RE: egusphere-2026-654
Title: Convolution Based Techniques for Computing Self Attraction and Loading in MOM6

Dr. Qiang Wang
Topic Editor
EGU Geoscientific Model Development

Dear Dr. Qiang Wang

Thank you for the reviews of our article referenced above. The article has been revised to address the comments of the Referees. The revised manuscript has been submitted. Here is a list of changes that we have made in response to the referee comments:

- The addition of section 2.6 in response to comments from Referee 2 and moving some of the content which was previously in section 4.2 to this section
- A number of citations have been added in response to referee comments; following Referee 1, special attention was paid to earlier literature in the geodesy community about using the convolution method for computing SAL
- Gibbs phenomenon and aliasing are now discussed in greater detail
- Cesàro summation is now discussed more carefully throughout the paper
- The CSFMM is introduced with more details
- A number of wording changes have been made in response to referee comments
- Some notational changes have been made to improve clarity

I have uploaded an updated draft of the paper, as well as an latex diff copy highlighting the changes between the updated draft of the paper and the previous version. In the following pages, detailed responses to the referee comments are provided.

Sincerely,
Anthony Chen, on behalf of co-authors:
He Wang
Robert Krasny
Brian Arbic

Referee 1:

—----------

It seems that most global ocean modelers who incorporate tidal forcing into their models now finally appreciate the necessity of implementing self-attraction and loading (SAL). But in light of HOW they are doing it, as described in this paper, it seems necessary that changes be made. This paper describes why and how. The paper will likely become influential in the modeling community. It thus deserves to be published.... eventually.

The second half of the paper is what will be convincing to modelers and justifies publication. The problem is with the first half of the paper which does not use, or even acknowledge, a large already existing literature on this topic. The first half needs a good bit of revision, to acknowledge the existence of that prior work, and also the authors should examine whether that prior work necessitates some changes to their own approaches. This prior work is mostly (maybe entirely) in geodetic journals, and it often addresses loading or gravity but not self-attraction, but the math is essentially the same.

—----------

We thank Referee 1 for acknowledging the main goal of the paper. We also thank them for their helpful comments on the paper. Following Referee 1’s suggestion, we have added references to and discussion of earlier work by the geodesy community on the use of convolutions for computing SAL. We also describe a key difference between this earlier work and ours–whereas their work tended to focus on computations of the SAL term from a small number of snapshots, here we are incorporating SAL inline in a model as it runs. This means that in our study the convolution SAL is computed at every model timestep, and hence computational efficiency and accuracy are both of prime importance. The requirement for computational efficiency and accuracy necessitates that some approximations be made, and error terms examined, as we discuss in the paper.

—----------

For example, the problem with Gibbs ringing is readily appreciated in the geodesy community, and for this and other reasons, nearly all geodetic loading computations use a convolution (Green's function) approach. Comparisons of spherical harmonic and convolution calculations, as presented here, have been done long ago. For example, a useful paper, not cited here, is:

Schrama, "Three algorithms for the computation of tidal loading and their numerical accuracy", J. Geodesy, 2005. doi: 10.1007/s00190-005-0436-3

Schrama actually does also address self-attraction.

There are popular, widely used, software packages that include many of the computations laid out here. I will mention:

- Duncan Agnew's SPOTL package

- Hilary Martens's LoadDef package

(both easily found via google).

These packages include things like computing Green's functions from loading Love numbers in numerically accurate ways. See especially Martens use of Kummer's series transformation. They (and other papers) discuss integrating over the singularity (or near-singularity) of the Green's functions and related matters -- I am pretty sure it does NOT involve skipping the integration at the origin as is done here (authors' statement after Eq (20)). Even Farrell's long 1972 paper does discuss some of this, too.

Many of the geodetic papers are also concerned with how to integrate over the ocean model near coastlines, which are not well mapped, but at least that problem does not arise here! (The border of the model ocean is known.)

In summary, I think the authors should start with Schrama's paper and at least read through the integration discussion in Agnew's manual. Martens's works are also useful and comprehensive. Then the authors can modify (if needed) some of their approaches, while acknowledging this previous literature. At that point, they can warn fellow ocean modelers of the need to account properly for SAL with the examples they give.

—----------

We thank Referee 1 for the suggested references, which we have added to the paper. We have read through these works and do not believe that modifications to our approach are needed.

The reason that we are using the form of the Green's function that we have chosen, based on asymptotic approximations of the Love numbers, is primarily for reasons of computational feasibility. In serial, performing a single SAL acceleration convolution calculation with the closed form gradient Green's function using the Cubed Sphere Fast Multipole Method (CSFMM) takes approximately 1.88 seconds. Using a truncated Legendre series of the Green's function is significantly slower. Summing 10 terms of the Green's function for an equivalent calculation takes 5.54 seconds, and summing 40 terms takes 33.6 seconds. While this may be feasible in an "offline" setting where data is being post processed, such an additional computational expense would not be feasible for use during a model simulation, in which the SAL term is being computed at every time step. This discussion is now included in section 2.4.

As for the singularity, there are many methods for handling this. We have chosen to skip the singularity, as it is consistent with treating the integral using a Cauchy principal value around the singularity. To understand this, first note that the SAL potential Green's function has an even singularity. The gradient of an even function is then an odd function, and the Cauchy principal value of an odd function around a singularity is 0. Of course, the model does not quite match the mathematics exactly, but the amount of error due to skipping the singularity is small because of this. This discussion is now included in section 2.5.

—----------

A few other items....

In any resubmission, I would appreciate having line numbers.

—----------

We have now done this. Thank you for the suggestion.

—----------

I don't understand the desire to break the symmetry of the Green's functions by including the gradient in it. The gradient can be done later, since it has to be done anyway for the grad \eta term. Giving up the symmetry in G is a high price to pay when one has to evaluate gradients later anyway. The SAL field is a smooth function (even over the land/ocean boundary) and easily differentiated, and I see no reason not simply to use that field and keep G symmetric. Maybe I'm missing something.

—----------

Breaking the symmetry of the Green’s function by taking the gradient doesn’t incur any additional computational cost in our situation as the CSFMM handles the nonsymmetric kernel without any issues. In addition, because of the use of the CSFMM, taking the gradient is important. The CSFMM uses barycentric Lagrange interpolation for the approximations that reduce the computational cost, but as a result, the output field from the CSFMM has small scale polynomial oscillations. Computing the SAL potential with the CSFMM and then taking the gradient would magnify these small oscillations, which would then adversely affect the solution. This has been added to Section 2.5.

—----------

Expressions (14) and the listed constant terms (a1, b0, b1) appear to me to be very poor approximations when compared with Farrell's Table A2 (notwithstanding different earth models). Farrell tabulates k_n'/n, and it is not a constant until n > 1000 or so.

—----------

We acknowledge that the low wavenumber behavior is likely not captured correctly as the approximations used only hold for high wavenumbers. However, this doesn’t seem to have a large impact on the solution. The reason for choosing this approximation is for reasons of computational efficiency, as previously discussed. As seen in Figures 2 through 5, the spatial structure of the SAL acceleration is quite similar with differences in magnitude in coastal regions. However, the large scale structure and amplitude in the open ocean is quite similar. Additionally, as seen in Tables 1 and 2, the error behavior does not appear to be negatively impacted by these approximations either. In the future, further research into other forms of the Green’s function is possible and this may lead to additional improvements as well. This is now discussed in 2.4.

—----------

Figure 3. It's somewhat worrying to me that the degree-1 terms differ. These depend of the reference frame (and thus bear on geocenter motion) and I'm surprised the red and blue curves do not match there. It's worth checking for some reference-frame or other inconsistency here.

—----------

Thank you for the comment. The mismatch is likely due to the errors present in the approximation of the Load Love Numbers used in deriving the closed form Green’s function, which are largest at low wavenumbers. However, it is our understanding that the M2 tide shouldn’t have a wavenumber 1 signal, and as such, any wavenumber 1 behavior would be a numerical artifact. If the referee has any additional comments, we would be interested in hearing them.

—----------
I was actually very surprised to read that the "standard" procedure is to use spherical harmonics at N=40, which is surely far too low. I had assumed that packages that use fast spherical harmonics were being used (e.g., spherepack3), which (e.g.) for a 0.5° grid would automatically take an expansion to N=360.
—----------
We note that the use of inline SAL in ocean tide models is relatively new, and a standard number of spherical harmonics to use has not been established. However, several prior works (Barton et al. 2022, Brus et al. 2023, Wang et al. 2024) on the use of inline SAL have chosen N=40 as a suitable choice.

While using up to degree N=40 spherical harmonics does not capture the full amount of information present in the sea surface height field, at the same time, increasing N past this value offers only very small improvements to the accuracy of the computed tides, as seen in Tables 1 and 2. In the open ocean, variations are large scale and correspond to low degree spherical harmonics, so N=40 suffices to capture the SAL accurately in the open ocean. Near coastlines, Gibbs errors mean that increasing the number of spherical harmonic components used does not improve the accuracy of the SAL acceleration.

Additionally, MOM6 (Wang et al. 2024) does not use a fast spherical harmonics library. Instead, it directly computes the spherical harmonic components via direct integration. MPAS-Ocean (Barton et al. 2022) previously used a fast spherical harmonics package, SHTns, and found similar results with regards to error saturation at N=40. The choice of N=40 was also partially motivated by runtime considerations as well. This is discussed in Section 2.1.
—----------
I agree that the problem with the "standard" SAL approach for ocean modeling is probably with the anomalous gradients near land. It might just be due to the fact that this is where the major dissipation from bottom friction occurs and so must be done accurately. Perhaps that is more or less equivalent to the "back effect".
—----------
Yes, we agree with this comment from the referee. Further investigation as to the reason for the improvement in the tidal error is possible in the future as well, as mentioned in Section 4.2.
—----------
Page 2, line 9: This reads as if my 1998 paper advocates using a scalar approximation, when in fact it does the opposite. I'd appreciate it if you move the (Ray,1998) reference to the end of the following sentence [after "different scales"].
—----------
Thank you for the comment. We have made this change.
—----------
In the final paragraph, I think the reference to earth's free oscillations is irrelevant to the present problem. The speed and frequency of free oscillations have nothing (directly) to do with the earth's response to tidal loading. The only anomaly like this that might eventually come into play is from an anelastic earth response, and new GNSS analyses are starting to bound this; e.g.,

perhaps at the mm level (in amplitude, so far not delay) in vertical deformation, according to Bos et al (2015; http://dx.doi.org/10.1002/2015JB011884)

—----------

Thank you for this comment. We have made changes to this section. The sentence in question has been removed.

Referee 2:

—----------

This paper is concerned with the representation of ocean self-attraction and loading (SAL) in global ocean models. In tide-resolving ocean models, the effect of the SAL is represented as a perturbation to the equilibrium tidal potential. However, this perturbation depends on the time-dependent distribution of ocean mass (which can be expressed in terms of the sea level anomaly), so the SAL cannot simply be computed once as an external forcing field and incorporated into the forcing. Instead, the sea level anomaly computed by the ocean model must be incorporated into the SAL forcing as the model is running.

The effect of a localized infinitesimal sea level anomaly on the surrounding gravitational potential at the surface of the Earth may be expressed by a convolution of the sea level anomaly field with a Green's function. Prior to this work, efforts to incorporate the SAL in global ocean models either took the SAL as a prescribed forcing computed from independent data, or they approximated the convolution as a truncated spherical harmonic expansion. The new element of this work is the approximation of the convolution directly on the model's computational grid using a quadrature based on the Cubed Sphere Fast Multipole Method. This seems like an important effort to accurately and efficiently compute the SAL from first principles in ocean models. The authors implemented their approach in MOM6, and they evaluate the success of their efforts by comparing the output of the simulation with independent estimates of the tidal sea level anomaly, based on a TPXO solution.

—----------

We thank Referee 2 for acknowledging the main goal of the paper. We also thank them for their helpful comments on our work. Following Referee 2’s suggestion, we now discuss the sources of error in the different SAL computation methods using single snapshots of sea surface height, before we delve into the errors in model runs, which as the Referee points out are affected by other aspects of the model

—----------

Unfortunately, I think this approach conflates several possible sources of error and makes it impossible to evaluate the quality or significance of their work. Since their focus is on accurately and efficiently computing the convolution of a function with a Green's function, they should simply evaluate this in isolation before using an ocean model as the basis for comparison. Of course, they ought to use the actual model grid and realistic test fields for \eta in their comparisons. Likewise, since it is essential to their application, they should evaluate the accuracy of the gradients of the SAL (the components of acceleration), as they have done nicely here.

But this evaluation should be done in isolation from the other model numerics which irreversibly fold in (possibly compensatory) errors due to the following:

- time-stepping stability and accuracy (What time-level scheme is used for SAL?)
- errors in the frictional parameterizations (wave drag and bottom friction)
- errors in coastline location and bottom topography
- errors due to the neglect of baroclinicity
- errors in the TPXO comparison data

—----------

We agree that the error values in Tables 1 and 2 contain error from a number of sources. Prompted by Referee 2, in our revised manuscript, before delving into the model runs with inline SAL, we compare the computed SAL gradient using the different methods in Section 2.6, which covers Figures 2 through 5. Figures 2 and 3 indicate broad large-scale agreement between the spherical harmonic method and the convolution approach, while Figures 4 and 5 shows some improvements in the computed SAL acceleration in coastal regions when using the convolution approach owing to the removal of Gibbs errors. More detailed discussions on the accuracy and efficiency of spherical convolutions are givent in Chen and Krasny (2025).

Additionally, directly evaluating the accuracy of the computed SAL field is challenging. The SAL potential or SAL acceleration cannot be observed or measured, so error comparisons have no ground truth to compare to. This is why we have chosen to evaluate the SAL calculation in the context of a tidal simulation, following the approach of Barton et al. (2022), Brus et al. (2023), and Wang et al. (2024). There are many papers in the tidal literature which keep other things the same and which improve one particular aspect (SAL computation, bathymetry, wave drag, etc.) while holding other aspects fixed; ours is no different from these previous efforts.

—----------

For these reasons I suggest "reject" for this manuscript. Nonetheless, I think this is an important topic, and the authors are capable and careful scientists, so I hope they will revisit their approach and publish on this topic in the future.

l11: antecedent of "which" is not clear

l12 and throughout: avoid parens within parens

l19: It seems like you mean "3-dimensional" or "depth-resolving" model, rather than "baroclinic model".

l20: should be plural "configurations", unless the 3 authors all used the same configuration

l23: is the phrase, "large scale", in the right place?

l24: mesoscale eddies aren't necessarily smaller scale than the internal tides

l39: Be more specific: "the computation of a convolution" → "the convolution of the sea surface height"

l60: "tidal forcing" --> "astronomical tidal forcing"

l73: omit "in line" or move it closer to the word it modifies, "computing"

l104: omit comma before "that"

l143: Capital "G" in the derivative? Or define exactly what is the lower-case "g_{SAL}" function.

l166: What is the "error computation"? What is compared to what?

—----------

We thank the Referee for the various suggestions. We have made these changes.

—----------

l76: I have not done these computations myself, but don't the Love numbers decay like 1/n? I don't find the phrase, "modified Load Love Numbers", used in Ferrell (1972). It appears that he compares the Greens function integrals with spherical harmonic expansions. Why aren't you referring to his work in more detail? I doubt I am expert enough in this subject matter to review it adequately. Please distinguish between the "primed" and "bare" Love numbers.

—----------

Based on the Load Love numbers in Wang et al. (2012), the h_n do not decay like 1/n. The modified Load Love Numbers are mentioned by Farrell (1972) as the mass-load Love Numbers shortly after equation 32, based on the convention of Munk and MacDonald (1960). The Love numbers without the prime are the body tide Love numbers, while the Love numbers with the prime are the load tide Love numbers.

—----------

l86: How is the computation of an integral "prone to aliasing"? What do you mean?

—----------

Aliasing occurs when high wavenumber signals in the underlying data are projected to low wavenumber signals after performing a Fourier transform. In the case of spherical harmonics, the discretized grid can cause aliasing due to alignment issues with the spherical harmonics being used. This is now discussed in Section 2.1.

—----------

l88: Doesn't \eta go smoothly to zero at the shoreline? I don't understand where is the discontinuity that would give rise to the Gibbs phenomenon.

l102: The sea level "jump" occurs over the scale of a grid cell, so it is hard to imagine scenarios in which it could be significant. And, doesn't it involve a minuscule mass of ocean, compared to the deep water, so its effects in the SAL should be negligible?

—----------

The discontinuity is present in the fact that while the real sea surface height goes smoothly to 0 at coastlines, the model grid is not fine enough to resolve this. Additionally, the tides are largest near coastlines. As a result, along a coastline, the first land point has a SSH of 0, while the first ocean point has a large nonzero SSH.

The globally averaged magnitude of the M2 primary lunar semidiurnal tide is about 27 centimeters, but in some coastal regions, this magnitude can be up to several meters. In the open ocean, the magnitude of the tide is much smaller. Thus, while the mass of the ocean in coastal regions is small, the amplitude of the tide is large and this ends up having an impact. Additionally, the gradient taken to convert the SAL potential into the SAL acceleration magnifies small scales. This can be seen in Figure 4 and 5.

—----------

l113: Maybe it would be useful to show some graphical examples of the Gibbs phenomena and it amelioration for a 1-dimensional example. I assume you implemeted Cesaro summation? If not, then omit this section.

—----------

Yes, we did implement Cesàro summation, with error results presented in Tables 1 and 2. We chose not to include a graphical example of Gibbs phenomenon for the 1 dimensional case as it is fairly standard and available widely online. We did include an illustration of Gibbs phenomena in the SAL case in Figures 4 and 5.

—----------

l125: Shouldn't you go back to Ferrell 1972 here and simply start with the Green's function? Since the function \eta is originally defined in physical space, it seems odd to develop these expressions from the convolution theorem and its spherical harmonics.

l133 and l137: It seems like you should go back to Ferrell 1972 and either use his Green's function and/or explain why your approximation in equation (14) is justified.

—----------

We have made changes in this section.

While the Green’s function proposed in Farrell (1972) is the exact Green’s function, it is not suitable for computational use as it is an infinite series. Farrell discusses various approaches for convergence acceleration and then truncates the series at a high wavenumber, and similar approaches are taken in Schrama (2005), Agnew (2012), and Martens et al. (2019). However, none of these works perform SAL computations in an ocean model, in which SAL must be computed at every timestep, and the runtime impact of using a truncated series was quite severe. In serial, performing a single SAL acceleration convolution calculation with the closed form gradient Green’s function using the CSFMM takes approximately 1.88 seconds. Using a truncated Legendre series of the Green’s function is significantly slower. Summing 10 terms of the Green’s function for an equivalent calculation takes 5.54 seconds, and summing 40 terms takes 33.6 seconds. While this may be feasible in an “offline” setting where data is being post processed, such an additional computational expense would not be feasible for use during a model simulation. Additionally, as seen in Tables 1 and 2, the approximations used in the derivation of the Green’s function appear to impact the error by only slight amounts.

—----------

l153: How does omitting i=j "remedy" anything? Don't you need to express the function as a sum of a singular kernel and a remainder, and then analytically approximate or evaluate the singular part? See Green and Nycander for an example.

—----------

We have chosen to skip the singularity, as this procedure is consistent with treating the integral using a Cauchy principal value around the singularity. To understand this, first note that the SAL potential Green’s function has an even singularity. The gradient of an even function is then an odd function, and the Cauchy principal value of an odd function around a singularity is 0. Of course, the model does not quite match the mathematics exactly, but the amount of error due to skipping the singularity is small because of this. This is now discussed in section 2.5.

—----------

l165: Since you are concerned with the accuracy of an approximation, the representation of a convolution over a spherical shell with its truncated spherical harmonic expansion (and with its approximation over the spatial grid), I think you should simply examine this directly for some candidate \eta fields. It seems like the implementation in MOM6 should come later, once you are convinced that the approximations are accurate enough.

l178: While your MOM6 application might benefit from comparison with TPXO9 tides, I think you should separately study the accuracy of the approximations of the convolution. There are many other sources of error in MOM6 and in TPXO9 which could obfiscate or cloud the interpretation of your approach to computing the SAL.

—----------

These comments are similar to comments previously discussed above. We have added a more detailed comparison of the various SAL calculations in Section 2.6, highlighting the various differences in the calculations, and how the use of the convolution removes various oscillatory errors coming from Gibbs phenomenon, aliasing, and finite differencing.