

Referee 1

This manuscript presents an updated version of the existing one-layer melt rate model LADDIE. The new model version is not just incremental; it includes translation to a new code language, a new mesh structure, parallelisation and full integration with an ice sheet model, making this a significant upgrade to the older version. Differences in the new model version are described thoroughly in a well-referenced and easy-to-follow section of the paper. The model is then assessed against LADDIE v1, multi-model means from ice-ocean coupled ensembles and satellite-derived melt rates in a detailed analysis, showing its strengths and honestly discussing potential drawbacks. Results are then presented of a coupled simulation using the MISOMIP+ protocol, compared to using a parameterised melt. The discussion section raises a few important points which I was pleased to see highlighted.

I enjoyed reading and learning about this work, and am very interested to see future applications of this new model version. I find the manuscript to be exceptionally well-written and contain a good level of detail in its presentation. I found very few instances where I felt the need to suggest improvement. This manuscript is certainly suitable for publication in GMD, and I recommend publication after a few minor revisions, detailed below.

We would like to thank the anonymous referee for their very positive review. We agree with the points raised and give a detailed reply below in blue.

Fig. 3: I don't find the blue/grey colourscale in panel (c) to be clear unless I zoom in on the image. Perhaps using the same colourmap as panel (f), or something similar, would be a better choice for easy visibility of the differences. It may also be beneficial to show the velocity vectors as well as the speed distribution.

We have indeed changed the colourmap of panel c to the same Red/Blue one of panel f. We also like the suggestion to include velocity vectors and have added those.

Line 242-3: It looks as if doubling from 16 to 32 does not make much difference, and to get the ~33% time reduction you would need to quadruple to 64. Is this common in other tests than the one shown in Fig.4?

This is indeed the case for this example, and we have seen similar behaviour for other examples. We have highlighted this explicitly (l. 245).

Line 245-6: Related to the above, 32 or 64 cores may produce the most rapid simulations, but looking at the results in Fig.4, I would question whether increasing the number of CPUs beyond 16 is justified by the fairly minimal gains in time. What would your thoughts be on the optimal number of cores for efficient use of computing resources?

This is a fair point, and we fully agree that in this example, the use of 64 cores is difficult to justify compared to 16 cores. We will mention explicitly that for this example, 16 cores would indeed be optimal, though this may differ between use cases (l. 242).

Line 333: While LADDIE has higher *mean* melt rates, it has a much lower *median* melt rate, which I think should be addressed in this section. In Fig.6 it is clear that LADDIE produces close to zero melt

over large areas of the ice shelves, where the multi-model mean from RISE has a baseline $\sim 1\text{m/yr}$ melt across many of these areas. This is particularly noticeable when comparing the melt rates of Amery ice shelf, but is a common occurrence. In many places the satellite estimates in Fig.7 also show this (although notably not on Amery, where estimates actually show more refreezing) Why does LADDIE not replicate this low, but non-zero, melting?

This is a good point, and also visible in Fig 5. We have added this paragraph to this section to highlight this discrepancy (l. 313 onwards):

We also find a large difference in median melt rates between the RISE ensemble and LADDIE. This median value reflects the typical melt rates in the relatively flat and shallow regions of ice shelves. A notable example is the Amery ice shelf (# 29), where 3D ocean models simulate melting on the order of 1 m/yr , whereas LADDIE simulates near-zero melt rates. A similar discrepancy is found in the idealised ISOMIP+ configuration (Fig. 5). We attribute this difference to underestimated temporal variability in LADDIE, which converges to a nearly static steady state, whereas ocean models retain a considerable temporal variability in ocean speeds. This variability can induce a net transport of heat towards otherwise cold shallow regions, along with non-zero friction velocities. A specific example of how variability causes enhanced melting is 'mode 3' melting, induced by seasonal intrusions of warm water into the shallow regions of the cavity, which is not represented by LADDIE. Because these regions typically contribute little to buttressing, we consider this discrepancy of lesser importance.

Referee 2

General comments

This paper, like the LADDIE model itself, aims to fill the gap between, on the one side, computationally expensive physics based 3D ocean circulation models and, on the other side, empirical and machine learning parameterisations and plume or overturning models. LADDIE V2.0 is a physics based 2D circulation model for ice shelf cavities, much cheaper to integrate than 3D ocean models (hence can be run at a finer resolution) and incorporating a more physically complete circulation than the simpler approaches. This is a glaring gap in the current generation of ice sheet models used for future projection (ISMIP6/7 models). LADDIE V2.0, and this paper, therefore represents an important contribution to the modelling toolkit and the literature.

My perspective: I am primarily an ice sheet modeller, not an ocean physics expert. So my review considers the demonstrated capabilities of LADDIE and its application to ice sheet modelling rather than the derivation of governing equations and the oceanographic assumptions that have gone into these.

The layout is very clear. The model is well described. An informative selection of comparisons are presented. The comparison between LADDIE and the quadratic melt rate in the coupled setup is very powerful.

We thank the referee, Rupert Gladstone, for this positive review. Please find a point-by-point reply to all comments.

I have only minor comments to provide.

I like the LADDIE logo, but I personally don't like it being included in the figures. It feels too much like advertising. Same for the UPSY and UFEMISM logos.

This is a very fair point and we have removed the logos from figures 1 and 2.

In principle, a paper should provide sufficient information for a reader to repeat the experiments. But sometimes there isn't enough specific information, especially in section 4 (see also my specific comments).

We have added additional detailed configuration including the Stokes approximation and the sliding law, and pointed out the complete configuration file available in the data repository (l. 403).

I've cloned the code out of curiosity. The LADDIE model itself is a very small part of the total repository, both in terms of file count and disk space. This isn't perfect, and the file counts of the repo are quite large (some HPC machines are very strict about this), but it is probably small enough not to cause problems.

This is indeed the result of a pragmatic choice, after several attempts with submodules, to minimise the overhead of model maintenance for our tiny group. Luckily we haven't received any complaints yet from other users in this regard.

I've scanned the code from a coupling perspective, assuming that the authors wish for LADDIE to be coupled to models other than UFEMISM in the future. The mesh information and partition information seems to be very well and clearly stored. The initialise and run routines look to be compatible with being called through a different model, though a wrapper to group the multiple initialisation routines might be needed for some couplers. For some couplers the MPI communicator might need to be passed to the submodels. This would need the "MPI_COMM_WORLD" instances changes to be able to use a communicator passed to LADDIE, which might not be the global MPI context. In short, some minor changes might be needed for coupling, but the code structures are mostly well designed for coupling.

This is very useful feedback, and we are glad to hear that the basic model structure is suited for coupling. We have added a note that we foresee minor code changes to be necessary for wider applications, and that we will attempt to implement these upon request from various groups (l. 487).

I couldn't see a description of which Stokes approximation was used in UFEMISM. My understanding is that several are implemented and DIVA is commonly used. Please clarify that in the text.

This is correct and should indeed be stated in the text. We have included these details (l. 403).

Shameless self promotion: There's a new study led by Qin Zhou that applies very fine resolution 3D ocean modelling to ice shelf channels under Fimbul. Could be relevant to your introduction and around pages 17-19.

"Channelized Topography amplifies melt-sensitivity of cold Antarctic ice shelves", Nature Comms, May 2026.

Thank you for pointing out this relevant paper. We have cited this in the intro (l. 28)

Specific comments

I didn't find an explanation for what the parameter A_h means in equations 2 and 3. Or K_h .

Indeed, these were missing from the table, we have included these in Table 1.

Figure 1 caption. Technically, triangles have zero thickness. It is probably better to say "triangular prism" than "triangle with finite thickness".

This is indeed correct and we have adapted it.

Line 106 "melting of" -> "melting or"

Agreed, and corrected.

Is there supposed to be an equation 16? It's showing blank for me.

Indeed, equations 16 and 26 are empty and shouldn't be there. We have removed those.

Page 8. I got quite confused by the notation here, until I realised that "n" has rather carelessly been allowed to wander from timestep counting to substep counting. Please make sure not to mention n in the phrases "the weights β_n " and "for substep n" at the start and end of the first paragraph following the equations. Please ensure that your terminology for counting timesteps is separate from the terminology for counting substeps.

Indeed there is a duplicate meaning of n here, and we have replaced β_n with β_k . In fact, substep n (or $n+0$) should be a valid notation to indicate when diffusive terms are computed, but to avoid ambiguity, we rephrased this to 'full timesteps n'.

Using stars for H is a viable way to indicate substeps, but then it becomes confusing that beta uses a different notation. H^{**} and β_2 both imply that we're in the second substep. Why not use stars for both? e.g. β^{**} for the second substep?

Beta represents weights of substeps to determine interpolated values of H between these substeps. Hence, stars should be avoided for beta values, and we use subscripts for beta to avoid confusion with the substeps (superscripts). Overall, we retain the notations of Lilly et al (2023), and keep stars for interpolation between substeps and $n, n+1/3$, etc for the actual substeps. β_k is then used as weights of the various substeps.

Figure 2. The figure seems to imply that the mesh is updated less frequently than every coupling interval, but the text states "At each coupling time step, the mesh is provided by UPSY". This looks contradictory to me.

This is indeed confusing and we removed the statement that the mesh is provided each time step (l. 411).

A note on accelerated forcing. It's not quite as simple as implied at the end of page 8, especially when rates of change are either directly passed between models or implied by the boundary conditions that are directly passed between models. I think you need to be clear about whether you intend the ocean model to equilibrate each timestep (which is kind of implied in your text but you don't talk about demonstrating this) or accelerate the rates of change.

Thank you for pointing out our misinterpretation of the accelerated coupling. We indeed use the rapid equilibration and assume an approximate equilibration during each coupling interval. We removed the comparison to the accelerated forcing to avoid confusion (l. 190).

If you pass temperature from ice to ocean, you're implying a rate of heat transfer. If your ocean model runs for a fraction of the full coupling interval, and you've not "accelerated" this heat transfer, you might get a different result compared to running the ocean model for the full coupling interval. Is your intention here that the ocean model runs long enough to equilibrate to the updated temperature, such that the rate of heat transfer doesn't matter?

Indeed, as stated above, we use equilibration rather than acceleration, which also applies to the heat transfer. We removed the reference to accelerated forcing to avoid confusion (l. 190).

Line 259 typo "het salt"

Fixed

Line 291 typo "Ice-shelf/ocen"

Fixed

It is very notable that the LADDIE median and mean melt rates are below and above the RISE range respectively. This warrants at least a sentence or two. I guess it is due to a spatial distribution with small regions of very high melt that lift the mean? Perhaps related to finer resolution than the RISE models?

This was also raised by the other referee. We included the following paragraph to explain the difference in median melt rates (after the already included paragraph devoted to mean melt rates) from l. 313 onwards:

We also find a large difference in median melt rates between the RISE ensemble and LADDIE. This median value reflects the typical melt rates in the relatively flat and shallow regions of ice shelves. A notable example is the Amery ice shelf (# 29), where 3D ocean models simulate melting on the order of 1 m/y, whereas LADDIE simulates near-zero melt rates. A similar discrepancy is found in the idealised ISOMIP+ configuration (Fig. 5). We attribute this difference to underestimated temporal variability in LADDIE, which converges to a nearly static steady state, whereas ocean models retain a considerable temporal variability in ocean speeds. This variability can induce a net transport of heat towards otherwise cold shallow regions, along with non-zero friction velocities. A specific example of how variability causes enhanced melting is 'mode 3' melting, induced by seasonal intrusions of warm water

into the shallow regions of the cavity, which is not represented by LADDIE. Because these regions typically contribute little to buttressing, we consider this discrepancy of lesser importance.

Looking at figures 5 and 6 LADDIE seems more ready than other models to drop melt rates toward zero far from grounding lines. Why is this? Who is "correct" here?

See paragraph above, we consider this a limitation of LADDIE, as the non-zero melt rates are also found in remote sensing products, and a physical argumentation can be provided why these melt rates should not be zero.

Line 339. Do you compare to the Davison dataset because you have more confidence in that data set? Or because it gives a closer match to the LADDIE melt rates? And yes, this question is intentionally tongue in cheek... but readers will be wondering why, if you don't give a reason for choosing that data set for your visual comparison. After all, the comparison to models was a multi-model mean, so why not a multi-data mean?

In fact, we compare to Paolo, which shows the worst agreement with LADDIE, rather than Davison. Davison and Rignot have no publicly available gridded datasets, and Paolo is considered to be an update of Adusumilli. We explain this choice and to avoid the impression of cherry picking, note that for this, Paolo would be the worst choice (l. 350).

Line 380. Would you like to say something about why satellite estimates don't work near the GL? Is it because they have to make the hydrostatic assumption? Or are there other reasons?

That is indeed the main reason, which we have added explicitly (l. 395).

Page 19/Figure 8. There's a very elegant pattern of curved high melt regions on the west side of the main trunk of the PIG ice shelf in LADDIE. I don't see this in the satellite estimates. Can you say anything about this? Perhaps at least say why it occurs in LADDIE, even if you don't know why it doesn't occur in the satellite estimates? Does it follow a pattern of curved channels in the geometry you use?

I think this refers to the southern shear margin (right hand side), which is heavily damaged. This breaks the assumption of continuity made in the processing of satellite data. We highlighted that these melt patterns by LADDIE are guided by basal topography, though a comparison to the satellite estimates cannot be made in these regions (l. 388).

Line 389. Do you mean applying an idealised ocean forcing to LADDIE or directly to a melt parameterisation?

The same idealised ocean forcing is applied to LADDIE and to the quadratic parameterisation. We have stated this more clearly (l. 407).

Line 390. I know you provide a reference, but could you say just slightly more about WARM?

Yes, we described this forcing profile in more detail here, which has in the meantime been renamed 'high-melt' (l. 410).

Line 390. What does the 300 years refer to? After the spinup? In which case how long is the spinup?

We added that the spinup without sub-shelf melting is done for 20,000 years (l. 405) and clarified that the 300 year simulation is the one with applied ocean forcing (l. 410)

Line 392. Can you quantify "significantly"?

We have added that this depends on the user-defined configuration and an internally computed 'mesh fitness coefficient' (l. 412)

Line 418. Typo "dependend"

Fixed

Line 418. Can you speculate on what might be different between the ice sheet models? Is it the Stokes approximation? The sliding law? Resolution?

We left this deliberately vague, as we have not figured this difference out yet, and consider this to be beyond the scope of this LADDIE model description manuscript. We prefer to avoid such speculation and leave this question to be resolved by a proper in-depth study.

Line 427. Could also be due to using an unstructured mesh instead of having a row of square grid cells perfectly aligned with the GL.

Yes this is a good point which we have included (l. 450)

Line 461. Typo? Not sure what "at" is doing in "repository at to optimise"

Fixed