

Interactive comment on “An Approach to Computing Discrete Adjoints for MPI-Parallelized Models Applied to the Ice Sheet System Model” by Eric Larour et al.

L. Hascoet (Referee)

laurent.hascoet@inria.fr

Received and published: 27 May 2016

General comments:

The paper describes the construction of the adjoint of the ice sheet simulator ISSM (written in C++) through the use of the overloading-based AD tool ADOL-C, and through the use of the MPI wrapper library AMPI (Adjoinable MPI). The paper presents the first meaningful gradients obtained with this adjoint ISSM, and discusses performance aspects. The main focus of the paper is about the programming strategy that was chosen to ease the algorithmic differentiation process with ADOL-C, the specific efficient tactic to propagate the adjoint through calls to library linear solvers, and the constraints and benefits of using the AMPI library to get an MPI-parallel adjoint.

C1

Although not a specialist of the application domain, I have the feeling that the main similar efforts by glaciology researchers are presented and compared adequately. The paper claims that the adjoint ISSM will allow for full exploitation of the enormous quantity of measurement data available. This is something one can believe in.

The strategies proposed in sections 2.2 to 2.4 go into details that shed an interesting light on ADOL-C (questions of location handling and contiguity that I was not aware of, thanks) and on the technicalities of the linear solver "trick". Many of these are of general interest for the AD people, overloaders or source-transformers alike.

Text is well written, in excellent English. My only minor concern is about its high level of technicity in Computer Science, which might make it hard to read at places, especially for readers specialized rather in geophysics models. But this high level makes it a profitable read for computer science (in particular AD) people.

To summarize, I like the paper. I think it is of value for developers of models, and for AD tool developers as well. Getting this value requires effort while reading, though.

Specific comment:

Still, we face this old dilemma between source-transformation AD (probably more efficient) applied to codes in old-fashioned languages (Fortran or C), versus overloading-based AD (slightly heavier in memory) applied to codes in modern languages (nicer to develop). Sentence at line 370 obviously gives the authors' opinion on that, but I'm wondering if things are really that terrible. After all, in other application domains, people do use source-transformation AD (when available).

On AMPI in this context, the paper might give more hints on what happens while interpreting the trace R(T) i.e. when an AMPI call is registered in T, R(T) calls AMPI primitives at that time too. Otherwise, the reader might wonder how this can work.

Technical corrections:

– Line 19 "alreay"

C2

- I don't understand the n^2 vs n at line 90. Am I missing something?
- Sentence at line 140 look ambiguous to me. Maybe a few commas or small words would help?
- Using * and ** in table 2 is unfortunate: guaranteed confusion with the C dereference operator.
- Paragraph at lines 269-275 is hard to follow. Can it be rephrased ?
- Line 368 "critical"

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-99, 2016.