

Interactive comment on “P-CSI v1.0, an accelerated barotropic solver for the high-resolution ocean model component in the Community Earth System Model v2.0” by Xiaomeng Huang et al.

E. Müller (Referee)

e.mueller@bath.ac.uk

Received and published: 25 August 2016

1 General comments

In their paper: "P-CSI v1.0, an accelerated barotropic solver for the high resolution ocean model component in the Community Earth System Model v2.0" the authors discuss the implementation of a new preconditioned iterative solver algorithm for the solution of the elliptic PDE which arises in implicit time stepping of the barotropic mode in ocean models. By using an iterative method which avoids global communications,

[Printer-friendly version](#)

[Discussion paper](#)



the scalability of the solver can be improved significantly on large core counts. The authors demonstrate that this leads to substantial performance improvements when running the model at high spatial resolution on more than 16,000 cores of the Yellowstone supercomputer.

Accurate and scalable models are absolutely essential for reliable predictions of the Earth's climate which have wide impact in the geoscience community. In the introduction the authors argue convincingly why the development of high resolution ocean models and of massively parallel elliptic solvers is necessary and their novel algorithm approach addresses an important bottleneck for scalability on large core counts (global parallel reductions). Since the barotropic solver accounts for a large fraction in the runtime, this work has a large impact for the numerical model they study and can also help to improve related models in atmospheric- and ocean- modelling. The work is put into context by referring to relevant related publications and the paper is very well written throughout, with the results supporting the theoretical analysis (in particular the theoretical performance analysis). The scientific results, in particular the use of a communication-avoiding iterative method as an alternative to "standard" Krylov-subspace methods such as CG, are very interesting and the benefits of the method are demonstrated convincingly by detailed numerical experiments.

As stated at the end of the introduction, this paper is based on a related conference proceedings publication [1], where the key ingredients of the algorithm and parallel scaling tests are described in detail. Compared to [1], the present GMD paper contains the following new material:

1. The properties of the discretised system and in particular the spectral radius of the matrix is derived for a simplified test setup (constant ocean depth). By a theoretical analysis of the convergence rate of the new P-CSI algorithm the authors argue that it converges as fast as the CG and ChronGear solvers which require additional global reductions

[Printer-friendly version](#)[Discussion paper](#)

2. Numerical estimates of the spectral radius for different aspect ratios and time step sizes are presented
3. The detailed convergence history of different solvers/preconditioners is studied

However, the key concepts of the solver/preconditioner setup and similar results are already given in [1] (for example for the convergence numbers, compare Fig. 6 in [1] and Figs. 9–11 in the present paper, Fig. 13–15 are a subset of results from [1] and as far as I can tell Fig. 10 and 12 are obtained with a similar setup as in [1]). I'm therefore slightly concerned as to whether this paper contains sufficient new material to for a new publication, in particular since I'm not sure how relevant the variations in time step size really are in practice - the time step size is largely fixed by the CFL limit in other model components, such as the advective time scale (see further discussion below). To publish the paper it has to be made clear that large parts of it consist of new results. The solver code is made available online, but I was not able to compile and run it since it requires installation of the full model.

2 Specific comments

- Early on in the introduction and in section 2.2 the authors mention that global reductions limit the performance of ChronGear and CG solvers, and this is one of the main motivations for using the P-CSI method. While this is clearly shown in Fig. 10, it might be good to already refer to numerical evidence here or quote numbers from [1]: which percentage of the runtime is spent in global reduction operations?
- In section 2 the authors derive the barotropic mode in the fundamental equations and then discretise it implicitly in time. At this point it might be good to briefly

[Printer-friendly version](#)[Discussion paper](#)

mention how this is related to other model components: how is the implicitly calculated height perturbation coupled back to the full equations? Are other parts of the equations (such as advection) solved explicitly? What are the typical time velocities/time scales (I assume that the implicit treatment of the barotropic mode is necessary since the gravity-wave velocity $\sqrt{g * H}$ is much larger than other velocities in the system, such as advective velocities - is this correct?).

- When showing strong scaling results such as in Figs. 3, 13-15 the number of unknowns per processor is relevant to assess the relative importance of halo exchange, could this information be added to the figures?
- Discussion in section 4.1: the barotropic CFL number due to gravity waves of speed $c_g = \sqrt{gH}$, $n_g = c_g \cdot dt/dx$, is a very important quantity and for a given resolution directly related to the time step size. However, typically the time step size dt is limited by other processes in the model. For example, if there is another process with typical speed c' , which is treated explicitly, then the related CFL number $n' = c' \cdot dt/dx$ is limited by $n' < O(1)$. For example in atmospheric models $c_{advection} \approx 10 \cdot c_{acoustic}$, and hence the n_g should not be larger than ≈ 10 . Could the authors include a discussion of this and also discuss physical limits on n_g by referring to other components (i.e. non-barotropic) of the model? I think this is very important since the CFL number has a significant impact on the solver performance. By using $c_g = \sqrt{9.81m/s^2 \times 4km} = 200m/s$ the large time step sizes in Fig. 5 seems to be completely unphysical if I assume that there is another explicitly treated process in the model which is $\approx 10\times$ slower than the gravity waves, but my intuition from atmospheric models might be misleading here and if the explicitly treated non-barotropic dynamics happens at much larger time scales then those large time steps make sense. This seems to be implied by the setup used for the 0.1 degree runs: assuming a depth of 4km, a time step size of 172.8s would lead to a CFL number of $\approx 10^4$. Since the condition number depends largely on the CFL number it would be good to see what the physically

[Printer-friendly version](#)[Discussion paper](#)

relevant values are.

- The condition $dt = dx/v$, which is imposed at the bottom of page 9 should be clarified. The authors refer to v as the "barotropic velocity" and then vary this between 2m/s and 200m/s. Should this v be some other velocity in the system which limits the time step size? The velocity relevant for the barotropic equation is the gravity wave speed $c_g = \sqrt{gH}$ which is 200m/s for a depth of 4km. Same question in section 5.1, where the authors fix $v = 2m/s$.
- It would help if the CFL number and (an estimate of the) condition number of the matrix are given for the realistic 0.1 degree run. Since the largest and smallest eigenvalue are estimated this information should be available.
- If the CFL numbers are very large (see previous points), then I really think that advanced preconditioners have the potential for improving the performance. Multi-grid preconditioners could reduce the iteration count from $O(100)$ to $O(10)$, so might pay off even if one preconditioner application is more expensive.
- In Fig. 5 it would be good to indicate the range of typical physical time step sizes for each resolution instead of just plotting a wide range of time scales.
- Page 13, line 394: while the matrix becomes more ill conditioned as the problem size increases, the condition $dt = dx/v$ will limit this growth, in fact the upper bound on the condition number is at the order of $\approx gH/v^2$.
- The theoretical analysis is carried out for a constant ocean depth H . How reasonable is this assumption and which impact do variations in H have?

3 typos/minor comments

- at several places in the paper "scaler" should be replaced by "scalar"

- to me "boundary communication" is a slightly unusual expression, I'd call this "halo exchange" since "boundary" could refer to a physical boundary in the global domain (such as the ocean-land interface).
- at the bottom of page 4: should this read "[...] the barotropic continuity Eq. (4) *has been* linearised [...]" ("is linearised" implies that another term has to be removed from (4) to obtain a linear equation, but (4) is already linear).
- bottom of page 8: "spectrum radius" -> "spectral radius"
- definition of $P_k(\xi)$ between Eqs. (19) and (20) on page 10: What are α and β here?
- in appendix A and B it might help if the global reduction operations in steps 2. and 5. of PCG and steps 3. and 4. of ChronGear are highlighted. Also, a sentence to the appendix which clarifies that the global reduction of ρ_k and σ_k in the ChronGear algorithm can be combined (thus halving the latency) might help.
- Fig. 4 does not add relevant information and should be removed
- Fig. 2: replace "sparse pattern" -> "sparsity pattern"

references

[1] Hu et al., 27th International Conference for High Performance Computing, Networking, Storage and Analysis (SC2015), 2015

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-135, 2016.

Printer-friendly version

Discussion paper

