

2 **Supplementary material for: Bathymetric influences on Antarctic ice-shelf**
3 **melt rates**

4

5

6

7 *Authors: D. N. Goldberg, T. A. Smith, S. H. K. Narayanan, and M. Morlighem*8 **Contents**9 **1 Modifications to the MITgcm adjoint** **1**10 **2 Resilient Adjoints** **2**11 **1 Modifications to the MITgcm adjoint**

12 The MITgcm, and in particular a configuration using the SHELFICE physics package for an Antarctic
 13 ice shelf, has been differentiated algorithmically⁴, and so no additional modifications were required for
 14 applications to ice sheet-ocean interactions. However, there are technical issues in using bathymetry
 15 as a control variable. For instance, fluid fractions at grid cell faces (see Section ?? of main text) are
 16 based on the minimum fraction of adjacent cells, leading to potential non-differentiability. We adopt
 17 the approach of⁵ of “smoothing” the min/max functions, but we note that this feature has not been
 18 used outside of bathymetric sensitivity studies.

19 Another computational challenge in treating bathymetry as a control variable lies with the implicit
 20 solve for the free surface at each time step⁶. The model solves the linear system $\mathbf{A}\eta = \mathbf{b}$ for η , where η
 21 is the free surface at the next time step, and \mathbf{b} is a field arising from the baroclinic step of the model.
 22 \mathbf{A} is a linear, self-adjoint operator on η and the propagation of sensitivity from η to b can be calculated
 23 analytically:

$$24 \quad \delta^* \mathbf{b} = \mathbf{A}^{-1} \delta^* \eta, \quad (1)$$

25 where $\delta^* \eta$ is the *adjoint sensitivity* of η and likewise for \mathbf{b} . This formulation is standard in the MITgcm

26 for adjoint based sensitivity analyses of any control variable except for fluid depth. However, the operator
 27 \mathbf{A} depends on ocean column depth, which in the present study is a control variable, and therefore the
 28 backward-propagation of sensitivities from η to \mathbf{A} must be considered as well.⁵ dealt with this issue
 29 by allowing the AD tool to differentiate the linear solver code; however, as it is an iterative solver,
 30 this approach requires storing intermediate variables at each solver iteration during every time step of
 31 the forward model, which hinders performance and does not scale well to high dimensional problems.⁵
 32 recommend, but do not implement, using the approach of², which augments Eqn. (1) with

$$33 \quad \delta^* \mathbf{A} = -\delta^* \mathbf{b} \eta^T. \quad (2)$$

34 In this work we implement this approach, obviating the need for the AD tool to differentiate the implicit
 35 solver.

36 **2 Resilient Adjoints**

37 Simulation of large models requires the use of high performance computing (HPC), generally with defined
 38 job time limits. For instance, standard batches on the ARCHER supercomputer have a walltime limit
 39 of 24 hours (there is a special queue for jobs that take up to 48 hours, but there are fewer resources
 40 available and generally longer wait times for this queue). Additionally, imposed time limits aside, longer
 41 computational jobs increase the risk of network or server errors leading to crashes. The MITgcm has
 42 a restart capability allowing to circumvent these limits: the “state” of the model is periodically saved
 43 to file, and new jobs can begin from this time stamp by reading the saved state. To restart the adjoint
 44 model, simulations must save both the forward and adjoint states – a capability referred to as *resilient*
 45 *adjoints*. A similar capability was previously implemented with TAF as *the Divided Adjoint* (DIVA).

46 Here we provide an overview of resilient adjoints, a strategy that enhances the default checkpointing
 47 scheme used by OpenAD. Checkpointing approaches store the state of the primal (forward) computation
 48 and reduce the amount of memory that is required to compute adjoints. By default, OpenAD uses
 49 binomial checkpointing for the time-stepping loop³. Consider a computation consisting of l timesteps,
 50 with c the number of checkpoints that can be stored. Figure S1 (top) illustrates binomial checkpointing
 51 for $l = 10$ and $c = 3$.

52 A two-level checkpointing approach can build upon this approach by converting the time stepping
53 loop into a loop nest containing l_2 outer iterations and l_1 inner iterations where $l = l_2 \times l_1^1$. The inner
54 loop uses binomial checkpointing as before; the outer loop uses periodic checkpointing. The left part
55 of Figure S1 (bottom) illustrates two level checkpointing for $l_2 = 5$, $l_1 = 10$ and $c_1 = 3$. The resilient
56 adjoints capability enhances two level checkpointing by storing to disk the adjoint state computed at
57 the end of each outer level iteration. To restart a computation at the granularity of an l_2 timestep then,
58 only the stored l_2 state checkpoints and the last adjoint checkpoint, if any, are required.

59 References

- 60 [1] Aupy, G., J. Herrmann, P. Hovland, and Y. Robert, Optimal multistage algorithm for adjoint
61 computation, *SIAM Journal on Scientific Computing*, 38, doi:10.1137/15M1019222, 2014.
- 62 [2] Giles, M. B., G. Corliss, C. Faure, A. Griewank, L. Hascoët, and U. Naumann, *On the Iterative*
63 *Solution of Adjoint Equations*, pp. 145–151, Springer New York, New York, NY, doi:10.1007/978-1-
64 4613-0075-5-16, 2002.
- 65 [3] Griewank, A., and A. Walther, Algorithm 799: Revolve: An implementation of checkpointing for the
66 reverse or adjoint mode of computational differentiation, *ACM Trans. Math. Softw.*, 26(1), 19–45,
67 doi:10.1145/347837.347846, 2000.
- 68 [4] Heimbach, P., and M. Losch, Adjoint sensitivities of sub-ice shelf melt rates to ocean circulation under
69 Pine Island Ice Shelf, West Antarctica, *Annals of Glaciol.*, 54, 59–69, doi:10.3189/2012/AoG60A025,
70 2012.
- 71 [5] Losch, M., and P. Heimbach, Adjoint sensitivity of an ocean general circulation model to bottom
72 topography, *Journal of Physical Oceanography*, 37(2), 377–393, doi:10.1175/JPO3017.1, 2007.
- 73 [6] Marshall, J., C. Hill, L. Perelman, and A. Adcroft, Hydrostatic, quasi-hydrostatic, and nonhy-
74 drostatic ocean modeling, *Journal of Geophysical Research: Oceans*, 102(C3), 5733–5752, doi:
75 10.1029/96JC02776, 1997.




Figure S1: Top: Binomial checkpointing schedule for $l = 10$ time steps and $c = 3$ checkpoints. Bottom Left: Two level checkpointing schedule for $l = 50$ with $(l_2 = 5)$ outer level iterations and $(l_1 = 10)$ inner level iterations. Periodic checkpointing is used in the outer level and binomial checkpointing shown by the dashed box is used at the inner level. Bottom Right: Enhanced two level checkpointing schedule with support for resilient adjoints through the writing and reading of the adjoint state at the outer level.




Figure S2: Perturbed beds (dotted contours) and corresponding perturbed barotropic stream functions (shading) in different regions of high sensitivity in Fig. 3 of the main text. (a) through (d) correspond to finite perturbations in locations (1) through (4) in Fig. 3(a) of the main text, respectively. Bathymetric perturbations plotted with $\delta R=10$ (Eqn. 3 of the main text) and 1m isolines. Isolines of unperturbed stream functions are also shown (solid where positive, dashed where negative; .05 Sv spacing).