Geosci. Model Dev. Discuss.,
doi:10.5194/gmd-2016-135-AC2, 2016

Geoscientific Model Development Discussions

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 reviewer:
First of all, 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.

1 Response to general comments

This paper presents a new solver and preconditioner for the barotropic mode of the POP ocean model, focusing on parallel performance. The new CSI solver is slower than the previous CG solver in serial, but it removes a parallel reduce operation which makes it significantly faster at high number of processes. The new EVP preconditioner reduces the number of iterations required by the solver, thus reducing both computation and communication time. The authors demonstrate that the new P-CSI solver significantly improves the efficiency of the POP model in massively parallel runs. The paper extends an earlier proceedings paper (Hu et al. 2015) by presenting: a more detailed description of the barotropic solver used in POP, analysis of the eigenvalues and condition number of the associated linear operator, analysis of the convergence rates of different solvers, and new estimates of the computational complexity of the solvers. It is however debatable whether the additional material merits another publication as most of the material originates from Hu et al. (2015). In addition the analysis of the condition numbers could be improved.
[Response]:
Yes. This paper extends our previous 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. However, we also hope that our work can be widely understood and accepted by climate modelers so that we expand the paper and submit to GMD by a completed new revision. This includes a large change in the review of barotropic mode and the current solver which may cause severe bottleneck in the large-scale computing; a new theoretical analysis in section 4 (provide a robust base for the approach) and a different view of the new results in section 5. Although some figures look similar, the stories behind these figures are totally different from the SC paper if you closely compare

them. For the major changes, please see our detailed reply to the reviewer 1. We will clearly introduce these differences in the revised manuscript if we have the opportunity.

2 Response to specific comments

(1) Section 4.1 Spectrum and condition number
As we are dealing with a 2D shallow water solver, this analysis would be clearer if it incorporated the non-dimensional barotropic CFL number. It would be useful to know what the CFL number of the 2D mode is in typical CESM runs, and use that as a basis for the analysis and idealized experiments.
Noting that $CFL = c\Delta t/\Delta x$, with $c = \sqrt{gH}$ being the speed of the gravity waves, ϕ in eq. (14)(and consequently the condition number) can be expressed as a function of CFL. Specifically $\phi = 1/(CFL_x CFL_y)$ and (using $\lambda_{min}, \lambda_{max}$ from line 273) the condition number is approximately $k = 4CFL^2 + 1$ for large time steps and aspect ratio =1, which shows the dependency clearly.
Similarly in Figs 5, 6 and 7 it would be useful to use the CFL number instead of the time step or 2D velocity. The value of the time step alone, for example, is not informative as it depends on how the idealized run was set up.
[Response]:
Thanks for your good suggestions. We will use the CFL numbers in Figs 5, 6 and 7 rather than time step sizes and 2D velocities in our revised manuscript. The non-dimensional barotropic CFL number will be used as the x-axis in Fig. 5 and the legends in Fig. 6 and Fig. 7.

(2) Section 3.2 A block EVP preconditioner
In the last paragraph the authors mention that the drawback of the EVP preconditioner is that it cannot be used to solve large problems due to propagation of errors. Does

this imply that EVP cannot be used at low processor counts? If so have the authors experimented with or documented the failure of EVP? Is it possible to derive a threshold problem size under which the EVP preconditioner is reliable?
[Response]:
Sorry for the confusing. This problem only refers to the original EVP approach (described in Roache, 1995). The standard EVP solver is already modified so that the method can be used for any domain size and any processor count using domain decomposition (e.g., Dietrich, 1975) regardless of parallelization. So it has no problem at all for any domain size. However, if the domain size is too large without using domain decomposition, the computation will be very slow (See the complexity analysis 4.3 when p=1). Using parallel domain decomposition can actually help and speed up the EVP solver. We will revise the sentence in line 230 to avoid the confusing and add the following two reference papers [3][4].
D. Dietrich (1975) 'Optimized Block-Implicit Relaxation ', Journal of Computational Physics, Vol.18, No.4 421-439.
Roache, P. J. (1995) Elliptic marching methods and domain decomposition, vol. 5, CRC press, 1995.

(3) Section 4.3 Computational complexity
These estimates of computational complexity are similar to those presented in Hu et al. (2015). Some of them are different however (e.g. eqns. 26, 27, 28) especially in terms of the computation time T_c, Why the difference?
[Response]:
As you pointed out, there is a minor difference in the new manuscript. In order to make our analysis more general and understandable, we remove the assumption of equal partition along longitude and latitude directions. Therefore, the meaning of N is changed from partition number in both directions to the total grid number. The $O(\cdot)$ notation is also changed for the consistency. We also correct a minor error in which the computation time in P-CSI (with diagonal preconditioning) should be $T_c = 12N^2/P$

instead of $T_c = 13N^2/P$ to be precise.

(4) line 400: Fig. 8 shows that the P-CSI solver converges slower compared to PCG. Why is it so? The analysis in Section 4.2 concluded that the convergence rate should be similar.
[Response]:
We illustrated that the P-CSI has the same theoretical lower bound of convergence rate as PCG and ChronGear at page 11, line 324 when the estimation of eigenvalues is appropriate ($k^{'} = k$). In practice, as we also illustrated at page 13, line 400, P-CSI usually converge slower than PCG with the same preconditioning. The reason is that P-CSI requires that $0 < \upsilon \leq \lambda_i \leq \mu (i = 1, 2, ..., N)$, which means that $k^{'} = \mu/\upsilon \geq \lambda_{max}/\lambda_{min} = k$. According to Eq. (18) and Eq. (24), P-CSI has a slower convergence rate than PCG unless the estimation of eigenvalues is optimal. Furthermore, the theoretical bound is often too conservative for PCG. In practice, an increase in the convergence rate is often observed as the problem size increases, which is known as superlinear convergence of the PCG method [2]. To clarify, this explanation will be added into our revised version.

(5) line 415: It is not clear how the grid is divided in blocks. It might be worth explaining this is Section 2.1.
[Response]:
Thanks for your suggestions. POP divides the horizontal domain into small blocks evenly and distributes them to processes. We assume that there are N and M grids along longitude and latitude respectively, and the global domain is divided into n*m small blocks with size of N/n*M/m. These blocks are distributed to processors using simple Cartesian strategy or Space-filling Curve method [1].
We will add these materials into Section 2.1 to make the grid partition clear in our revised version.

3 Response to technical corrections

(1)line 109: ρ_0 is the constant reference density, not the actual water density.
[Response]: Thanks for your corrections. We will modify this error in our revised version.

(2) line 122: Meaning of the last sentence is unclear, please elaborate/reformulate.
[Response]: Thanks for your corrections. We will rewrite this sentence in our revised manuscript as follows.
"When we directly integrate the continuity equation from the bottom to the surface, we will get a form $\int_{-H}^{\eta} dz(\bigtriangledown \cdot u + \partial w/\partial z) = \partial\eta/\partial t + \bigtriangledown \cdot (H + \eta)U - q_w = 0$ under the surface boundary condition $w(\eta) = d\eta/dt - q_w = \partial\eta/\partial t + u(\eta) \cdot \bigtriangledown\eta - q_w$. The term including η inside the divergence leads to a nonlinear elliptic system, so many mature numerical methods such as conjugate gradient method cannot handle this problem. To avoid this, POP linearizes the barotropic continuity equation by modifying the boundary condition from $w(\eta) = d\eta/dt - q_w = \partial\eta/\partial t + u(\eta) \cdot \bigtriangledown\eta - q_w$ to $w(\eta) = \partial\eta/\partial t - q_w$ and dropping a small term. More details can refer to [1]."

(3) line 170: typo: scalar
[Response]: Thanks for your corrections. We will correct "scaler" with "scalar" in our revised manuscript.

(4) line 285: Here the authors assume that the time step satisfies the CFL condition, i.e. CFL$\leq$1, where the velocity v is chosen rather arbitrarily. It would be better to use CFL numbers typical to CESM applications.
[Response]: Thanks for your corrections. To avoid the influence of experiment setup, we will use three non-dimensional barotropic CFL numbers (0.01, 1 and 5) as the legend of Fig. 7. The CFL number in realistic 0.1 degree run of POP is about 3.46, so we choose 0.01-5 as the range of CFL number.

(5) fig 7: what is the grid aspect ratio used in this test?
[Response]: The grid aspect ratio used here is 1, and we will add this information to make our manuscript rigorous.

(6) line 335: T_c is not defined in this paper.
[Response]: The T_c represents the computation complexity, we will add this definition in our revised manuscript.

(7) line 513: typo: scalar
[Response]: Thanks for your corrections. We will replace "scaler" with "scalar" in our revised manuscript.

(8) figs. 16 and 17: these figures are not mentioned in the manuscript
[Response]: Thanks for your corrections. We will remove these two figures in our revised version.

4 References

[1] 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)."ÂăRep. LAUR-01853Âă141 (2010).
[2] Beckermann, Bernhard, and Arno BJ Kuijlaars. "Superlinear convergence of conjugate gradients."ÂăSIAM Journal on Numerical AnalysisÂă39.1 (2001): 300-329.
[3] D. Dietrich (1975) 'Optimized Block-Implicit Relaxation ', Journal of Computational Physics, Vol.18, No.4 421-439.
[4] Roache, P. J. (1995) Elliptic marching methods and domain decomposition, vol. 5,

CRC press, 1995.

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