev.*, 143:4631–4644, 2015.

[5] P. Smith, A. Fowler, and A. Lawless. Exploring strategies for coupled 4d-var data assimilation using an idealised atmosphere-ocean model. *Tellus A*, 67(0), 2015.