

We would like to expressly thank Editor Ham, Prof Christianson, and anonymous review for their helpful comments and support of our work. In this document, we address the comments made by the reviewers.

Additionally, it bears mention that due in part to the suggestion that the norm under which adjoint convergence is evaluated should be modified, and due in part to the fact that the computing cluster on which some the experiments in the initial submission were run was replaced by one with differing specifications, all experiments for this paper were re-run. This then accounts for the slightly differing timing estimates in Table 5, where it can be seen that virtually all run times are shorter than their counterparts in the first submission. Also, the time step and # of time steps of the Smith Glacier experienced was slightly modified; this was not for any reason relating to code changes, but simply a wish to have a longer simulation and a shorter time step, which was took less time on the new cluster. Note the shading is different than in the initial submission because (a) the colorscale is different and (b) the duration is longer.

Reviewer 1

This paper presents an application of the adjoint fixed point iteration proposed by Christianson (1994) to the MITgcm ice flow model. The paper describes motivates the algorithm, describes the basic principles and the implementation and demonstrates its usability on two problems. Overall, the article is well written and demonstrates the benefits of Christianson's algorithm over applying AD "naively".

Thank you very much for your comments!

**** Section 5: It would be interesting to state the solver tolerances, and the number of required Picard iterations.***

The forward and reverse tolerances are both 10^{-8} – stated as well in the original manuscript (line 374). Note that, while this has not changed in the revision, a change has been made to the code in response to referee 2's review (see below) on the importance of using the conjugate norm to establish convergence of the adjoint loop; so for the adjoint this tolerance reflects reduction of the 1-norm and not the sup-norm. Convergence criteria in the forward problem are unchanged.

We now give the average forward and reverse iteration counts as well.

**** Lines 463ff: The performance of the LU-solver will degenerate with the size of the problems, as direct solvers typically scale worse than well-preconditioned linear solvers. Hence, this sentence should be phrased more carefully.***

Thank you, we make note of this now.

**** Figure 2: From the text description it is expected that for a forward tolerance of 10^{-9} , the error would be 0 (as it is assumed as the ground truth). Adding this data point results in a big jump from 10^{-9} to 10^{-8} . What is the reason for this?***

The error presented is absolute error (sup-norm), i.e. not scaled by the norm of the adjoint field. (We now qualify maximum error by maximum *pointwise* error, we feel this is sufficient to avoid confusion). As adjoint values are $O(10^5)$, errors can be considerable even if relative errors are not. Scaling the sup-norm of the error by the sup-norm of the adjoint field would yield a value of $\sim 10^{-9}$ for the leftmost point; and similarly if we considered L2-norms. Still, the linear dependence shown would remain – and it is the linear dependence that we attempt to show with this figure.

Note that the rightmost point (10^{-4}) is now removed. It was determined that in this experiment, the forward problem was not well-enough converged for the adjoint model to produce a zero-order-correct answer, and we have left this point off in our revision as we feel the four points shown are still sufficient for our argument.

*** Table 5: It would be interesting and usefull also list (and discuss) the required Picard iteration numbers that were used in the forward/adjoint solves.**

This is done now (see above).

*** The references should be checked for correct spelling (e.g. names in titles should be capitalized)**

Done, thanks for catching these...

*** Lines 469-484: This paragraph is essential and I would have liked to read it earlier. The question that this paragraph addresses is: The adjoint equations are linear, so why does one need to perform a (computationally expensive) Picard iteration at all?**

Thank you for the suggestion – this paragraph now appears in the introduction with a brief restatement in the discussion/conclusions section.

Reviewer 2 (Christianson)

General comment

I really like this paper.

It's a nice application of an established fixed-point iteration method to a new area, explained in a way that should facilitate use of the approach elsewhere. The iterator used is interesting, because it is sufficiently contractive to converge in a relatively small number of iterations (unlike those used in CFD, for example) but is not super-linear (unlike, say, Newton constructors), and so adjoint iteration is required and the choice of adjoint start-point is potentially significant.

Thank you!

There are two audiences for this paper: geoscientists looking to apply AD efficiently to their specific problems; and those already familiar with AD wanting to apply this particular approach in another application domain. It may be worthwhile to insert additional references to the standard AD literature in order to help the second group get the most out of this paper; and to help the first group learn more about AD, which in turn will give them access to techniques which originated in other application domains.

Thank you for the suggestion. The citations you suggest below have been added, and in the introduction as soon as AD is mentioned we give reference to a few general texts (and a community website)

Specific comments

line 87: the mechanical adjoint was originally proposed by J.C. Gilbert, "Automatic Differentiation and Iterative Processes", Optimization Methods and Software 1(1) (1992), 13-22, and it might be useful to cite the discussion in Gilbert's paper, as well as that in Christianson 1994. That the mechanical adjoint doesn't always actually solve the adjoint fixed-point problem accurately - or at all - was pointed out by Gilbert: the quick test whether it did is to check if the adjoints corresponding to u_0 are close to zero, where u_0 is the starting value for the forward iterations.

Thank you – a reference to the Gilbert paper has been added

line 155: This would be a good point to insert some references to standard AD literature: as well as the excellent book by Griewank and Walther already in the Reference list, there is a brisk introductory survey paper (available open-access) by BartholomewBiggs et al, "Automatic Differentiation of Algorithms", JCAM 124(1-2) (2000), 171-190.

Thank you for the suggestion – references added.

line 164: this observation remains true even when the matrix is not self-adjoint.

We still believe that a self-adjoint matrix allows for greater ease. To put in context, see Goldberg and Heimbach (2013), Parameter and state estimation with a time-dependent adjoint marine ice sheet model, *The Cryosphere*, 7(6), 1659-1678, eqs 7-8 (the paper is now referenced). Were the matrix not self-adjoint, the subroutine that implements these equations would need to transpose the matrix stored to tape.

line 193: it doesn't have to be the Euclidean norm, contraction with respect to any operator norm will do!

We have removed "Euclidean".

Formally, we realise it should be specified that this operator norm is that induced by the norm in which the fixed-point problem converges. On the other hand, since norms in a finite-dimensional vector space are equivalent, then as long as we fix the problem size the choice of norm should not be important. We hope you agree and we do not press the matter further in the text (though see below for line 250).

line 202: "uses a fixed point loop to calculate (7)" - not quite. The fixed point loop in Christianson (1994) deliberately calculates (10) rather than (7). This is for two reasons: to avoid repeatedly adding numerically (very) small terms to big ones; and in order to allow a "warm start" by using an "arbitrary" initial value of $\delta^* w$ that is close to the fixed point. It may be worth moving equation (10) earlier in the paper and pointing out explicitly that : (a) equation (7) converges with n to the value of $\delta^* \hat{a}$ that corresponds to the fixed point of equation (10) ; (b) equation (10) converges to the correct fixed point regardless of what starting point for $\delta^* w$ is actually used ; and (a) equation (7) corresponds to the result of calculating $\delta^* \hat{a}$ after iterating equation (10) precisely n times starting from $\delta^* w = \delta^* u$; Table 4 seems to assume starting at $\delta^* w = \delta^* u$, but there is no need for this restriction.

line 245: see above discussion on line 202.

Thank you for the correction. While we see the mistake made (eq 7 assumes a specific initial condition) we feel it is correct to say that your algorithm constructs the truncated infinite series within brackets in (7), albeit not as an "end-product", so the wording is amended to reflect this. We feel it is clear enough from (6) that (7) converges to $\delta^* \hat{a}$, so this is not made explicit.

The discussion around (10) is modified slightly. Attn is brought to the fact this is equivalent to step 9 of Alg 3.1 of BC94, and further we show that the result of n iterations with arbitrary $\delta^* w_0$ will converge to the prefactor in (10), showing the point you make that convergence is indep. of the initial guess.

line 250: it would be nice to know what norm is being used for the forward convergence: logically the adjoint norm should be used for the reverse convergence. (For example, the sup norm should

be used in reverse if the 1-norm is used forwards; the euclidean norm is self-adjoint, etc.) To first order, the error in the calculated value of the cost function J is the inner product of the error in u (from the forward pass) with the converged value of $\delta^ w$. This inner product is bounded by any vector norm of the error in u multiplied by the corresponding adjoint vector norm of $\delta^* w$: Christianson, "Reverse accumulation and implicit functions", *Optimization Methods and Software*, 9 (4) (1998), 307-322.*

Thank you for addressing this. In our first draft (paper and code) the sup-norm was used to evaluate convergence in all loops.

Again, due to the equivalence of norms in finite-dimensional spaces we argue this should not matter too much in practice. Still, we have updated the code to use the conjugate norm (i.e. $1/p+1/q=1$ if $1 < p < \infty$, otherwise sup-norm \leftrightarrow 1-norm) to evaluate convergence of the adjoint loop. All the results presented in the revision implement this change, with the sup-norm used for the forward iteration. This is made clear in the paper.

(We point out that Fig 2 plots the sup-norm of the error, but results look similar in the 1- and 2-norms.)

line 393: the point about indirect solvers being more efficient in large dimensions is a good one, but (as well as having the best derivative values) the final forward iteration also generally has the best pre-conditioner.

This is a good point, though this statement presupposes the type of preconditioner being used – but we do find that as the forward iteration proceeds, the number of required CG iterations for a given accuracy drops. We choose the phrasing:

"Even without the L-U optimization, however, the BC94 algorithm ensures all linear solves in the adjoint model correspond to the converged state of the fixed-point problem. In practice, this matrix is relatively well-conditioned, leading to better performance of the Conjugate Gradient solver."

Technical comments

line 54: applying the chain rule to the numerical values

done

line 76: correspond to a discretization of the correct

done

line 123: of a nonlinear operator F to obtain u :

done

line 165: - this analytic approach allows invocation

done

line 198: required to ensure convergence of Φ to a fixed point

done

line 229: undone at the end of each iteration. Once convergence is reached, storing takes place as normal in the POSTLOOP phase.

done

line 232: simplest to replace certain specific sections of OpenAD-transformed code

we chose to reword slightly to bring the OpenAD template mechanism to light

line 253: would not require changes to this subroutine [obviously it will affect what the subroutine does!]

agreed! But we went with a slightly different modification

line 285: i'm really not clear why these are uniformly set to zero

This was simply to more easily define the experiment, as we say now: "(in reality, there would be "background" melting to be perturbed, and changes to these melt rates would elicit responses of similar magnitudes, but background melting is zero for the sake of simplicity)"

line 300: presumably $m_{i,j}^{fp}$ is the value obtained using BC94?

Yes, noted now

line 330: state the range from Figure 2 explicitly here.

done

It would also be useful to have iteration counts for forward and reverse convergence (rather than having to deduce them from Table 5.)

Presumably you are referring to table 5, and not the Fig 2 experiments, we agree these values are relevant to table 5 and are now included; however, they are less relevant to figure 2 (and furthermore the experiments would need to be run again to get this information).

line 340: in reverse order relative to forward computation.

done

line 354: recover variable values from the forward computation, so that they can be used in the adjoint computation.

done

line 359: only one level of checkpoints is required.

Done, thank you, it was worded awkwardly before

line 442: closer to the forward sweep Fig 1(d) caption: useful to know how the 2nd order differencing was done.

In the caption we refer the reader to eq 17 (now 18).

Fig 2: seems to have an outlier at 10^{-4} . Any idea why?

We determined that the adjoint model essentially was not converging at this high a forward tolerance, and decided to remove this point from the analysis. We feel that the remaining datapoints still support our argument.

Table 5: what is the significance of the red and blue entries?

This is to highlight the memory difference between the mechanical and fixed-point adjoint approaches, and to highlight the performance gain of the L-U optimization. We have added a note to the caption.

An optimized treatment for algorithmic differentiation of an important glaciological fixed-point problem.

Daniel Goldberg¹, Sri Hari Krishna Narayanan², Laurent Hascoet³, and Jean Utke⁴

¹Univ. of Edinburgh, School of GeoSciences, Edinburgh, United Kingdom

²Maths. and Comp. Science Division, Argonne National Lab, Argonne, IL, USA

³INRIA, Sophia-Antipolis, France

⁴Allstate Insurance Company, Northbrook, IL, USA

Correspondence to: D N Goldberg (dan.goldberg@ed.ac.uk)

Abstract. We apply an optimized method to the adjoint generation of a time-evolving land ice model through algorithmic differentiation (AD). The optimization involves a special treatment of the fixed-point iteration required to solve the nonlinear stress balance, which differs from a straightforward application of AD software, and leads to smaller memory requirements and in some cases shorter computation times of the adjoint. The optimization is done via implementation of the algorithm of Christianson [1994] for reverse accumulation of fixed-point problems, with the AD tool OpenAD. For test problems, the optimized adjoint is shown to have far lower memory requirements, potentially enabling larger problem sizes on memory-limited machines. In the case of the land ice model, implementation of the algorithm allows further optimization by having the adjoint model solve a sequence of linear systems with identical (as opposed to varying) matrices, greatly improving performance. The methods introduced here will be of value to other efforts applying AD tools to ice models, particularly ones which solve a “hybrid” shallow ice / shallow shelf approximation to the Stokes equations.

1 Introduction

In recent decades it has become clear how little we understand about the processes governing ice sheet behavior (Vaughan and Arthern, 2007), and the complexity that is required in numerical ice sheet models in order to understand this behavior (Little et al., 2007; Lipscomb et al., 2009). The representation of poorly-understood processes in ice sheet models leads to large, poorly-constrained parameter sets, the size of which might potentially scale with the size of the numerical grid. It is vital that there be a means to relate the outputs of an ice sheet model back to these parameters, both comprehensively and efficiently. However, the simplest method of sensitivity assessment – running the model multiple times while varying each parameter in isolation – quickly becomes intractable because of the complexity of the models. Consider, for instance, a dynamic model of the Antarctic

Ice Sheet, which takes several days to run on a supercomputing cluster, and contains several hundred
25 thousand parameters pertaining to the spatially varying frictional and geothermal properties of the
bed over which it slides. Assessing the sensitivity of the model to this parameter field by the method
described above would not be feasible.

Adjoint models provide a means to assess these sensitivities in a way which is independent of the
number of parameters. The adjoint of an ice sheet model simultaneously calculates the derivatives of
30 a single model output (often called a *cost function*) with respect to all model parameters – or rather,
the *gradient* of the cost function with respect to the parameter set, or *control variables*. Note that the
latter computation more naturally lends itself to scientific inquiry, as

- this single output can be one of societal interest, for instance the contribution of an ice sheet
to sea level over a given time window; and
- 35 – an investigator is unlikely to solely be interested in just one of these (potentially) several
hundred thousand poorly-constrained parameters.

The adjoint model is essentially the linearization of the model, only the information is propagated
backward in time (or rather in reverse to computational order). As such, the original model is often
referred to as the *forward model*. Essentially, it is this backward-in-time propagation that allows for
40 simultaneous calculation of these derivatives, regardless of the dimension of the parameter set.

One of the earliest instances of the use of the adjoint of an ice sheet flow model was that of
MacAyeal (1992), in which a control method was developed to optimally fit a model to observed
velocities through adjustment of bed friction parameters. The ice flow model used in this study was
a depth-integrated approximation to the shear-thinning Stokes equations, appropriate to ice shelves
45 and weak-bedded streams (*MacAyeal*, 1989). Moreover, it was a “static” model, i.e. it consisted only
of the nonlinear stress balance governing ice velocities, and did not evolve the ice geometry or tem-
perature. The method has since been used in a number of applications (e.g., *MacAyeal et al.*, 1995;
Rommelaere, 1997; *Vieli and Payne*, 2003; *Larour et al.*, 2005; *Khazendar et al.*, 2007; *Sergienko et al.*,
2008; *Joughin et al.*, 2009). Similar methods have been applied to “higher-order” approximations
50 (*Pattyn et al.*, 2008), or to the Stokes equations themselves (e.g., *Morlighem et al.*, 2010; *Goldberg and Sergienko*,
2011; *Petra et al.*, 2012; *Perego et al.*, 2014; *Isaac et al.*, 2015).

More recently, algorithmic differentiation (AD) tools have been applied to ice sheet models for
adjoint model generation. AD tools differentiate models by ~~differentiating elemental operations and~~
applying the chain rule to their numerical values (e.g., *Forth et al.* (2012); *Naumann* (2012), also see
55 www.autodiff.org). They have been applied extensively to atmospheric and ocean codes (*Errico*,
1997; *Heimbach et al.*, 2002; *Heimbach*, 2008). The use of AD offers ease of differentiation of the
model. For instance, the majority of the adjoint models mentioned in the previous paragraph ignore
the dependence of nonlinear ice viscosity on strain rates, producing an “approximate” set of adjoint
equations which have the same form as the forward model, allowing for code reuse. At the same time,

60 this “approximate” adjoint ignores terms in the model gradient without knowing whether they are negligible. While the “full” adjoint model involves equations distinct from the forward model, the use of AD avoids having to write the code to solve them. Another advantage is modularity. Modifying, for example, the specific form of strain-rate dependence of viscosity in an ice sheet model would then require invasive changes to an analytically-derived set of adjoint equations. When generating
65 the adjoint through AD, these changes are automatic. Furthermore, AD tools are invaluable when dealing with time-dependent or multiphysics models, where model complexity makes it very difficult to generate adjoint code “by hand”. In fact, to date the only time-dependent ice sheet adjoint models have been generated through the use of AD (*Heimbach and Bugnion, 2009; McGovern et al., 2013; Goldberg and Heimbach, 2013; Larour et al., 2014*).

70 For clarity we will draw a distinction between the partial differential equations (PDEs) that comprise a mathematical model of a physical system, and the computational model that discretizes these equations. The PDEs represent an operator, the linearization of which has an adjoint (the *continuous* adjoint), which can be discretized ~~*Goldberg and Sergienko (2011)*~~*(Goldberg and Sergienko, 2011)*. Alternatively, the computational model can be differentiated directly. We focus on this *discrete* adjoint in this paper. As mentioned above, a discrete adjoint model can be thought of as the reverse order
75 computation of the original model ~~*Griewank and Walther (2008); Heimbach and Bugnion (2009)*~~*(Griewank and Walther, 2008;* but an important subtlety is that this discrete adjoint may not necessarily correspond to ~~the correct continuous~~ *a discretization of the correct* adjoint, a subtlety which bears on the accuracy of ice sheet adjoint models.

80 Most ice flow models solve a nonlinear elliptic system of partial differential equations (PDEs) for ice velocity, and these equations require an iterative fixed-point approach. (Here “most ice flow models” is taken to mean *all* ice flow models, except those which make the Shallow Ice Approximation (SIA, *Hutter (1983)*). The SIA strictly applies only to slow-moving ice frozen at its base, and not the fast-flowing ice streams at the Antarctic and Greenland margin which currently exhibit variability.)

85 We refer to this fixed-point iteration as the Forward Fixed Point Iteration (**FFPI**). Ice sheet models of this type, to which AD tools have been applied previously, simply step backward through the FFPI (*Goldberg and Heimbach, 2013; Larour et al., 2014; Martin and Monnier, 2014*). This strategy is sometimes referred to as the *mechanical adjoint* (*Griewank and Walther, 2008*). The mechanical adjoint of a fixed-point solution is in fact the iterative solution of a distinct fixed-point problem, whose
90 convergence differs from that of the forward loop (~~*Christianson, 1994*~~*(Gilbert, 1992; Christianson, 1994)*, and to which we refer as the Adjoint Fixed Point Iteration (**AFPI**). As such the mechanical adjoint could potentially perform too many iterations, thereby wasting resources; or too few iterations, resulting in decreased accuracy. In fact, in some cases the mechanical adjoint can be inaccurate regardless, as we show in Section 4.1. Additionally, the mechanical adjoint can lead to burdensome
95 memory and/or recomputation loads as discussed in Section 3. *Martin and Monnier (2014)* show

accuracy can be maintained by truncating the iteration in the mechanical adjoint, but do not provide a robust, situation-independent way of doing so.

It should be pointed out that some authors have implemented ice model adjoint generation without any iteration within the adjoint model. Depending on the approximation to the Stokes momentum balance used, the adjoint stress balance can be derived directly from the equations involved (Perego et al., 2014; Isaac et al., 2015). The result is a linear elliptic equation that can be solved without iteration, but which leads to a linear system that is far less sparse than in the forward model. Additionally, the equations must potentially be re-derived if the model physics are changed. Moreover, not all such approximations to the Stokes balance allow such an approach. “Hybrid” stress balances, which solve two-dimensional approximations to the Stokes balance and are appropriate for both fast-sliding and slow creeping flow, are increasing in popularity due to low computational cost but reasonable agreement with the First Order approximation [e.g. Goldberg (2011); Schoof and Hindmarsh (2010); Cornford et al. (2013); Arthern et al. (2015); W. Lipscomb, pers. comm]. Our ice model implements such a hybrid stress balance. Differentiating such a balance at the equation level is possible but very tedious, and leads to very complicated expressions that depend strongly on discretization (Goldberg and Sergienko, 2011), both undesirable properties.

Christianson (1994) provides a mathematical strategy for finding the adjoint of a fixed-point problem via direct solution of a related fixed-point problem. The convergence of this related problem can be directly evaluated, avoiding the problem of too many or too few iterations. A novelty of the approach is that only information from the converged state of the forward loop is used for the adjoint computation, permitting additional efficiency gains. In this paper we present an application of the AD software OpenAD (*Utke et al.*, 2008) to the MITgcm time-dependent glacial flow model (*Goldberg and Heimbach*, 2013). A different AD tool has previously been applied to this ice model, so here we focus on the implementation of the Christianson algorithm (henceforth called **BC94**) – an innovation which is observed to yield substantial improvements in performance.

2 Fixed-point problem

The forward model to which AD methods are applied is that of *Goldberg* (2011), which is a “hybrid” of two low-order approximations to the nonlinear Stokes flow equations that govern ice creep over timescales longer than a day (*Greve and Blatter*, 2009). These are the Shallow Ice Approximation, appropriate for slow-flowing ice governed by vertical shear deformation, and the Shallow Shelf Approximation (SSA; *Morland* (1987); *MacAyeal* (1989)), appropriate for fast-flowing ice governed by horizontal stretching and shear deformation. The hybrid equations have been shown appropriate in both regimes, and represent considerable computational savings over the Blatter-Pattyn equations (*Blatter*, 1995; *Pattyn*, 2003; *Greve and Blatter*, 2009), as they require the solution of a two-dimensional system of elliptic PDEs rather than a three-dimensional one.

We do not discuss the details of the model here, as they are given in detail in *Goldberg* (2011) and in *Goldberg and Heimbach* (2013). Rather, we focus on its FFPI. Conceptually, the model algorithm can be divided into two components: prognostic (time-dependent) and diagnostic (time-independent). In the MITgcm land ice model, the prognostic component comprises an update to ice vertical thickness (H) through a depth-integrated continuity equation, as well as an update of the surface elevation and, implicitly, the portion of the model domain where ice is floating in the ocean rather than in contact with its bed. The diagnostic component solves the FFPI for ice velocities based on the current thickness profile. Mathematically this step can be understood as the inversion of a nonlinear operator F to obtain \mathbf{u} :

$$140 \quad F(\mathbf{u}, \mathbf{a}) = \mathbf{f}. \quad (1)$$

Here \mathbf{u} is a vector representing horizontal depth-averaged velocities u and v . F is the discretization of a nonlinear elliptic PDE in depth-averaged velocity. \mathbf{a} represents the set of material parameters that determine the coefficients of the PDE: ice thickness (H), basal friction rheological parameters (C), and ice rheological parameters (A). \mathbf{f} is the discretization of driving stress (*Cuffey and Paterson*, 2010), or the depth-integrated hydrostatic pressure gradient (which is determined by ice thickness). In this model (and in many others) the nonlinear elliptic equation is solved by a sequence of solutions of linear elliptic operators, where the operators depend on the result of the previous linear solve:

$$\mathbf{u}_{(m+1)} = (L\{\mathbf{u}_{(m)}, \mathbf{a}\})^{-1} \mathbf{f} \equiv \Phi(\mathbf{u}_{(m)}, \hat{\mathbf{a}}), \quad (2)$$

150 where, in the definition of Φ , $\hat{\mathbf{a}}$ represents the augmentation of the set \mathbf{a} to include \mathbf{f} . L is a linear operator constructed using $\mathbf{u}_{(m)}$, the current iterate of \mathbf{u} , and the parameters $\hat{\mathbf{a}}$. Note that $\hat{\mathbf{a}}$ will differ for each time step through the dependence on ice thickness, which is updated by the prognostic component of the model. In general, the ice rheological parameters depend on ice temperature, which is advected and diffused over time. Our ice model does not have a thermomechanical component, but once developed, it will not affect the algorithm we present in this paper.

Eq. (2) is our FFPI mentioned previously. In practice the iteration is truncated when subsequent iterates agree in some predefined sense, but in theory will converge to a unique solution $\mathbf{u}_*(\hat{\mathbf{a}})$. In the process of computing the adjoint to the ice model, $\frac{\partial \mathbf{u}_*}{\partial \hat{\mathbf{a}}}$ must be found, either directly or indirectly. The focus of this paper is an efficient, scalable method of computing this object.

160 **3 Forward model and “mechanical adjoint”**

Here we give a brief overview of how the model and its mechanical adjoint are constructed. For further details the reader should consult *Goldberg and Heimbach* (2013). Table 1 contains a high-level pseudocode version of the ice model time stepping procedure. Superscripts denote time step

indices. First the velocity solve (CALC_DRIVING_STRESS and the following loop) finds ice veloc-
165 ities based on current ice thickness and material parameters; then the prognostic component updates
thickness (ADVECT_THICKNESS). The function Φ comprises the construction of the linear sys-
tem L (including the nonlinear dependence of the matrix coefficients on the previous iterate) and its
solution.

Table 2 gives an overview of our implementation of the mechanical adjoint. Here we introduce
170 some notation: for a given computational variable X , the *adjoint* to X , which formally belongs to the
dual tangent space at X , is denoted $\delta^* X$ (e.g. [Heimbach and Bugnion, 2009](#)) (e.g. [Heimbach and Bugnion \(2009\)](#);
[see also Bartholomew-Biggs et al. \(2000\); Griewank and Walther \(2008\)](#)). The algorithm evolves the
adjoint variables (e.g., $\delta^* H$) backward in time. These adjoint variables carry with them the sensitiv-
ities of the model output to the corresponding forward variables, and the sensitivities are eventually
175 propagated back to the input parameters. Note that the adjoints of the individual (pseudo-) subrou-
tines are given and correspond to the (pseudo-) subroutines of the forward model, mirroring the way
the adjoint is actually constructed. Just like the forward model, the adjoint contains an inner loop
– this is a specific implementation of the AFPI, which will be discussed in further detail below.
As the computation of Φ involves the solution of a linear system of equations, the adjoint of Φ in-
180 volves the solution of the adjoint of that system. Since the matrix $L\{\mathbf{u}_{(m)}, \mathbf{a}\}$ is self-adjoint, it is
easier to calculate this result analytically than for an AD tool to differentiate the linear solver code
~~–allowing~~ (e.g. [Goldberg and Heimbach, 2013](#)). This allows for invocation of external “black box”
libraries that cannot be differentiated by the tool. This ~~strategy is used by other applications of AD~~
~~to analytical approach allows invocation of AD for~~ ice models (e.g., [Martin and Monnier, 2014](#)).

185 An important point to be made is that the inner loop in Table 2 is executed as many times as
the corresponding inner loop in the forward model ($lastm^{[n]}$), without any checks of convergence.
This could lead to under- or over- convergence, as stated previously. Another important aspect is
that at each reverse time step, and, importantly, at each iteration of the FFPI, the state of the forward
model is required. In particular, every matrix $L\{\mathbf{u}_{(m)}, \mathbf{a}\}$ must be stored (or recomputed), along with
190 other intermediate variables within the fixed-point loop. The storage and recovery steps are shown
explicitly in tables 1 and 2 – and can lead to burdensome memory loads depending on the number
of fixed-point iterations taken at each time step.

The mechanical adjoint of our model was first generated using TAF (Transformation of Algo-
rithms in Fortran; [Giering et al. \(2005\)](#)), but has subsequently been generated via OpenAD with
195 little further difficulty.

4 Fixed point treatment

[Christianson \(1994\)](#) presents an algorithm (BC94) for calculating the adjoint of a fixed-point prob-
lem that addresses the shortcomings given above, namely the dependence of the termination of the

adjoint loop on that of the forward loop, and the requirement to store variables at each iteration of
 200 the adjoint loop. Additionally it provides the opportunity for further optimization when applied to a
 higher-order ice sheet model, as discussed below.

4.1 Mathematical basis

For a rigorous mathematical analysis of BC94 the user is asked to consult the original paper. Here
 we give a brief overview of its mathematical basis. In terms of Eq. (2), consider the converged state
 205 of the fixed point problem:

$$\mathbf{u}_* = \Phi(\mathbf{u}_*, \hat{\mathbf{a}}). \quad (3)$$

Consider a total differential of this equation:

$$\delta \mathbf{u}_* = \frac{\partial \Phi}{\partial \mathbf{u}}(\mathbf{u}_*, \hat{\mathbf{a}}) \delta \mathbf{u}_* + \frac{\partial \Phi}{\partial \hat{\mathbf{a}}}(\mathbf{u}_*, \hat{\mathbf{a}}) \delta \hat{\mathbf{a}}. \quad (4)$$

Rearranging gives

$$210 \quad \delta \mathbf{u}_* = \left[I - \frac{\partial \Phi}{\partial \mathbf{u}} \right]^{-1} \frac{\partial \Phi}{\partial \hat{\mathbf{a}}} \delta \hat{\mathbf{a}}. \quad (5)$$

If the ~~Euclidean~~ operator norm of the square matrix $\partial \Phi / \partial \mathbf{u}$ is less than unity then the above is
 equivalent to

$$\delta \mathbf{u}_* = \left(I + \partial \Phi / \partial \mathbf{u} + (\partial \Phi / \partial \mathbf{u})^2 + (\partial \Phi / \partial \mathbf{u})^3 + \dots \right) \frac{\partial \Phi}{\partial \hat{\mathbf{a}}} \delta \hat{\mathbf{a}}. \quad (6)$$

Note that in the above series, $\partial \Phi / \partial \mathbf{u}$ is always evaluated at the converged solution \mathbf{u}_* . The above
 215 condition on the norm of $\partial \Phi / \partial \mathbf{u}$ will not hold in general – but since this is one of the conditions
 required ~~for convergence to ensure convergence of Φ~~ to a fixed point, we can expect that it will be
 satisfied at \mathbf{u}_* .

From eq. (6) we obtain the desired *adjoint* operator, approximated by a truncated series of length
 n :

$$220 \quad \delta^* \hat{\mathbf{a}} = \left(\frac{\partial \Phi}{\partial \hat{\mathbf{a}}} \right)^T \left[I + \left(\frac{\partial \Phi}{\partial \mathbf{u}} \right)^T + \left(\left(\frac{\partial \Phi}{\partial \mathbf{u}} \right)^T \right)^2 + \dots + \left(\left(\frac{\partial \Phi}{\partial \mathbf{u}} \right)^T \right)^n \right] \delta^* \mathbf{u}_*. \quad (7)$$

The algorithm of Christianson (1994) ~~uses essentially constructs the operator within brackets~~
~~in (7) via~~ a fixed-point loop ~~in order to calculate~~, the convergence criterion of which determines
 the truncation length n . This loop represents an implementation of the AFPI, distinct from the one
 implemented by the mechanical adjoint. In order to make this distinction explicit, the operator in eq.
 225 (7) can be written

$$\sum_{i=0}^n \left(\frac{\partial \Phi}{\partial \hat{\mathbf{a}}} \right)^T \prod_{k=n+1-i}^n \left(\frac{\partial \Phi}{\partial \mathbf{u}} \right)^T, \quad (8)$$

where it is understood that in the $i = 0$ term the product sequence evaluates to the identity. It is straightforward to check that the mechanical adjoint (cf Table 2) effectively computes the operator

$$\sum_{i=0}^n \left(\frac{\partial \Phi_{(n-i)}}{\partial \hat{\mathbf{a}}} \right)^T \prod_{k=n+1-i}^n \left(\frac{\partial \Phi_{(k)}}{\partial \mathbf{u}} \right)^T, \quad (9)$$

230 where $\partial \Phi_{(k)}/\partial \mathbf{u}$ and similar terms indicate that the gradient is calculated using the variables that have been stored at forward iteration k , rather than at the converged solution. It is apparent that this expression can differ from eq. (7) if some iterates are far from the fixed point, or if the gradient of Φ is sensitive to \mathbf{u} . In fact, it has been observed in certain cases that a poor choice of initial iterate can lead to inaccurate adjoint calculation. Furthermore, in the mechanical adjoint, the truncation length
 235 depends on the number of forward iterations, which may not be related to the convergence of this series. A scheme which truncates this series based on the size of the truncated terms will be more robust.

4.2 Implementation in OpenAD

Tables 3 and 4 give an overview of our implementation of BC94 in the MITgcm ice model using
 240 OpenAD. High-level changes to the code were necessary, but the subroutines that comprise the action of the operator Φ were left unchanged. As shown in 3, rather than calling Φ directly, the loop implementing the FFPI calls a subroutine called PHISTAGE with an argument `phase` which has values `PRELOOP`, `INLOOP`, or `POSTLOOP`. Just before the fixed-point loop PHISTAGE is called with `PRELOOP`, which does nothing (that is, nothing in forward mode). Within the loop, PHISTAGE
 245 is called with argument `INLOOP`, which essentially has the same effect as the call to Φ in the original ice model time stepping algorithm. After the loop is converged, PHISTAGE is called with argument `POSTLOOP`, which calls Φ one more time (which, if the iteration is converged, should have negligible effect). Of key importance is that any storing of variables that takes place within the call to Φ in the `INLOOP` phase is *undone* at the end of each iteration, ~~unless~~. Once convergence is reached:
 250 ~~In other words, exactly one “iteration’s worth” of storage occurs during the time step, storing takes place as normal in the `POSTLOOP` phase.~~

The reason for the addition of this layer PHISTAGE is rooted in the nature of OpenAD source transformation. To implement BC94 using this tool, it was found to be simplest to ~~replace OpenAD-transformed code with handwritten code, which can be done at the subroutine level using *templates* files~~
 255 apply the *template* mechanism provided by OpenAD, that lets the end-user provide a customized differentiation of specific sections of the code by means of a *template*, hand-written once and for all. Such a template was written for PHISTAGE in order to implement the pseudocode in tables 3 and 4. The subroutine thus serves as a “layer” which does not affect the diagnostic ice physics represented by the function Φ or the prognostic physics implemented outside of the FFPI loop. Thus the modularity offered by
 260 the AD approach is not lost.

Table 4 shows how the adjoint model is constructed, making use of the OpenAD-generated adjoint code for Φ . In adjoint mode, the calls to PHISTAGE happen in reverse order. The variable \mathbf{w} is a placeholder with no real role in the forward computation, but the adjoint of the call to PHISTAGE in the POSTLOOP phase assigns to $\delta^*\mathbf{w}$ the adjoint ~~values~~ of velocity resulting from
 265 AD_ADVECT_THICKNESS, in other words $\delta^*\mathbf{u}_*$. In the ~~INLOOP-INLOOP~~ phase $\delta^*\mathbf{w}$ is updated according to the equation \div

$$\delta^*\mathbf{w}_{(m+1)} = \delta^*\mathbf{w}_{(m)} \left(\frac{\partial\Phi}{\partial\mathbf{u}} \right)^T + \delta^*\mathbf{u} \quad (10)$$

where m indicates the AFPI iteration step. (In the table ~~the~~ the subscript indices are left off for clarity.) ~~This loop iteratively constructs the truncated infinite series.~~ This assignment is equivalent to step 9 of Algorithm 9.1 of Christianson (1994). Given that $\delta^*\mathbf{w}$ is initialized to $\delta^*\mathbf{u}_*$, it can be seen that $\delta^*\mathbf{w}_{(n)}$ is equivalent to the argument of $\left(\frac{\partial\Phi}{\partial\mathbf{a}}\right)^T$ in eq. 7 (or rather, its action on $\delta^*\mathbf{u}_*$). Christianson (1994) observes furthermore that if the convergence criteria are met, any other initial $\delta^*\mathbf{w}_{(0)}$ will converge to $\delta^*\hat{\mathbf{a}}$ for a sufficient n . This property can be used to implement a warm start of the algorithm when a good initial guess of $\delta^*\mathbf{w}$ is available. We did not test this idea for
 275 our present experiments. Finally, the adjoint-mode call to PHISTAGE with PRELOOP represents the operation of $\left(\frac{\partial\Phi}{\partial\hat{\mathbf{a}}}\right)^T$ on the result.

The introduction of the variable \mathbf{w} represents the bulk of the modifications that were necessary to implement the algorithm using OpenAD. The only additional modification is a handwritten evaluation of convergence of $\delta^*\mathbf{w}$: we terminate when the relative reduction in the ~~sup-norm-norm~~
 280 the change in $\delta^*\mathbf{w}$ is below a fixed tolerance. The norm in which convergence is evaluated is the conjugate norm to that used in the forward iteration: that is, if forward convergence is evaluated in the L_p norm, then adjoint convergence is evaluated in the L_q norm, where $\frac{1}{p} + \frac{1}{q} = 1$ (and the L_1 norm is conjugate to the sup -norm). Though all norms are equivalent in a finite-dimensional vector space, this feature is added for completeness, motivated by the fact that the error in the derivative is
 285 bounded by the inner product of the error in the forward iteration and the error in the reverse iteration (Christianson, 1998). In the results presented in this paper, convergence of the forward iteration is evaluated in the sup -norm (thus adjoint convergence is evaluated in the L_1 norm).

We emphasize that all of these modifications are at the level of the “wrapper” PHISTAGE, which does not contain any representation of model physics (and hence changes to the model physics would
 290 not ~~impact~~ require changes to this subroutine nor ~~its handwritten adjoint code~~ to its adjoint template).

4.3 Optimization of linear solve

As mentioned previously, evaluating Φ involves the solution of a large (self-adjoint) linear system, and thus the adjoint of Φ involves the solution of a linear system with the same matrix (assuming the same values of \mathbf{u} and $\hat{\mathbf{a}}$). In the mechanical adjoint model, within a given time step, this matrix
 295 differs with each iteration of the adjoint loop; however, in BC94, only the right hand side differs.

This invariance suggests the use of a linear solver whose cost can be amortized over a number of solves, such as an L-U decomposition or an algebraic multigrid preconditioner, the internal data structures of which only need be constructed once. In this study, we consider only an L-U solver.

5 Test Experiment

300 A simple experimental setup was developed to test the accuracy, performance, and convergence properties of the implementation of BC94. The setup consists of an advancing ice stream and shelf in a rectangular domain $(x, y) \in [0, 80\text{km}] \times [0, 40\text{km}]$. We prescribe an idealized bedrock topography R and initial thickness h_0 . R does not vary in the along-flow (x -) direction and forms a channel through which the ice flows, prescribed by

$$305 \quad R(x, y) = -600 - 300 \times \sin\left(\frac{\pi y}{40\text{km}}\right), \quad (11)$$

while initial thickness is given by

$$h_0(x, y) = \begin{cases} 300 \text{ m} + \min\left(1, \left(\frac{x-50 \text{ km}}{62 \text{ km}}\right)^2\right) \times 1000 \text{ m} & 0 \leq x < 50 \text{ km} \\ 300 \text{ m} & 50 \text{ km} \leq x \leq 70 \text{ km}. \end{cases} \quad (12)$$

Where $x > 70 \text{ km}$, there is open ocean (until the ice shelf front advances past this point). Where ice is grounded, a linear sliding governs basal stress:

$$310 \quad \tau_b = -Cu \quad (13)$$

where $C = 25 \text{ Pa (a}^{-1}\text{m)}$. The Glen's Law coefficient (which controls the ice stiffness) is given by $8.5 \times 10^{-18} \text{ Pa}^{-3} \text{ a}^{-1}$, corresponding to ice with a uniform temperature of $\sim -34^\circ\text{C}$. At the upstream boundary, ice flows into the domain at $x = 0$ at a constant volume flux per meter width of $1.5 \times 10^6 \text{ m}^2/\text{a}$. At $y = 0$ and $y = 40 \text{ km}$ no-flow conditions are applied. Velocity, thickness and grounding line are plotted in Fig. 1(a). Further details of the equations are given in *Goldberg and Heimbach (2013)*.

In the experiment, a cost function J is defined by running the model forward in time for 8 years, and evaluating the summed square velocity at the end of the run. That is,

$$J = \sum_{i,j} u(i, j)^2 + v(i, j)^2 \quad (14)$$

where i and j indicate cell indices in the x - and y -directions, respectively, and u and v are cell-centered surface velocities. Unless specified otherwise time step is 0.2 years and grid resolution is 2000 m, so $1 \leq i \leq 40$ and $1 \leq j \leq 20$. The control variable consists of basal melt rate m , defined for each cell and considered constant over a cell and in time (and nonzero only where ice is floating), and set uniformly to zero in the forward run, even under floating ice [\(in reality, there would be "background" melting to be perturbed, and changes to these melt rates would elicit responses of similar magnitudes, but background melting is zero for the sake of simplicity\)](#). Fig. 1(b) plots the

adjoint sensitivities of m , or alternatively $\partial J/\partial m_{ij}$, where m_{ij} is melt rate in cell (i, j) . The field shows broad-scale patterns that are physically sensible: in the margins of the ice shelf toward its front, thinning through basal melting will weaken the restrictive force on the shelf arising from tangential stresses at the no-slip boundaries. The driving force for flow is proportional to ice shelf thickness, and so in the center of the shelf thinning leads to deceleration. Meanwhile, ice shelf velocities are very insensitive to melting at the center of the ice shelf front.

We find that the results of the mechanical adjoint and of the adjoint model implementing BC94 (which we henceforth refer to as the “fixed-point adjoint”) are almost identical, with a relative difference no larger than 10^{-6} over the domain (not shown). However, the adjoint sensitivities should also be compared against a “direct” computation of the gradient, i.e. one which does not involve the adjoint model. In this case $\partial J/\partial m_{ij}$ is approximated through finite differencing, by perturbing m_{ij} by a finite amount and running the forward model again. This calculation is carried out for each cell (i, j) . Fig. 1(c) plots $disc_{fd}$, given by

$$disc_{fd} = \frac{\delta^* m_{ij}^{fp} - \delta^* m_{ij}^{cd}}{\delta^* m_{ij}^{cd}} \frac{\delta^* m_{ij}^{fp} - \delta^* m_{ij}^{cd}}{\delta^* m_{ij}^{fp}}, \quad (15)$$

where $\delta^* m_{ij}^{fp}$ is obtained through the BC94 algorithm, while $\delta^* m_{ij}^{cd}$ is a centered-difference approximation:

$$\delta^* m_{ij}^{cd} = \frac{1}{2\epsilon} (J(m_{ij} + \epsilon) - J(m_{ij} - \epsilon)), \quad (16)$$

and $J(m_{ij} + \epsilon)$ indicates that the meltrate at cell (i, j) only is perturbed by ϵ . ϵ is set to 0.01 m/a uniformly.

$disc_{fd}$ is seen to become quite large, on the order of $\sim 1\%$ in some parts of the domain, warranting further examination. An implicit assumption in the discrepancy measure $disc_{fd}$ is that the finite difference approximation has negligible error, which may not be the case. We can estimate where this finite-difference error will be large: from a Taylor series expansion, and ignoring round-off error (which we do not attempt to estimate), the error in approximating the adjoint sensitivity of m_{ij} by finite difference is roughly proportional to the second derivative $\partial^2 J/\partial(m_{ij})^2$. As a proxy for this quantity we plot in Fig. 1(d) the 2nd-order difference of J :

$$\Delta^2 J_{ij} = J(m_{ij} + \epsilon) + J(m_{ij} - \epsilon) - 2J \quad (17)$$

~~Aside from the left hand boundary, this~~ This measure appears to correlate well with $disc_{fd}$. ~~Thus we can at least partly attribute the pattern of discrepancy in Fig. 1(c) to errors in the finite difference approximation, aside from the central part of the ice shelf front. Here, the large relative errors are likely due to the small magnitude of the adjoint sensitivities.~~ We emphasize that (17) is not an accurate measure of the second derivative – which is obviously not achievable through finite differencing if first-order derivatives are inaccurate – but is simply meant to give an indication of its magnitude.

5.1 Truncation errors

360 The analysis of *Christianson* (1994) suggests the error of the calculated adjoint depends linearly on both the *reverse truncation error* and the *forward truncation error*. The reverse truncation error is the difference between the final and penultimate iterates in the adjoint loop, i.e. the error associated with terminating the loop after a finite number of iterations. That is, referring to Table 4, if m iterations are carried out, the reverse truncation error is equal to

$$365 \quad \alpha \|w_m - w_{m-1}\|, \tag{18}$$

where α is related to the gradient of Φ at the fixed point. The norm here is the *sup*-norm, because this is the norm on which our convergence criterion is based.

While a tight bound for α will vary with each time step, it can be expected that the reverse truncation error will vary linearly with the convergence tolerance and we do not address it further. However, we investigate the dependence on forward truncation error as follows. A sequence of adjoint model runs is carried out with increasingly smaller tolerances (from 10^{-5} to 10^{-8}) for the forward fixed-point iteration loop. The tolerance of the reverse loop is kept at a small value (10^{-8}). The adjoint sensitivities corresponding to the smallest forward tolerance (10^{-9}) are assumed to be “truth”; error is estimated by comparison with these values. Fig. 2 plots the maximum pointwise error in the adjoint calculation over the domain against the forward tolerance, which is a good measure of the forward truncation error. Within a range of forward truncation error the dependence is nearly linear, although this dependence appears to become weaker as forward truncation error becomes smaller.

5.2 Performance

Here we evaluate the relative performance of the mechanical and fixed-point adjoint models in terms of both timing and memory use. The results are presented in Table ??, but we must first briefly discuss how the OpenAD-generated adjoint computes sensitivities for a time-dependent model. As mentioned in the introduction, adjoint computation takes place in reverse order relative to forward computation. This presents an issue, because at each time step in this reverse computational mode, the adjoint model requires knowledge of the full model state at the corresponding forward model time step. In general, keeping the entire trajectory (including intermediate variables) of a time-dependent model run in memory is not tractable. Therefore efficient adjoint computation is a balance between recomputation (beginning from intermediate points in the run known as “checkpoints”), storage of checkpoint information on disk, and keeping variables in memory (in data structures called “tapes”). The “store” and “restore” commands in tables 1-4 refer to tape manipulation. For further information on adjoint computation see *Griewank and Walther* (2000) and *Griewank and Walther* (2008).

390 In our implementation this amounts to an initial forward run with no taping (aside from the final time step), but writing of checkpoints to disk. This initial run is referred to below as the “forward sweep”. Afterwards the “reverse sweep” begins, beginning with the final time step. The reverse

sweep consists of an ~~initial-initial~~ adjoint computation for the final timestep. As reverse computation
395 proceeds, the model is restarted from checkpoints to recover ~~variables used in~~ variable values from
the forward computation, so that they can be used in the adjoint computation. The details of this pro-
cess are important because they determine how many extra forward time steps (without taping) must
be taken. These plain time steps set up the computation of a subsequent time step in “tape mode”,
i.e. they write intermediate variables to tape during computation. This is followed immediately by a
400 time step computation in “adjoint mode”. In the model runs we consider, ~~no extra plain checkpoints~~
~~are~~ only one level of checkpoints is required. A run of 40 time steps, then, will consist of nearly
40 time steps in “plain mode” (no taping, but with checkpoint writing), 40 time steps in tape mode,
and 40 time steps in adjoint mode. Even if adjoint time steps and writing to disk and to tape are
negligible, such a run will still take about twice as long as the forward model.

405 In Table ?? we compare run times for the forward and reverse sweeps for the mechanical and
fixed-point adjoints of our test problem, at multiple grid resolutions. We also give run times for
the “untouched”, or “plain” model, i.e. code which has not been transformed by OpenAD. The
difference between this time and the forward sweep represents writing checkpoints to disk, taping in
the final time step, and any other extra steps or changes (e.g. modified variable types) caused by the
410 transformation.

We ~~also show the~~ additionally give the average number of iterations per time step. In the “plain”
runs this number is the average number of Picard iterations per time step in the forward model, which
does not change for the adjoint runs. For adjoint runs, the average iteration count for the adjoint loop,
i.e. the loop represented in Table 4, is given. We do not give a value for the mechanical adjoint, as the
415 number of adjoint iterations is set by the number of forward iterations. Note that while the average
forward iteration count grows significantly between the 80x40 and 160x80 runs, the same is not true
for the adjoint runs.

Also reported is the maximum length of the double tape in memory. There are different tapes
for different variable types: integer, double, logical and character. The double tape is observed to
420 require the most memory in our tests. However, due to storage of loop indices, the integer tape is
nonnegligible, requiring between 20% (in the largest test) ~~to~~ and 50% (in the smallest test) of the
memory required by the double tape. The numbers reported represent an upper bound, as our system
of reporting tape lengths has a granularity of $16 \times (1024)^2$ elements. Thus all BC94 runs have double
memory loads between 8 MB and 136 MB, but more exact figures cannot be given. Memory load
425 is per processor – which is why, in the mechanical adjoint runs, the double tape length increases
four-fold from the 40x20 run to the 80x40 run, but not from the 80x40 run to the 160x80 run. In this
case, domain decomposition decreases the per-processor tape length; but on the other hand, the tape
grows with the *maximum* forward iteration count (rather than the average), which is about twice as
large for the 160x80 run as the others.

430 In all cases, the forward and adjoint fixed-point tolerance thresholds are set to 10^{-8} . As resolution increases, stability considerations require smaller time steps, so the number of time steps doubles when cell dimensions are halved. The simulations are run on Intel Xeon 2.672.40GHz cpus and the number of cores used is displayed. Unless otherwise specified, the Conjugate Gradient solver from the PETSc library (<http://www.mcs.anl.gov/petsc>) with IL-U preconditioner is used to invert
435 all matrices. —

The results show that without further optimization, the BC94 algorithm does not offer large timing performance gain over the mechanical adjoint. The forward sweep is slightly shorter, but the reverse sweep is roughly the same. However, the memory load is far less, only going up to (at most) 136 MB in the high resolution run where the mechanical adjoint uses 2.762.95 GB. This provides a possible
440 explanation for the forward sweep of the mechanical adjoint being slower: it is overhead associated with the additional memory allocation. As even at the highest resolution this is still a modestly-sized problem, it is likely that certain setups of the model on certain machines would quickly reach memory limits and either crash or beginning swapping memory, significantly affecting performance.

Substantial timing performance gains are not seen until the L-U optimization described in Section 4.3. As discussed, this optimization is made possible by the BC94 algorithm. At the highest
445 resolution tested, the reverse sweep takes 4031% less time, and overall the model run is 3022% shorter. The performance gain is due to the fact that in a time step, the direct L-U decomposition is only done once, and subsequent linear solves are by forward- and back-substitution, which are far less expensive operations. As indirect solvers such as Conjugate Gradients are typically faster
450 than direct matrix solvers, it is unclear what relative performance gain would be at even higher resolutions; but in the three resolutions tested, ~~relative performance improves with resolution~~ as well as in the realistic experiment in Section 6, a noticeable improvement was observed. Even without the L-U optimization, however, the BC94 algorithm ensures all linear solves in the adjoint model correspond to the converged state of the fixed-point problem. In practice, this matrix is relatively
455 well-conditioned, leading to better performance of the Conjugate Gradient solver.

We mention that the BC94 algorithm has recently been implemented in the AD tool Tapenade, through a different user interface that relies on directives inserted in the code rather than on the OpenAD templating mechanism. It has not been tested on an ice flow model but on two other CFD codes, without our linear solver optimisation part. Their performance results are in line with ours,
460 with a minor run-time benefit but a major reduction of memory consumption (*Taftaf et al.*, 2015).

6 Realistic Experiment

In addition to idealized experiments, the fixed-point adjoint has been tested in a more realistic setting. Smith Glacier in West Antarctica is a fast-flowing ice stream that terminates in a floating ice shelf. In recent years, high thinning rates of Smith have been observed (*Shepherd et al.*, 2002;

465 *McMillan et al.*, 2014), and this is thought to be related to, or even caused by, thinning of the adjacent ice shelves by submarine melting (*Shepherd et al.*, 2004). Here we examine this mechanism using the fixed-point adjoint. To initialize the time-dependent model, we choose a domain and a representation of the bedrock elevation and ice thickness in the region from BEDMAP2 (*Fretwell et al.*, 2013) and constrain the hidden parameters of the model (basal frictional coefficient field and depth-averaged ice temperature) according to observed velocity using methods that have become standard in glaciological data assimilation (e.g., *Joughin et al.*, 2009; *Favier et al.*, 2014). The observed velocities come from a dataset of satellite-derived velocity over all of Antarctica (*Rignot et al.*, 2011).

475 Using the bed and thickness data, and the inferred sliding and temperature fields, the model is stepped forward for ~~5 years with 0.2~~ 10 years with 0.125 year time steps (80 time steps). The simulation is run on ~~24~~ 60 cpus. As with our test experiment, submarine melt rate is used as the control variable. The cost function, rather than being a measure of velocity, is the loss of *Volume Above Floatation* (VAF) in the domain at the end of the ~~5~~ 10 years. VAF is essentially the volume of ice that could contribute to sea level change, and is often used to assess the effects of ice shelf thinning on grounded ice *Dupont and Alley* (2005). It is given by

$$480 \quad \text{VAF} = \sum_i \text{HAF}(i) \Delta x \Delta y, \quad (19)$$

$$\text{HAF}(i) = \left(h(i) + \frac{\rho_w}{\rho} R(i) \right)^+, \quad (20)$$

where i is cell index, h is thickness, ρ and ρ_w are respectively ice and ocean density, R is bedrock elevation, and the “+” superscript indicates the positive part of the number. We use $\rho = 918 \text{ kg/m}^3$ and $\rho_w = 1028 \text{ kg/m}^3$. A key aspect is that any floating ice does not contribute to VAF.

485 The results are shown for the ice shelves connecting to Smith Glacier in Fig. 3, overlain on grounded ice velocities (adjoint melt rate sensitivities are zero where ice is grounded). It is interesting to note where the sensitivities are largest, along the margins of the ice shelves and also along the boundary between the two main sections of the ice shelf. The mechanism is similar to that of our test experiment: the margins are where shear stress is exerted, and thinning here will lessen the backforce on grounded ice. The boundary between the two sections of the ice shelf likely plays a similar role in the ice shelf force balance, as velocity shear is large in this area (not shown).

Regarding accuracy, the finite-difference approximation to the gradient cannot be found for every ice shelf cell. However, we compared the adjoint sensitivity to the finite difference approximation at 4 arbitrary locations, and relative discrepancy was on the order of 10^{-5} . In terms of performance, this is a much larger setting than even the highest resolution examined in the test problem. The 500 m cell size leads to approximately 200,000 ice-covered cells in the domain (which means the matrices involved, which incorporate both x - and y - velocities, have 400,000 rows and columns). The forward sweep has a run time of ~~1700~~ 1150 seconds and the reverse sweep ~~2340 seconds~~ 1778 seconds. Without using the L-U optimisation, the reverse sweep is 2765 seconds. (Multiple runs on

500 the same cluster give similar timing results.) ~~Only the fixed-point adjoint with an L-U solver for the adjoint loop is considered.~~ The timing results are encouraging, indicating that the ~~reverse sweep timing comes closer the forwardsweep timing with larger-scale simulations~~ relative forward/adjoint timing observed in the test problem carries over to large-scale, realistic problems.

7 Discussion and conclusions

505 The fixed-point algorithm of *Christianson* (1994) has been successfully applied to the adjoint calculation of a land ice model. The algorithm is very relevant to the model code, as the bulk of the model’s computational cost is the solution of a nonlinear elliptic equation through fixed-point iteration. As many land ice models solve a similar fixed-point problem – particularly those intended to simulate fast-flowing outlet glaciers in Antarctica and Greenland – the methodology introduced
510 here has potential for the application of algorithmic differentiation techniques to other ice models. The implementation of the algorithm replaces a small portion of AD-generated code by handwritten code. However, this is done such that it does not interfere with the modularity offered by AD approach, and it does not require revision as model physics change.

The algorithm offers two advantages over the more straightforward “mechanical adjoint,” i.e. the
515 application of AD without intervention. First, the code solves the true adjoint to the fixed point iteration, rather than an approximation (*c.f.* Eq. 9). This avoids inaccurate results arising from “bad” initial guesses, and ensures proper convergence of the fixed-point adjoint. Second, the memory requirements do not increase with the number of adjoint iterations as they do with the mechanical adjoint. In the case of OpenAD, the effect on timing performance is small; but for machines with
520 limited memory or for larger problems, the large memory load associated with the mechanical adjoint will be a serious issue.

In the context of our ice model, the nature of the algorithm allows for further optimization, as it replaces the sequential solve of linear systems with differing matrices to a sequence of solves with the same matrix. Replacing the Conjugate Gradient solver of the forward model with a direct
525 L-U solver in the adjoint model leads to further performance improvement. The ratio of the reverse sweep to forward sweep, which is roughly the ratio of the run times of adjoint and forward models, decreases from 2.6 for the smallest problem considered to 1.4 for the largest. In the case where only a single time step is taken (not discussed above), no checkpoints are necessary, and the duration of the reverse sweep can be as little as 0.3 times the forward sweep.

530 It should be ~~pointed out that some authors have implemented ice model adjoint generation without any iteration within the adjoint model. Depending on the approximation to the Stokes momentum balance used, the adjoint stress balance can be derived directly from the equations involved (Perego et al., 2014; Isaac et al., 2015).~~ The result is a linear elliptic equation that can be solved without iteration, but which leads to a linear system that is far less sparse than in the forward model. Additionally, the equations must

535 potentially be re-derived if the model physics are changed. Moreover, not all such approximations
to noted, however, that the Stokes balance allow such an approach. “Hybrid” stress balances, which
solve two-dimensional approximations to the Stokes balance and are appropriate for both fast-sliding
and slow-creeping flow, are increasing in popularity due to low computational cost but reasonable
agreement with the First Order approximation e.g. *Goldberg (2011); Schoof and Hindmarsh (2010); Cornford et al. (2013); Arthern*
540 *W. Lipscomb, pers. comm.* Our ice model implements such a hybrid stress balance. Differentiating
such a balance at the equation level is possible but very tedious, and leads to very complicated
expressions that depend strongly on discretization (*Goldberg and Sergienko, 2011*), both undesirable
properties performance gain depends on the amortization of the L-U decomposition over the adjoint
iteration loop. If the L-U decomposition degrades in performance relative to the Conjugate Gradient
545 solve (a potential for large problems) or the number of iterations decreases, this gain could be lost.

As mentioned in the introduction, it is possible to differentiate the stress balance of an ice model at
the differential equation level rather than the code level. Such approaches, however, (a) cannot make
use of forward equation solvers, (b) remove somewhat the modularity of the AD approach, and (c)
are not suitable for “hybrid” models, which are becoming popular due to their balance between
550 generality and computational expense. Thus we argue that our application of the Christianson fixed-
point algorithm in our algorithmically differentiated ice model framework represents a contribution
to land ice modeling in general.

8 Code availability

All code necessary to carry out the experiments is publicly available through the MITgcm, Ope-
555 nAD and PETSc websites. Please see the supplement to the paper for detailed instructions regarding
installation and running of experiments.

9 Acknowledgements

This work was made possible in part through a SAGES (Scottish Alliance for Geoscience, Environ-
ment and Society) travel grant for early career exchange, NERC grant NE/M003590/1, ARCHER
560 Embedded CSE support grant eCSE03-09, and by a grant from the U.S. Department of Energy,
Office of Science, under contract DE-AC02-06CH11357. We are grateful for valuable input from
Editor Ham, Referee Christianson and one anonymous referee. Additionally the authors are grateful
for valuable input from B. Smith, J. Brown and P. Heimbach.

References

- 565 Arthern, R. J., R. C. A. Hindmarsh, and C. R. Williams, Flow speed within the Antarctic ice sheet and its controls inferred from satellite observations, *Journal of Geophysical Research: Earth Surface*, 120(7), 1171–1188, doi:10.1002/2014JF003239, 2014JF003239, 2015.
- Bartholomew-Biggs, M., S. Brown, B. Christianson, and L. Dixon, Automatic differentiation of algorithms, *Journal of Computational and Applied Mathematics*, 124(1–2), 171 – 190, 570 doi:http://dx.doi.org/10.1016/S0377-0427(00)00422-2, numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations, 2000.
- Blatter, H., Velocity and stress fields in grounded glaciers: a simple algorithm for including deviatoric stress gradients, *Journal of Glaciology*, 41, 333–344, 1995.
- Christianson, B., Reverse accumulation and attractive fixed points, *Optimization Methods and Software*, 3(4), 575 311–326, doi:10.1080/10556789408805572, 1994.
- Christianson, B., Reverse accumulation and implicit functions, *Optimization Methods and Software*, 9(4), 307–322, doi:10.1080/10556789808805697, 1998.
- Cornford, S. L., D. F. Martin, D. T. Graves, D. F. Ranken, A. M. Le Brocq, R. M. Gladstone, A. J. Payne, E. G. Ng, and W. H. Lipscomb, Adaptive mesh, finite volume modeling of marine ice sheets, *J. Comput. Phys.*, 580 232(1), 529–549, doi:10.1016/j.jcp.2012.08.037, 2013.
- Cuffey, K., and W. S. B. Paterson, *The Physics of Glaciers*, 4th ed., Butterworth Heinemann, Oxford, 2010.
- Dupont, T. K., and R. Alley, Assessment of the importance of ice-shelf buttressing to ice-sheet flow, *Geophys. Res. Lett.*, 32, L04,503, 2005.
- Errico, R. M., What is an adjoint model?, *BAMS*, 78, 2577–2591, 1997.
- 585 Favier, L., G. Durand, S. L. Cornford, G. H. Gudmundsson, O. Gagliardini, F. Gillet-Chaulet, T. Zwinger, A. Payne, and A. M. L. Brocq, Retreat of Pine Island Glacier controlled by marine ice-sheet instability, *Nature Climate Change*, 4, 117–121, doi:10.1038/nclimate2094, 2014.
- Forth, S., P. Hovland, E. Phipps, J. Utke, and A. Walther (Eds.), *Recent Advances in Algorithmic Differentiation, Lecture Notes in Computational Science and Engineering*, vol. 87, Springer, Berlin Heidelberg, 590 doi:10.1007/978-3-642-30023-3, 2012.
- Fretwell, P., et al., Bedmap2: improved ice bed, surface and thickness datasets for Antarctica, *The Cryosphere*, 7(1), 375–393, doi:10.5194/tc-7-375-2013, 2013.
- Giering, R., T. Kaminski, and T. Slawig, Generating efficient derivative code with TAF adjoint and tangent linear euler flow around an airfoil, *Future Gener. Comput. Syst.*, 21(8), 1345–1355, 595 doi:10.1016/j.future.2004.11.003, 2005.
- Gilbert, J. C., Automatic differentiation and iterative processes, *Optimization Methods and Software*, 1(1), 13–22, doi:10.1080/10556789208805503, 1992.
- Goldberg, D. N., A variationally-derived, depth-integrated approximation to a higher-order glaciological flow model, *Journal of Glaciology*, 57, 157–170, 2011.
- 600 Goldberg, D. N., and P. Heimbach, Parameter and state estimation with a time-dependent adjoint marine ice sheet model, *The Cryosphere*, 7(6), 1659–1678, doi:10.5194/tc-7-1659-2013, 2013.
- Goldberg, D. N., and O. V. Sergienko, Data assimilation using a hybrid ice flow model, *The Cryosphere*, 5, 315–327, doi:10.5194/tc-5-315-2011, 2011.

- Greve, R., and H. Blatter, *Dynamics of Ice Sheets and Glaciers*, Springer, Dordrecht, 2009.
- 605 Griewank, A., and A. Walther, Algorithm 799: Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation, *ACM Trans. Math. Softw.*, 26(1), 19–45, doi:10.1145/347837.347846, 2000.
- Griewank, A., and A. Walther, *Evaluating Derivatives. Principles and Techniques of Algorithmic Differentiation*, Vol. 19 of *Frontiers in Applied Mathematics*, 2nd ed., SIAM, Philadelphia, 2008.
- 610 Heimbach, P., The MITgcm/ECCO adjoint modeling infrastructure, *CLIVAR Exchanges*, 13(1), 13–17, 2008.
- Heimbach, P., and V. Bugnion, Greenland ice-sheet volume sensitivity to basal, surface and initial conditions derived from an adjoint model, *Annals of Glaciol.*, 50, 67–80, 2009.
- Heimbach, P., C. Hill, and R. Giering, Automatic generation of efficient adjoint code for a parallel Navier-Stokes solver, in 'Computational Science ICCS 2002, Vol. 2331, part 3 of *Lecture Notes in Computer Science*, edited by J. J. Dongarra, P. M. A. Sloot, and C. J. K. Tan, pp. 1019–1028, Springer-Verlag, 2002.
- 615 Hutter, K., *Theoretical Glaciology*, Dordrecht, Kluwer Academic Publishers, 1983.
- Isaac, T., N. Petra, G. Stadler, and O. Ghattas, Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the antarctic ice sheet, *Journal of Computational Physics*, 296, 348 – 368, doi:http://dx.doi.org/10.1016/j.jcp.2015.04.047, 2015.
- 620 Joughin, I., S. Tulaczyk, J. L. Bamber, D. Blankenship, J. W. Holt, T. Scambos, and D. G. Vaughan, Basal conditions for Pine Island and Thwaites Glaciers, West Antarctica, determined using satellite and airborne data, *Journal of Glaciology*, 55, 245–257, 2009.
- Khazendar, A., E. Rignot, and E. Larour, Larsen B ice shelf rheology preceding its disintegration inferred by a control method, *Geophys. Res. Lett.*, 34, L19503, doi:10.1029/2007GL030980, 2007.
- 625 Larour, E., E. Rignot, I. Joughin, and D. Aubry, Rheology of the Ronne Ice Shelf, Antarctica, inferred from satellite radar interferometry data using an inverse control method, *Geophys. Res. Lett.*, 32, L05,503, doi:10.1029/2004GL021693, 2005.
- Larour, E., J. Utke, B. Csatho, A. Schenk, H. Seroussi, M. Morlighem, E. Rignot, N. Schlegel, and A. Khazendar, Inferred basal friction and surface mass balance of the Northeast Greenland Ice Stream using data assimilation of ICESat (Ice Cloud and land Elevation Satellite) surface altimetry and ISSM (Ice Sheet System Model), *The Cryosphere*, 8(6), 2335–2351, doi:10.5194/tc-8-2335-2014, 2014.
- 630 Lipscomb, W., R. Bindschadler, E. Bueler, D. M. Holland, J. Johnson, and S. Price, A community ice sheet model for sea level prediction, *EOS Trans. AGU*, 90, doi:10.1029/2009EO030004, 2009.
- 635 Little, C. M., et al., Toward a new generation of ice sheet models, *EOS Trans. AGU*, 88, 578–579, 2007.
- MacAyeal, D. R., Large-scale ice flow over a viscous basal sediment: Theory and application to Ice Stream B, Antarctica, *Journal of Geophysical Research-Solid Earth and Planets*, 94, 4071–4087, 1989.
- MacAyeal, D. R., The basal stress distribution of Ice Stream E, Antarctica, inferred by control methods, *Journal of Geophysical Research*, 97, 595–603, 1992.
- 640 MacAyeal, D. R., R. A. Bindschadler, and T. A. Scambos, Basal friction of Ice Stream E, West Antarctica, *Journal of Glaciology*, 41, 247–262, 1995.
- Martin, N., and J. Monnier, Adjoint accuracy for the full Stokes ice flow model: limits to the transmission of basal friction variability to the surface, *The Cryosphere*, 8(2), 721–741, doi:10.5194/tc-8-721-2014, 2014.

- McGovern, J., I. Rutt, J. Utke, and T. Murray, ADISM v.1.0: an adjoint of a thermomechanical ice-sheet
 645 model obtained using an algorithmic differentiation tool, *Geoscientific Model Development Discussions*,
 6(4), 5251–5288, doi:10.5194/gmdd-6-5251-2013, 2013.
- McMillan, M., A. Shepherd, A. Sundal, K. Briggs, A. Muir, A. Ridout, A. Hogg, and D. Wingham, In-
 creased ice losses from Antarctica detected by CryoSat-2, *Geophysical Research Letters*, 41(11), 3899–3905,
 doi:10.1002/2014GL060111, 2014GL060111, 2014.
- 650 Morland, L. W., Unconfined ice-shelf flow, in *Dynamics of the West Antarctic Ice Sheet*, edited by C. J. V. der
 Veer and J. Oerlemans, pp. 99–116, Reidel Publ Co, 1987.
- Morlighem, M., E. Rignot, G. Seroussi, E. Larour, H. Ben Dhia, and D. Aubry, Spatial patterns of basal drag in-
 ferred using control methods from a full-stokes and simpler models for Pine Island Glacier, West Antarctica,
Geophys. Res. Lett., 37, L14,502, doi:10.1029/2010GL043853, 2010.
- 655 Naumann, U., *The Art of Differentiating Computer Programs: An Introduction to Algorithmic Differentiation*,
 no. 24 in Software, Environments, and Tools, SIAM, Philadelphia, PA, 2012.
- Pattyn, F., A new three-dimensional higher-order thermomechanical ice-sheet model: basic sensitivity, ice-
 stream development and ice flow across subglacial lakes, *Journal of Geophysical Research-Solid Earth and
 Planets*, 108, doi:10.1029/2002JB002329, 2003.
- 660 Pattyn, F., et al., Benchmark experiments for higher-order and full-Stokes ice sheet models (ISMIP HOM), *The
 Cryosphere, Volume 2, Issue 2, 2008, pp.95-108*, 2, 95–108, 2008.
- Perego, M., S. Price, and G. Stadler, Optimal initial conditions for coupling ice sheet models to earth system
 models, *Journal of Geophysical Research: Earth Surface*, 119(9), 1894–1917, doi:10.1002/2014JF003181,
 2014.
- 665 Petra, N., H. Zhu, G. Stadler, T. J. Hughes, and O. Ghattas, An inexact Gauss-Newton method for inversion of
 basal sliding and rheology parameters in a nonlinear Stokes ice sheet model, *Journal of Glaciology*, 58(211),
 889–903, doi:doi:10.3189/2012JoG11J182, 2012.
- Rignot, E., J. Mouginot, and B. Scheuchl, Ice flow of the Antarctic Ice Sheet, *Science*, 333(6048), 1427–1430,
 doi:10.1126/science.1208336, 2011.
- 670 Rommelaere, V., Large-scale rheology of the Ross Ice Shelf, Antarctica, computed by a control method, *Journal
 of Glaciology*, 24, 694–712, 1997.
- Schoof, C., and R. C. A. Hindmarsh, Thin-film flows with wall slip: An asymptotic analysis of higher order
 glacier flow models, *Quart. J. Mech. Appl. Math.*, 63, 73–114, 2010.
- Sergienko, O. V., R. A. Bindschadler, P. L. Vornberger, and D. R. MacAyeal, Ice stream basal conditions from
 675 block-wise surface data inversion and simple regression models of ice stream flow: Application to Bind-
 schadler Ice Stream, *Journal of Geophysical Research*, 113, F04,010, doi:10.1029/2008JF001004, 2008.
- Shepherd, A., D. J. Wingham, and J. Mansley, Inland thinning of the Amundsen Sea sector, West Antarctica,
Geophys. Res. Lett., 29, L1364, doi:10.1029/2001GL014183, 2002.
- Shepherd, A., D. J. Wingham, and E. Rignot, Warm ocean is eroding West Antarctic Ice Sheet, *Geophys. Res.
 680 Lett.*, 31, L23,402, 2004.
- Taftaf, A., L. Hascoët, and V. Pascual, Implementation and measurements of an efficient Fixed Point Adjoint,
 in *EUROGEN 2015, ECCOMAS, GLASGOW, UK*, 2015.

- Utke, J., U. Naumann, M. Fagan, N. Tallent, M. Strout, P. Heimbach, C. Hill, D. Ozyurt, and C. Wunsch, OpenAD/F: A modular open source tool for automatic differentiation of Fortran codes, *ACM Transactions on Mathematical Software*, 34, 2008.
- 685
- Vaughan, D. G., and R. Arthern, Why is it hard to predict the future of ice sheets?, *Science*, 315(5818), 1503–1504, doi:10.1126/science.1141111, 2007.
- Vieli, A., and A. J. Payne, Application of control methods for modelling the flow of Pine Island Glacier, West Antarctica, *Annals of Glaciol.*, 36, 197–204, 2003.

Table 1. Pseudocode version of forward model time-stepping procedure.

```

FOR  $n$  = initialTimeStep TO finalTimeStep
  // Constructs  $\hat{\mathbf{a}}$  from  $H^{[n]}$  :
  CALL CALC_DRIVING_STRESS ( $H^{[n]}$ )
   $m = 0$ 
  REPEAT UNTIL CONVERGENCE OF  $\mathbf{u}$ 
     $\mathbf{u} = \Phi(\mathbf{u}, \hat{\mathbf{a}})$ 
     $m = m+1$ 
    store  $L$ ,  $\mathbf{u}$  and other variables
   $lastm^{[n]} = m$ 
  // Finds  $H^{[n+1]}$  from continuity equation with  $\mathbf{u}$ :
  CALL ADVECT_THICKNESS ()

```

Table 2. Pseudocode version of mechanical adjoint.

```

FOR  $n$  = finalTimeStep DOWNTO initialTimeStep
  // Constructs  $\delta^* H^{[n]}$  and  $\delta^* \mathbf{u}^{[n]}$  from  $\delta^* H^{[n+1]}$ 
  // via the adjoint of the continuity equation :
  CALL AD_ADVECT_THICKNESS ()
  REPEAT  $lastm^{[n]}$  TIMES
    restore  $L$ ,  $\mathbf{u}$  and other variables
     $\delta^* \hat{\mathbf{a}} = \delta^* \hat{\mathbf{a}} + \delta^* \mathbf{u} \left( \frac{\partial \Phi}{\partial \hat{\mathbf{a}}} \right)^T$ 
     $\delta^* \mathbf{u} = \delta^* \mathbf{u} \left( \frac{\partial \Phi}{\partial \mathbf{u}} \right)^T$ 
  // Updates  $\delta^* H^{[n]}$  from  $\delta^* \hat{\mathbf{a}}$  :
  CALL AD_CALC_DRIVING_STRESS ( $\delta^* H^{[n]}$ )

```

Table 3. Pseudocode version of modified forward model for BC94.

```
FOR  $n = \text{initialTimeStep}$  TO  $\text{finalTimeStep}$ 
  // Constructs  $\hat{\mathbf{a}}$  from  $H^{[n]}$  :
  CALL CALC_DRIVING_STRESS ( $H^{[n]}$ )
   $\mathbf{u} = \text{initial guess}$ 
  CALL PHISTAGE (PRELOOP,  $\mathbf{w}$ ,  $\mathbf{u}$ ,  $\hat{\mathbf{a}}$ )
  REPEAT UNTIL CONVERGENCE OF  $\mathbf{u}$ 
    CALL PHISTAGE (INLOOP,  $\mathbf{w}$ ,  $\mathbf{u}$ ,  $\hat{\mathbf{a}}$ )
  CALL PHISTAGE (POSTLOOP,  $\mathbf{w}$ ,  $\mathbf{u}$ ,  $\hat{\mathbf{a}}$ )
  // Finds  $H^{[n+1]}$  from continuity equation with  $\mathbf{u}$ :
  CALL ADVECT_THICKNESS ()

SUBROUTINE PHISTAGE (phase,  $\mathbf{w}$ ,  $\mathbf{u}$ ,  $\hat{\mathbf{a}}$ )
  IF (phase==PRELOOP)
    // do nothing

  ELSE IF (phase==INLOOP)
    save tape pointer
     $\mathbf{u} = \Phi(\mathbf{u}, \hat{\mathbf{a}})$ 
    // Makes sure no storage is done :
    restore tape pointer

  ELSE IF (phase==POSTLOOP)
     $\mathbf{u} = \Phi(\mathbf{u}, \hat{\mathbf{a}})$ 
    store  $L$ ,  $\mathbf{u}$  and other variables
```

Table 4. Pseudocode version of fixed-point (BC94) adjoint.

```

FOR n = finalTimeStep DOWNTO initialTimeStep

    // Constructs  $\delta^* H^{[n]}$  and  $\delta^* \mathbf{u}$  from  $\delta^* H^{[n+1]}$ 
    // via the adjoint of the continuity equation :
    CALL AD_ADVECT_THICKNESS ()
    CALL AD_PHISTAGE (POSTLOOP,  $\delta^* \mathbf{w}$ ,  $\delta^* \mathbf{u}$ ,  $\delta^* \hat{\mathbf{a}}$ )
    REPEAT UNTIL CONVERGENCE OF  $\delta^* \mathbf{w}$ 
        CALL AD_PHISTAGE (INLOOP,  $\delta^* \mathbf{w}$ ,  $\delta^* \mathbf{u}$ ,  $\delta^* \hat{\mathbf{a}}$ )
    CALL AD_PHISTAGE (PRELOOP,  $\delta^* \mathbf{w}$ ,  $\delta^* \mathbf{u}$ ,  $\delta^* \hat{\mathbf{a}}$ )
     $\delta^* \mathbf{u} = 0.0$ 
    // Updates  $\delta^* H^{[n]}$  from  $\delta^* \hat{\mathbf{a}}$  :
    CALL AD_CALC_DRIVING_STRESS ( $\delta^* H^{[n]}$ )

SUBROUTINE AD_PHISTAGE (phase,  $\delta^* \mathbf{w}$ ,  $\delta^* \mathbf{u}$ ,  $\delta^* \hat{\mathbf{a}}$ )

    IF (phase==POSTLOOP)
         $\delta^* \mathbf{w} = \delta^* \mathbf{u}$ 

    ELSE IF (phase==INLOOP)
        save tape pointer
        restore  $L$ ,  $\mathbf{u}$  and other variables
         $\delta^* \mathbf{w} = \delta^* \mathbf{w} \left( \frac{\partial \Phi}{\partial \mathbf{u}} \right)^T + \delta^* \mathbf{u}$ 
        // Makes sure converged state is reused :
        restore tape pointer

    ELSE IF (phase==PRELOOP)
         $\delta^* \hat{\mathbf{a}} = \delta^* \mathbf{w} \left( \frac{\partial \Phi}{\partial \mathbf{a}} \right)^T$ 

```




Figure 1. (a) Surface speed (shading) in the test experiment. The flow direction is from right to left, and the white portion of the figure is where the ice shelf has not advanced to the end of the domain. Black contours give thickness spaced every 200 m and the white contour is the grounding line. (b) Adjoint sensitivities of ice speed to basal melt rates. (c) (log) relative discrepancy between adjoint sensitivities and the gradient calculated via finite differencing. (d) 2nd order (log) second order differencing of cost function J (see eq. 17).

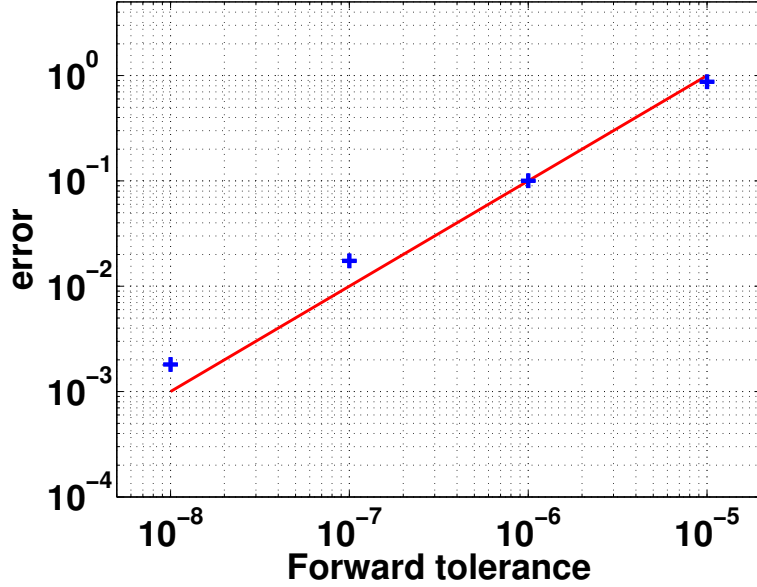


Figure 2. Maximum error in fixed-point adjoint calculation versus tolerance of forward loop. The red line indicates linear dependence.

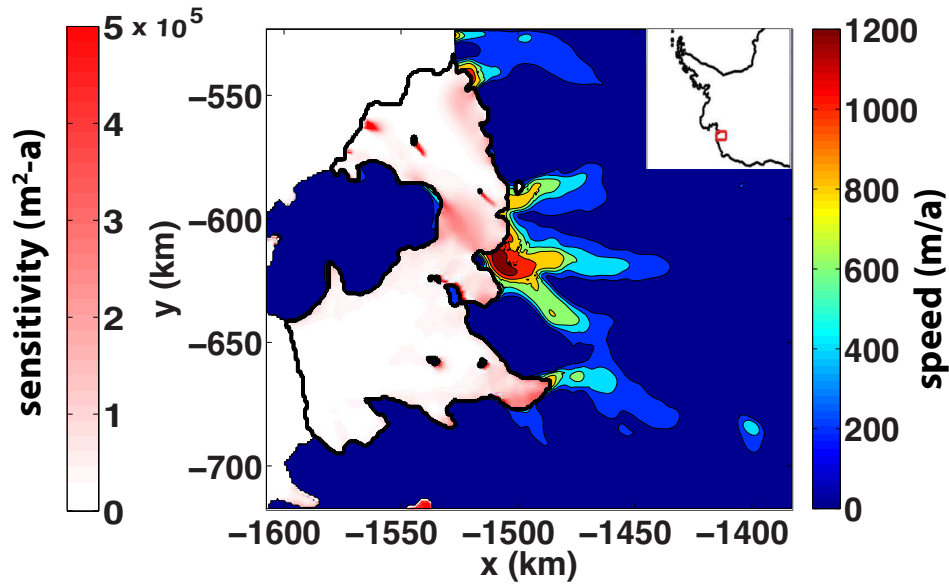


Figure 3. Adjoint sensitivity of loss of Volume above Floatation (VAF) to basal melting under the ice shelves adjacent to Smith Glacier (location shown in inset). Filled contours give modeled ice velocity where ice is grounded; red-white shading gives adjoint melt rate sensitivity under ice shelves. The thick black contour denotes the boundary of the ice shelves.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.