

Detailed responses to reviewer 1 (reviewer comments are included in black, responses in blue font)

General comments

The authors present a nested regional ocean model of Halifax Harbour and part of the Scotian shelf which is validated against measurements. A simple dissolution model is implemented and pulse releases of an alkaline effluent are modelled, consisting of a mix of dissolved and particulate alkalinity. The subsequent changes in alkalinity and DIC (from the induced CO₂ uptake) are evaluated and analyzed.

Overall the manuscript is well laid out, focused and easy to follow. The simulations presented establish an important standard of rigor for future OAE deployments in other areas. I recommend publication.

Response: We appreciate the positive assessment and the constructive feedback. We have addressed all the comments as described in the detailed responses below.

Specific comments

Comment:

1. The authors show that alkalinity addition inside a natural enclosed harbour enables a substantial fraction of the theoretically maximal CO₂ uptake to occur quickly and within the simulation domain, due to the long residence time and relatively shallow waters. As pointed out in L556-559, this makes MRV much easier both experimentally and from a simulation perspective. Of course the flipside of this is that a confined body of water which does not quickly spread any added ΔTA over large ocean areas will also limit the total sustained alkalinity addition rate in that area, limiting scaling of OAE.

It would be useful to add an estimation of this in the manuscript. For a rough, first pass estimate, perhaps one could assume that the response of ΔTA and ΔDIC are roughly additive and linear with respect to addition rate. Then, for each of the three locations, one could calculate what the maximum addition rate would be which would raise the maximal ΔpH to some acceptable limit (what that limit is is of course arbitrary, but perhaps something conservative like +0.1 or +0.05 units would be illustrative).

Another approach would be perhaps to examine the export rate of alkalinity out of the simulation boundary and try to estimate what sustained alkalinity addition rate (rather than a pulse) could be achieved, again within some ΔpH or ΔTA limit set within the domain.

A discussion of this and the trade-offs of release locations would be useful to the reader to understand better what sort of scale OAE can achieve.

Response: These are interesting and highly relevant comments. Regarding the trade-off between measurability of the alkalinity signal in the harbour (which is helped by the high residence time of the system) and the risk of breaching regulatory and environmental

thresholds (which is elevated in this system because of the long residence time) we would like to refer to Wang et al. (2025) where this was investigated using an “exposure” metric and many simulations with release from Tufts Cove and Mill Cove of different feedstocks at various dosing rates. Wang et al (2025) used a similar physical setup with a particle dissolution model. They found that exposure rate is much higher when feedstock is released from Mill Cove, especially in summer (higher residence time) for slow dissolving/fast sinking particulate feedstock. In their experiments the addition of slow sinking, slow dissolving particles at Tufts Cove resulted in lowest exposure.

We did discuss Wang et al. (2025) results regarding exposure risks in Section 6.3 but will expand this discussion in the revised manuscript.

Regarding the implications for scalability, this is the subject of a comprehensive analysis in a forthcoming manuscript.

Comment:

2. L317 $k_{\text{diss}}TA_p$ term:

The treatment of dissolution as an exponential decay process (i.e. $dTA_p/dt = -k TA_p$) was surprising at first glance. Usually dissolution of particular matter is treated with a shrinking core model, where the dissolution rate has units of $\text{mol cm}^{-2} \text{s}^{-1}$, the radius of particles shrinks linearly and fully dissolves in a finite amount of time. For a very narrow (as indicated in L335, “a particle size of $12\mu\text{m}$ ”) or uniform distribution of particle sizes I believe an exponential dissolution curve is only a mediocre fit.

I can see that an exponential model could perhaps capture the behaviour of a gaussian or log-normal distribution of particle sizes, but a short discussion of this and a justification of the choice of model here would be helpful.

Response: In Section 5.1 of the revised manuscript, we will discuss the choice of an exponential decay for the dissolution of particles and in Section 5.2 we will provide more details on the choice of the dissolution parameter in the experiments with particulate feedstocks.

Comment:

3. L317 w_pTA_p term:

It’s unclear to me how the sinking term is applied. As written it looks like there is an exponential decay, i.e. each time step some fraction of TA_p is lost to sinking from any given simulation grid voxel. What happens to that TA_p ? Does it get added to the cell below, until the bottom cell is reached after which it disappears into the sediment? Or does the model assume the sunk particles are removed completely (i.e. they sink out entirely at a rate of w_pTA_p from anywhere in the column?). As currently written it seems more like it’s the latter, as there is no term that accounts for sinking particles that arrive from a cell above (I was expecting a second term like $+w_pTA_p^{z=i-1}$)

Please clarify how the sinking mechanism is implemented and justify its construction.

The sinking rate is stated as 5.5 m^{-1} later (L337) but that can't be w_p since the units wouldn't be right (w_p should have units of inverse time, like k_{diss}). How is w_p calculated from the 5.5 m^{-1} ?

Response: We will clarify this part in the revised manuscript. The confusion originates from a typo in the units of w_p , which should have been m s^{-1} , and the formulation of sinking in Eq. 8, which should have been $w_p \frac{\partial T_{Ap}}{\partial z}$.

Sinking occurs between the vertical layers of the model. As the reviewer mentioned, sinking is a source term to the layer below and a loss term from the layer above (except for the surface layer). In the bottom layer, particulate material accumulates and a fraction is lost through the incorporation into the sediment (Eq. 14, see response to comment below).

Comment:

4. L326 The treatment of sediment loss in layer N is a little unclear. It says a term is “added” to $\partial \Delta T_A / \partial t$? Or does this replace the regular dissolution term in $\partial \Delta T_A / \partial t$ (last term in Equation 9) ? It might be clearer here to just rewrite the full Equation 9 (and perhaps Equation 8) in the case of the bottom cell, for clarity.

It's also confusing to me that the loss of T_{Ap} due to sinking/burial is already explicitly treated in equation 8 using w_p and then it's treated again here with the θ_{loss} term. Is θ_{loss} a constant ? Or is it calculated from w_p ?

Response: We will update Eqs 8 and 9 for clarity. Sinking and burial are two different processes (see also response above). In the bottom layer, sinking material is immediately resuspended and therefore accumulates. Instead of a sinking loss term there, we assume that a constant fraction of T_{Ap} present in the bottom layer and thus in contact with the sediment is incorporated into the sediment through bioturbation and thus lost from the system.

Comment:

5. L424 The comparison of H2 and H3 is very interesting and suggests perhaps a resolution as high as H3 isn't necessary. A similar comparison of H1 vs H2 would also be useful if the releases can be reasonably implemented at the coarsest level. Even if the release location would have to be assumed to be wider or poorly matched in terms of exact location, injection of the same amount of alkalinity in the coarsest model could be interesting to determine to what extent the H2 level is required.

Response: See also response to Comment 3 by Reviewer 3. H1 was designed to provide reasonably accurate boundary conditions to H2 on the Scotian Shelf but remains quite coarse (~760 m) with respect to the Halifax Harbour. For comparison, H2 has a

resolution of 150 m. The shape and circulation of the Harbour is not at all well resolved in H1, which affects residence time. We note that the purpose of this model is to accurately represent the dynamics in Halifax Harbour and the purpose of the manuscript is to describe this.

The Reviewer's suggestion to compare at the coarser resolutions is a good one for a study that is aimed at assessing how resolution affects transport and dispersion of alkalinity and DIC on the shelf. However, we wouldn't use H1 for this purpose because it only covers a small portion of the shelf (note that alkalinity, once it leaves Halifax Harbour, is transported out of the H2 domain in just 10 to 14 days). We are working with a larger-scale ROMS model, described in Ohashi et al. (2024, <https://gmd.copernicus.org/articles/17/8697/2024/>), that has a resolution similar to H1 and intend to compare this with coarser-resolution models.

Comment:

6. L769 It was a surprise to read here that the sediment loss term was set to zero. I feel like this should have been mentioned earlier, perhaps even right when the loss term(s) are introduced in L317ff. Is both w_p and θ_{loss} set to zero or just the latter? If it's just the latter, does the model currently just settle all the particles on the floor and let them dissolve from there until completely dissolved?

Response: For clarity, the loss term set to zero will be mentioned in Section 5.1. in the revised manuscript.

For an explanation on sinking and sediment loss, see response to Comment 3 above. Indeed, all particles settle to the bottom in the current setup and dissolve. Particles are also transported by current and can be resuspended by vertical mixing. We chose to set sediment loss to zero because unfortunately we do not have information on the influence of the sediment on Brucite particles in the Halifax Harbour, as mentioned in L569.

Comment:

7. L120: I assume the conversion factor is 1025 kg m^{-3} , not 1.025 kg m^{-3} (remove dot or change dot to comma)

Response: We did assume a water density of 1025 kg m^{-3} but the conversion factor is 1.025 because we convert μmol to mmol . We will clarify this conversion in the revised manuscript.

Comment:

8. L243 In equation (3), it appears that the parameter “c1” is duplicate as a coefficient to t and as an exponent. Likely it is meant to be c_2 instead?

Response: Indeed, this is a typo, it should be c_2 .

Specific comments

Comment:

9. L325 change to “is added that mimics” or “is added to mimic”

Response: Done.

Comment:

10. L331 “1.29 ml s⁻¹”, exponentiate the “-1”

Response: Done.

Comment:

11. L475 In such cases,

Response: Done.

Comment:

12. Fig.1D consider using a different color scheme for the bathymetry as the scale is different.

Response: We feel that an alternate color scheme is not necessary since the color bar is available in Figures 1c and 1d. However, since the color bar is missing from Figure 1a (same as in Figure 1c) we will add it in the revised manuscript.

Comment:

13. Figs. 3, 5,6,7, 10: Is it possible to indicate the release location in these plots with a small black arrow or similar. I know they are shown in Fig 1 D, but it would be very helpful to have that info on each of the other plots too.

Response: Yes, we will add a dot showing the release location in these plots.

Comment:

14. Figure 7: It would be nice to add a horizontal dashed line to the two graphs indicating the theoretical maximum uptake (at your CO₂ efficiency of 0.89) to get a sense for what fraction of the ultimate uptake occurs within the simulation domains.

Response: Yes, we will add an additional y-axis on the right indicating the realized uptake (this information is also available in Figure 9).

Comment:

15. Fig S4-S8 The observations of the depth profiles are sparse enough in time that it's difficult to assess visually how closely the corresponding model predictions match. Perhaps, for each observation time and depth simply make a scatter plot against the corresponding prediction value ? Could be color coded by depth perhaps to see if correlation is better at surface vs depth.

Response: This information is somewhat already synthesized in the statistics (Table 1) but for a visual comparison we will also add 1:1 plots in the supporting material of the revised manuscript.