

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.

Xiaomeng Huang et al.

hxm@tsinghua.edu.cn

Received and published: 7 September 2016

Dear Dr. Müller:

We would like to express our sincere appreciation to your valuable feedback. Your comments are highly insightful and enable us to significantly improve our manuscript. The following pages are our point-by-point responses to each of your comments.

[Printer-friendly version](#)

[Discussion paper](#)



1 Response to 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, 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 discretized 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.

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 am 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.

[Response]:

Thank you for your highly valued comments that our study has a large impact for the numerical model development and can help to improve related models in atmospheric and oceanic modelling.

[Printer-friendly version](#)

[Discussion paper](#)



This paper is an extended work originally presented in the 27th International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2015) as we indicated in Section I. Most of the audiences in the SC conference are supercomputing specialists. Therefore, we simplified the background of the ocean model and focused on the design of algorithm, scalability tests and efficiency in the SC paper. We introduced our solver with some pseudo-code and used a lot of computer terminologies since the main readers are computer scientists. In order to expand the influence of our work to more general readers, we decide to extend our paper to GMD which is an outstanding academic exchange platform for climate modelers. Therefore, we made a lot of changes to the content and structure of the GMD manuscript comparing with the SC paper. Although you have mentioned part of them in the general comments, we summarize our major changes as follows:

(1) We enriched the review of barotropic mode and introduced two solvers including PCG and ChronGear adopted in the original POP. We believe that this important part will help other climate modelers to comprehensively understand a general large-scale computing problem in POP, MOM, MITgcm, FVCOM, OPA models etc. They can associate their own climate models with our new solver through the detailed barotropic equations. In SC paper, we only give a simple version of barotropic mode and the ChronGear solver.

(2) We combined sections 3 and 4 in SC paper into the section "Design of the P-CSI solver" in GMD manuscript. In order to make the climate modelers instead of computer scientists to better understand the new solver, we rewrote most of sentences and moved the pseudo-code and the procedure of preconditioned into the appendix for interested readers. We also avoided the use of obscure computer jargon to make it more readable.

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

(3) After our presentation at SC conference last year, many helpful advices were gathered. Some specialists pointed out that we should provide more information about the universal applicability of our new solver in different situations and different applications. Therefore, in the GMD manuscript, a new section 4 is dedicated to the characteristics of the new approach. We theoretically analyzed the characteristics of P-CSI through the associated eigenvalues and their connection with the convergence rate. Based on different solvers/preconditioners, we derived the properties of the discretized system and, in particular, the spectral radius of the matrix by theoretical analysis of the Spectrum, condition number and the convergence rate. We concluded that P-CSI can converge as fast as the CG and ChronGear solvers which require additional global reductions. In SC paper, we only presented the computational complexity.

We agree that there are some similar materials in the GMD manuscript and SC paper in section 5 because this section presents the actual computational evaluation about the actual performance. We will make some adjustments as follows:

(1) Fig. 9 is pretty new (not in the SC paper) and verifies the theoretical analysis of the convergence rate of different barotropic solvers, we will merge it with Fig. 11 which is also new and confirms the improvement of convergence rate. We will add some slopes on Fig. 11 to reflect the condition numbers improved in the different pre-conditioners.

(2) The material we discuss about Figs. 10, and 12 is very similar to that in SC paper. I believe it is still very important to emphasize that the global reduction is the major bottleneck in the large-scale computing ($>O(10^3)$ processor counts) in the GMD paper. The P-CSI with EVP preconditioner enhances the performance by significantly reducing the iteration number so that the time for computation part is further reduced. In the revised version, we will further combine Fig. 10 and

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

12 and reemphasize this point by shortening the discussion and an overall comparison.

(3) We will eliminate Fig. 13 because it illustrates the problem of ChronGear + Diagonal and doesn't add much. Figs. 14 and 15 will be reduced into one figure. In the SC paper, we presented the analysis for barotropic solver only. Here, we will extend to the overall performance of model by including the timing for other components. This information is very important because the performance of 1/10 degree POP is not clearly documented yet. So, we will put the total time (all components). This will make a big difference and show the major advantage of our new approach.

According to your concern about the unphysical time step size in Fig.5, we will use a typical value for the first baroclinic wave speed of 2m/s and a typical gravity wave velocity of 200m/s to reanalyze the solver property. In the 0.1 degree ocean simulation, the CFL number is about $c\Delta t/\Delta x \approx 3.46$, thus the corresponding time step size are 17280s and 172.8s. We also believe it will make more senses to use CFL number as the x-axis in some plots (see the following responses) instead of time-step (which is also connected by dx). Therefore, all relevant figures will be changed accordingly.

In conclusion, after making the above modifications, we believe that this GMD manuscript will be an expanded, more fully developed, and more refined version of the conference paper. These new results are confirmed by a valuable theoretical analysis. All figures have completely different stories and key messages comparing the SC paper. We will clearly introduce these differences in the revised manuscript if we have the opportunity.

[Printer-friendly version](#)

[Discussion paper](#)



2 Response to specific comments

(1) 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?

[Response]:

In the introduction, we will add the following sentence,

"For example, when around four thousand cores are used, the global reduction in PCG and ChronGear takes approximately 74% and 68% of the whole barotropic time, respectively [1]. This situation will get worse with more cores."

Also, Fig. 10 will be merged with Fig. 12 so that we can reemphasize the advantage of our proposed method. This discussion will be shortened.

(2) 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 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?).

[Response]:

Thanks for your good suggestion. In the beginning of section 2.1, we will briefly introduce the following time scheme.

"POP uses the splitting technique to solve the barotropic and baroclinic system [1]. All terms in the Eq. (1) use the explicit scheme except the implicit treatment of barotropic mode and semi-implicit treatment of Coriolis and vertical mixing terms. The implicit

treatment of barotropic mode is necessary to simulate the fast gravity waves with the speed of $\sqrt{g * H} \approx 200\text{m/s}$ so that we can use the same time step as the baroclinic mode which has a velocity scale less than 2 m/s [1]. The time step of the 0.1 degree POP model is 172.8s."

Yes, you are correct that the implicit approach is necessary to match the same time step between the barotropic and baroclinic modes. Hope the above brief discussion can give the readers an overall description.

(3) 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?

[Response]:

This advice is valuable and we will add this information to Fig. 3 by including the new axis labeled as "number of grids per core" . We will eliminate Fig.13 because it just says the problem of ChronGear + Diagonal. The Figs. 14 and 15 will be reduced into one figure. We will refer to Fig. 3 for this information.

(4) Discussion in section 4.1: the barotropic CFL number due to gravity waves of speed $c_g = \sqrt{gH}$, $n_g = c_g dt/dx$, is a very important quantity and for a given resolution directly related to the time step size. However, typically the time step size 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 $n' = c' dt/dx$ is limited by $n' < O(1)$. For example in atmospheric models $c_{advection}^{10} \cdot c_{acoustic}$, 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 $c_g = \sqrt{9.81\text{m/s}^2 \times 4\text{km}} = 200\text{m/s}$ 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

Printer-friendly version

Discussion paper



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 relevant values are.

[Response]:

Thanks for your corrections. Considering the limitation of CFL condition, the time step size in Fig. 5 is beyond the physical range indeed. The original purpose of this selection is to make it more intuitive to readers that the time step sizes have a large influence on the condition number of the coefficient matrix without taking the physical consistency into account. We will redraw Fig.5 with the CFL number as the x-axis and make the values more credible in physics. In 0.1 degree ocean simulation, the time step size is $\Delta t=172.8\text{s}$ (500 steps per simulation day), thus for the barotropic mode, the CFL number is about $c \cdot \Delta t/\Delta x \approx 3.46$ where $c = 200\text{m/s}$ is the gravity wave velocity in the barotropic mode and $\Delta x = 10000\text{m}$ is the horizontal grid length. If $c' = 2\text{m/s}$ is a typical value for the first baroclinic wave speed, the CFL number is less than $c' \cdot \Delta t/\Delta x \approx 0.035$. CFL numbers varying from 0.01 to 5 will be used in Fig. 5 to cover more physically relevant cases in POP.

(5) 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 $v = 2\text{m/s}$.

[Response]:

Thanks for your corrections. We will use the typical value for the first baroclinic wave

Printer-friendly version

Discussion paper



speed 2m/s and the typical gravity wave velocity $c_g = \sqrt{gH} = 200m/s$ for a depth of 4km as the lower bound and the upper bound of the velocity range, respectively. Besides, we will use the non-dimensional barotropic CFL number instead of velocity as the legend of Fig. 7 to show the dependency.

(6) 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.

[Response]:

We will add this information in Section 4.1.

"In 0.1 degree realistic run, the CFL number is about $c \cdot \Delta t / \Delta x \approx 3.46(c = 200m/s, \Delta t = 172.8s, \Delta x = 10000m$ is the typical gravity wave speed, time step and spatial resolution, respectively) and the condition number is about 250."

(7) If the CFL numbers are very large (see previous points), then I really think that advanced preconditioners have the potential for improving the performance. Multigrid preconditioners could reduce the iteration count from $O(100)$ to $O(10)$, so might pay off even if one preconditioner application is more expensive.

[Response]:

As you point out, any advanced preconditioner which can quickly reduce the iteration count will be very useful to improve the performance. In fact, the EVP solver is a direct fast solver so that it has this capability and is the main reason we choose here. Furthermore, as we indicated in lines 66-77 of section 1, the multigrid method is a well-known scalable and efficient approach to solve the elliptic systems too. However, some related works confirmed that the geometric multigrid in global ocean models does not always scale ideally because of the presence of complex topography (land particularly), non-uniform or anisotropic grids ([2], [3], [4], [5], [6]). These constraints lead to an elliptic system with variable coefficients defined on an irregular domain in POP and complicate the modeling system. The algebraic multigrid (AMG) is an

Printer-friendly version

Discussion paper



alternative to the geometric multigrid to handle the complex topography. However, the AMG setup in the parallel environment is more expensive than the iterative solver in climate modelling, which makes it unfavorable as a preconditioner [2]. On the contrary, the EVP preconditioner is simple enough and can effectively reduce the condition number of coefficient matrix by about 5 times in both 1 and 0.1 degree cases, which leads to a reduction of 2/3 iterations. Therefore, we use the direct EVP solver. Even so, more study about the preconditioner in practical climate models will be very useful and we will take it as our future work.

(8) 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.

[Response]:

We will replace the time step sizes in the x-axis of Fig. 5 with the non-dimensional CFL number and make the values more credible in physics.

(9) 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 $\approx gH/v^2$.

[Response]:

Thanks for the corrections. We will add the following sentences in line 394 to make our revised version more rigorous.

"As shown in Fig. 8, when the problem size increases, the coefficient matrix becomes more poorly conditioned until it reaches the upper bound at the order of gH/v^2 ."

(10) 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?

[Response]:

The purpose of this assumption is to succinctly demonstrate the properties of the sparse matrix used in the POP. This assumption is not convincing in physics indeed,

Printer-friendly version

Discussion paper



while it provides a bound for cases with various H . We will expand our discussion to support variations in H , which should lead to similar results.

3 Response to typos/minor comments

(1) at several places in the paper "scaler" should be replaced by "scalar"

[Response]: We will replace "scaler" with "scalar" in our revised version.

(2) 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).

[Response]: We will replace "boundary communication" with "halo exchange" in our revised version.

(3) 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).

[Response]: We will change "is linearized" into "has been linearized" in our revised version.

(4) bottom of page 8: "spectrum radius" -> "spectral radius"

[Response]: We will revise "spectrum radius" with "spectral radius" in our revised version.

(5) definition of $P_k(\xi)$ between Eqs. (19) and (20) on page 10: What are α and β here?

[Response]: We will introduce the meaning of α and β in our revised version. The α is

Printer-friendly version

Discussion paper



the aspect ratio and β is the reciprocal of α which are defined in Eq. (13).

(6) 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.

[Response]: We will highlight these steps and add some sentences to clarify the global reduction operations. In line 515, we will add "these inner products use two global reduction operations."; In line 535, we will add "The inner products in ρ_k and σ_k use two global reduction operations. However, these two global reductions can be combined into one operation thus halving the latency."

(7) Fig. 4 does not add relevant information and should be removed.

[Response]: We will remove Fig.4 and related paragraphs in our revised version.

(8) Fig. 2: replace "sparse pattern" -> "sparsity pattern"

[Response]: We will replace "sparse pattern" with "sparsity pattern" in our revised version.

4 References

- [1] Hu et al., 27th International Conference for High Performance Computing, Networking, Storage 395 and Analysis (SC2015), 2015
- [2] Müller, E. H. and Scheichl, R.: Massively parallel solvers for elliptic partial differential equations in numerical weather and climate prediction, Quarterly Journal of the Royal Meteorological Society, 140, 2608-2624, 2014.
- [3] Matsumura, Y. and Hasumi, H.: A non-hydrostatic ocean model with a scalable

multigrid Poisson solver, Ocean Modelling, 24, 15-28, 2008.

[4] Kanarska, Y., Shchepetkin, A., and McWilliams, J.: Algorithm for non-hydrostatic dynamics in the regional oceanic modeling system, Ocean Modelling, 18, 143-174, 2007.

[5] Fulton, S. R., Ciesielski, P. E., and Schubert, W. H.: Multigrid methods for elliptic problems: A review, Monthly Weather Review, 114, 943-959, 1986.

[6] Stüben, K.: A review of algebraic multigrid, Journal of Computational and Applied Mathematics, 128, 281-309, 2001.

[7] Tseng, Y.-h. and Ferziger, J. H.: A ghost-cell immersed boundary method for flow in complex geometry, Journal of computational physics, 192, 593-623, 2003.

[8] Smith, R., et al. "The Parallel Ocean Program (POP) Reference Manual Ocean Component of the Community Climate System Model (CCSM) and Community Earth System Model (CESM)." *Rep. LAUR-01853* (2010).

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