

Authors’ response to reviewers for “Semi-Lagrangian advection in the NEMO ocean model”

June 26, 2020

To begin, we are thankful for the review comments provided by the reviewers Florian Lemarié (reviewer #1) and Mike Bell (reviewer #2). Collectively, the reviews have highlighted areas where the manuscript was lacking, particularly with respect to its clarity and presentation. We have revised the manuscript with this feedback in mind, making changes to incorporate or otherwise address this feedback.

Because both reviewers were in broad agreement about the weaknesses of the paper and did not provide conflicting feedback, a summary of changes is listed below by section. Additionally, the revised manuscript and a “diff” of changes will be submitted after this comment is finalized.

Title and abstract

Both reviewers noted that the original title of the article was too bold, to put it bluntly. Additionally, the executive editor pointed out the GMD requirement that papers which refer to development for a single model must mention the model name and version number in the title. Consequently, the title is now “Development of a semi-Lagrangian advection scheme for the NEMO ocean model (3.1).”

Reviewer #2 additionally noted that the original abstract gave the impression that the method developed in this article is “very similar to the SISL algorithms used in atmospheric models.” The first paragraph of the abstract is now revised to hopefully make the distinction more clear, and the second paragraph is reworded for simplicity.

Introduction

The introduction has been greatly expanded in response to the reviewers’ comments. In textual order:

- Reviewer #2 noted that the leading paragraphs on coupled modeling seemed to be a tangential motivation. This has the awkward characteristic of being tangential yet true – the idea of applying semi-Lagrangian advection to the ocean model at CMC came out of realizing the computational cost of running coupled forecast systems. This section has been revised and modestly expanded to make the practical focus more clear.
- Both reviewers note that the discussion of why semi-Lagrangian advection might help the timestep size in ocean models was lacking. This is now more comprehensibly discussed in new section 1.1, which draws the suggested direct contrast between timestep-limiting factors in atmosphere and ocean circulation models. Grid stretching is now a subsection to this discussion, which (at Reviewer #2’s request) now also includes a brief description of the ocean flows in the gridpoint-clustered portion of the Canadian Arctic Archipelago.
- The “existing work” subsection (now 1.2) more directly engages with the literature on conservation-preserving semi-Lagrangian methods (noting this as not implemented but a future possibility, and without this interpretation semi-Lagrangian advection has a finite-difference interpretation) and ALE coordinates.

The other points raised by reviewer #1 (the “two other levels of constraints”) are generally agreed to but addressed in the main body of the text as the issues arise.

Time discretization

- Both reviewers note that the notation in this section was awkward. Consequently, we have entirely revised this section to use a simpler notation of f^B (“before”), f^N (“now”), and f^A (“after”) that should be familiar to readers from the leapfrog context, adapting it to semi-Lagrangian advection to add f^D (“departure”). This also resulted in small changes to the notation in subsequent sections for consistency.
- This section is also reorganized to separate the semi-Lagrangian advection (2.1) from its reconciliation with the leapfrog algorithm (2.2) to clarify the changes in perspective.
- Reviewer #1 also expressed doubts about whether semi-Lagrangian advection as-defined was robust to the Asselin timestepping filter. This analysis is now present in the new subsection (2.3), and the Asselin filter does not negatively affect the stability of semi-Lagrangian advection as-implemented.

Interpolation

- Reviewer #2 remarked that the discussion of two-dimensional interpolation was cumbersome and verbose. What was formerly the non-numbered “two-dimensional application” subsection has been removed, with the comment briefly summarized and placed just before the slope-limiting discussion.
- Reviewer #2 also noted that the discussion of vertical advection was confusing, especially the claims about discontinuous derivatives. Subsection 3.2 has been revised and reworded.
- We also revised the subsequent discussion of vertical slope-limiting to clarify (at reviewer #2’s note of confusing language) why it is necessary at the bottom boundary in the presence of partial cells.
- A reference (Turkington et al., 1991) has been added for the numerical example on this section, at reviewer #2’s request. To our knowledge this is not a standard test-case in the semi-Lagrangian literature, but nonlinear generalizations of this approach are a standard technique for calculating the profiles of nonlinear internal gravity waves.
- Reviewer #1 noted that the description of the advection constraint for the Eulerian/leapfrog numerical example in this section was “fuzzy,” and so we have adopted the more precise definition. This led to no practical difference in the calculation, since the maximum Courant number in the domain is reached at the top and bottom boundaries where the vertical velocity is zero. (This did, however, lead to a discovery of a small bug in the code that generated this figure, which used the wrong vertical mode number to calculate wave-induced horizontal velocities for the purposes of evaluating the Courant number. This has been addressed in the submitted code repository and the figure regenerated; there is essentially no difference in results.)

Additionally, we replaced “CFL number” in the paper with “Courant number” throughout, since the latter concept is indeed the intended use of the term.

- Reviewer #1 also inquired about the performance of the Eulerian/leapfrog method in this section with a maximum Courant number close to 1. We investigated this over the range 0.2–0.99 and found little difference in error compared to the exact solution; this is mentioned in-text rather than by adding more lines to figure 3.

Trajectory calculation

- Reviewer #2 notes that the discussion about extrapolating into the boundary is confusing, and reviewer #1 asks whether this semi-Lagrangian method faces a Lipschitz stability condition. These are the same issue: the problem of extrapolating into the boundary arises only when a calculated trajectory would cross that of a fluid parcel that begins and remains on the (no normal flow) boundary. Consequently, we have revised the first part of subsection 4.1 to make this connection.

Numerical results

- Reviewer #2 requested more clarification on the inconsistency between the semi-Lagrangian advection and the Eulerian application of forcing. This amounts to an $O(\Delta t)$ approximation in the integral form of the semi-Lagrangian advection equation, and this is now noted in the discussion in subsection 5.1.
- Reviewer #1 notes that NEMO’s TVD scheme is really a “tracer variance dissipation” scheme. This has been changed throughout with a citation to Lévy et al. (2001) at the first mention.
- Both reviewers remarked on the relatively short timestep used in the ORCA025 runs of section 5.2. New footnote 8 has been added to provide more context; in brief the ice/ocean drag parameter is increased following Roy et al. (2015), which makes the problem more apparent for the operational forecasting configuration than for typically-presented runs. At the same time, we wanted to maintain the same physical parameterizations between the operational configuration and the runs presented in this paper. Addressing this problem would be ideal and is the focus of ongoing work, but the runs of section 5.2 took long enough to complete on the shared supercomputing resources that they cannot be practically be repeated in GMD’s peer-review timeframe even if a solution were immediately at hand.
- Reviewer #1 asked about the number of iterations taken to find trajectories and the effect of trajectory truncation. This is now discussed further in section 5.2; the mean number of trajectory iterations per cell for the semi-Lagrangian tracer run was 1.004, so truncated trajectories were truly exceptional. The performance cost of trajectory iteration is also addressed in the conclusions.
- Reviewer #2 requested expanded commentary on the MOC and circumpolar current results, which we have provided. Because these runs do report preliminary results, we want to be cautious about reporting false confidence that semi-Lagrangian advection causes physically-relevant changes in results that may not in fact be robust, but we agree that we erred on the terse side here.
- Reviewer #2 also requested a look at the mean global temperature profile at the start and end of the simulation. This is the new figure 9, with brief discussion at the end of section 5.

Conclusions

- Reviewer #2 requested a longer summary of the achievements, particularly one that highlights new algorithms. This is now added at the beginning of section 6, where we have added a list that highlights the core algorithms of this paper.
- Both reviewers had questions about the performance of the method and its parallel implementation. This is now dealt with in the conclusions, under the new (non-numbered) “performance and implementation” subsection.
- Reviewer #2 requested a deeper look at the application to climate simulations, and consequently we have expanded the discussion in the commentary on the results. The temperature profile results (specifically temperature stability in deep water) seem to be encouraging for climate applications, but we reserve a full recommendation for a future day when either temperature/salinity drift is fully characterized (and found to be acceptable) or conservation is explicitly added.
- Reviewer #2 also requested a brief discussion of how this algorithm might apply to the RK3 timestepping algorithm used in upcoming versions of NEMO. Since this is very much “future work” for both NEMO and semi-Lagrangian advection, we have added this discussion to the conclusions.

References

- Lévy, M., Estublier, A., and Madec, G.: Choice of an advection scheme for biogeochemical models, *Geophysical Research Letters*, 28, 3725–3728, <https://doi.org/10.1029/2001GL012947>, URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2001GL012947>, 2001.

Roy, F., Chevallier, M., Smith, G. C., Dupont, F., Garric, G., Lemieux, J.-F., Lu, Y., and Davidson, F.: Arctic sea ice and freshwater sensitivity to the treatment of the atmosphere-ice-ocean surface layer, *Journal of Geophysical Research: Oceans*, 120, 4392–4417, <https://doi.org/10.1002/2014JC010677>, URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014JC010677>, 2015.

Turkington, B., Eydeland, A., and Wang, S.: A Computational Method for Solitary Internal Waves in a Continuously Stratified Fluid, *Studies in Applied Mathematics*, 85, 93–127, <https://doi.org/10.1002/sapm199185293>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sapm199185293>, 1991.

~~Semi-Lagrangian advection in~~ Development of a semi-Lagrangian advection scheme for the NEMO ocean model (3.1)

Christopher Subich¹, Pierre Pellerin¹, Gregory Smith¹, and Frederic Dupont¹

¹Environment and Climate Change Canada, 2121 route Transcanadienne ouest, Dorval Québec, Canada H9P 1J3

Correspondence: Christopher Subich (Christopher.Subich@canada.ca)

Abstract. ~~As model resolutions~~ As resolutions of ocean circulation models increase, the ~~Courant-Frederichs-Lewy (CFL) number based on advective motion becomes the limiting factor in setting the timestep of time-explicit circulation models~~ advective Courant number – the ratio between the distance travelled by a fluid parcel in one timestep and the grid size – becomes the most stringent factor limiting model timesteps. Some atmospheric models ~~escape~~ have escaped this limit by using an implicit or semi-implicit semi-Lagrangian formulation of advection. ~~This formulation calculates~~, which calculates materially-conserved fluid properties along ~~parcel~~ parcel trajectories which follow the fluid motion and end ~~, for each timestep~~, at prescribed grid-points. Unfortunately, this formulation is not straightforward in ocean contexts, where the irregular, interior boundaries imposed by the shore and bottom orography are incompatible with traditional trajectory calculations.

This work ~~is the first application~~ describes the adaptation of the semi-Lagrangian method ~~to an~~ as an advection module for an operational ocean model. ~~In this context, we solve the difficulty posed by~~ We solve the difficulties of the ocean's ~~irregular, interior~~ internal boundaries by calculating parcel trajectories using a time-exponential formulation. ~~This formulation~~, which ensures that all ~~trajectories that are solutions to a fixed-point iteration have an origin point in the valid domain~~, and it does not ~~require any prescribed extrapolation of the fluid velocities into the invalid (land) portion of the domain~~. We parcel trajectories remain inside the ocean domain despite strong accelerations near the boundary. Additionally, we derive this method in a way that is compatible with the leapfrog timestepping scheme used in the NEMO-OPA (Nucleus for European Modelling of the Ocean, Océan Parallélisé) ocean model, and we present simulation results for a simplified test-case of flow past a model island and for 10-year free runs of the global ocean on the quarter-degree ORCA025 grid.

Copyright statement. TEXT

1 Introduction

Recent work by Smith et al. (2018) has shown that over the medium term (up to seven days), a coupled forecasting system involving ocean, ice, and atmospheric models can significantly improve forecasting skill over forecasts that ~~assume persistence of initial conditions~~ persist initial ocean and ice conditions over the atmospheric forecast period. While this is an exciting development for the future of numerical weather prediction, ~~it creates a combined computational problem out of models and~~

~~forecasting systems systems that have previously been run independently.~~ coupling adds a new dimension to the computational cost. Developing a deployable forecast system, especially with regional or ensemble components, requires exploiting every reasonable opportunity for optimization. One straightforward optimization is to maximize the admissible timestep of the ocean component, and we intend to improve the ocean timestep limit by implementing a semi-Lagrangian advection module into the popular NEMO-OPA (Nucleus for European Modelling of the Ocean, Océan Parallélisé; Madec (2008), version 3.1) model, used in this coupled system. This module is intended as a drop-in replacement for the model's other advection modules, and in particular it does not interfere with NEMO's time-stepping algorithm (leapfrog).

~~In particular, these coupled forecasts require close integration between the atmospheric and ocean components in order to exchange information. For reasons of computational efficiency, we want each model to run with its largest admissible timestep, but to support coupling we can usually only achieve this if the models have the same timestep or closely-related timesteps. While different factors cause the most stringent timestep restrictions in atmosphere and ocean circulation models, these models can adopt similar methods to alleviate the restriction~~

1.1 Timestep constraints in the ocean

A numerical model with an explicit time-marching scheme must generally limit its timestep to satisfy a Courant-Frederichs-Lwey (CFL) condition: information must not propagate more than a discretization-defined maximum number of cells in a single step, leading to a maximum stable Courant number. For hyperbolic systems such as the Euler equations (for the atmosphere) or hydrostatic equations (as implemented by NEMO-OPA), the information propagation speeds are controlled by the admissible wave modes of the systems, which become characteristic curves.

In the atmosphere, the most restrictive wave mode is that corresponding to sound waves. These waves are fast compared to atmospheric motions, and in response atmospheric models generally treat sound waves either implicitly or through sub-cycling, especially in the most restrictive vertical direction. The second most stringent restriction comes from simple advection by winds in the upper atmosphere. At the Canadian Meteorological Centre, ~~our coupled numerical weather forecast the atmospheric forecasting system (and atmospheric component of the coupled forecasting system)~~ uses the GEM (Geophysical Environmental Multiscale; Girard et al. (2014)) model ~~for its atmospheric component and the~~, which addresses this timestep restriction through a semi-Lagrangian treatment of advection (Robert, 1982).

In the ocean, the hydrostatic equations already eliminate sound waves, but the model is left with the problem of surface gravity waves. Here, NEMO-OPA ~~(Nucleus for European Modelling of~~ takes a similar approach to that used by atmospheric models for sound waves, by either treating the surface pressure gradient in a time-implicit manner (with a linearized free surface, used in this work), or by sub-cycling. The ocean lacks any direct equivalent to the atmosphere's strong upper-air winds, and so advection by the background velocity and internal gravity wave modes compete as the next most limiting factor for the maximum stable timestep. Lemarié et al. (2015) finds that the ~~Ocean, Océan Parallélisé; Madec (2008)) model for the ocean component. The GEM model already includes several features to provide for a long model timestep, and we seek to adapt some of these to the widely-used NEMO-OPA model (version 3.1).~~ Courant number associated with vertical advection

is more limiting than that associated with internal (baroclinic) gravity waves at resolutions of $\frac{1}{2}^\circ$, and the Courant number associated with horizontal advection catches up with that of gravity waves at resolutions of $\frac{1}{4}^\circ$ and finer.

1.2 Stretching in ocean grids

60 Grid stretching

~~While ocean currents are much slower than the fastest upper-atmosphere winds that contribute to the Courant-Frederichs-Lewy (CFL) restriction in atmospheric model¹, global ocean models in particular suffer from grid stretching. The~~ In order to cover the entire ocean in a single, continuous domain, global NEMO-OPA model configurations typically use grids based on the ORCA “tripolar” grid commonly used in global NEMO-OPA model configurations (Madec and Imbard, 1996; Murray, 1996)
65 (Madec and Imbard, 1996; Murray, 1996). This grid is defined in the northern hemisphere by an elliptical coordinate system, where ~~in the northern hemisphere circles of a~~ the latitude-like coordinate ~~are~~ is defined by ellipses with a shared pair of foci ; ~~and hyperbolas of a~~ and the longitude-like coordinate ~~are defined by~~ is defined by the hyperbolas orthogonal to these ellipses; ~~these~~. These coordinates match continuously at the equator to lines of latitude and longitude ~~on~~ in a Mercator projection. By placing the foci of the ellipses on land, the grid contains no singularities in the ocean domain.

70 Unfortunately, this placement causes an abundance of small grid cells in the north polar region, especially in the Canadian Arctic Archipelago. Figure 1 depicts this situation at a nominal $\frac{1}{4}^\circ$ resolution: the grid point spacing of 25-30km near the equator falls to 3-4km in the archipelago. ~~Currents in these narrow straits contribute ten times as strongly to a lateral CFL timestep restriction, compared to currents in the equatorial regions~~ The areas in figure 1 with the narrowest grid spacing are also shallow seas, with depths of 200m or less and non-tidal currents of 15–30cm/s. This grid stretching is of particular
75 concern when adapted to regional models such as that of Lemieux et al. (2016), which refine this grid while retaining its tripolar structure to use conforming boundary conditions.

The coordinate system is also stretched in the vertical direction. Using the z-level grid option of the NEMO-OPA model, layers near the surface are spaced much more closely together than layers nearer the ocean bottom, in order to provide adequate resolution of the mixing layer. ~~While ocean currents tend to follow nearly horizontal paths, residual upwelling or~~
80 ~~downwelling can still cause a vertical CFL restriction to be binding.~~ This stretching enhances the impact of vertical advection on the vertical Courant number, even if vertical-current magnitudes are low in absolute terms; Lemarié et al. (2015) notes that vertical advection provides a tighter bound on the timestep than horizontal advection.

~~Although the ocean models arrive at a binding CFL condition via a different route than for atmospheric models, semi-Lagrangian advection can alleviate this restriction in both cases. This method traces~~ Semi-Lagrangian advection alleviates both vertical and
85 horizontal Courant number restrictions by tracing fluid parcels in a Lagrangian, fluid-following coordinate system, ~~defined such that at~~. This coordinate system is defined so that the end of each timestep the fluid parcels arrive ~~at the prescribed grid.~~

¹ Atmospheric models are also limited by the timestep associated with sound waves, but it is common to process at least vertically-propagating sound waves implicitly to alleviate this restriction. A comparable limit in ocean models is that arising from surface gravity waves, and here NEMO-OPA uses either an implicit or time-splitting approach for similar reasons.

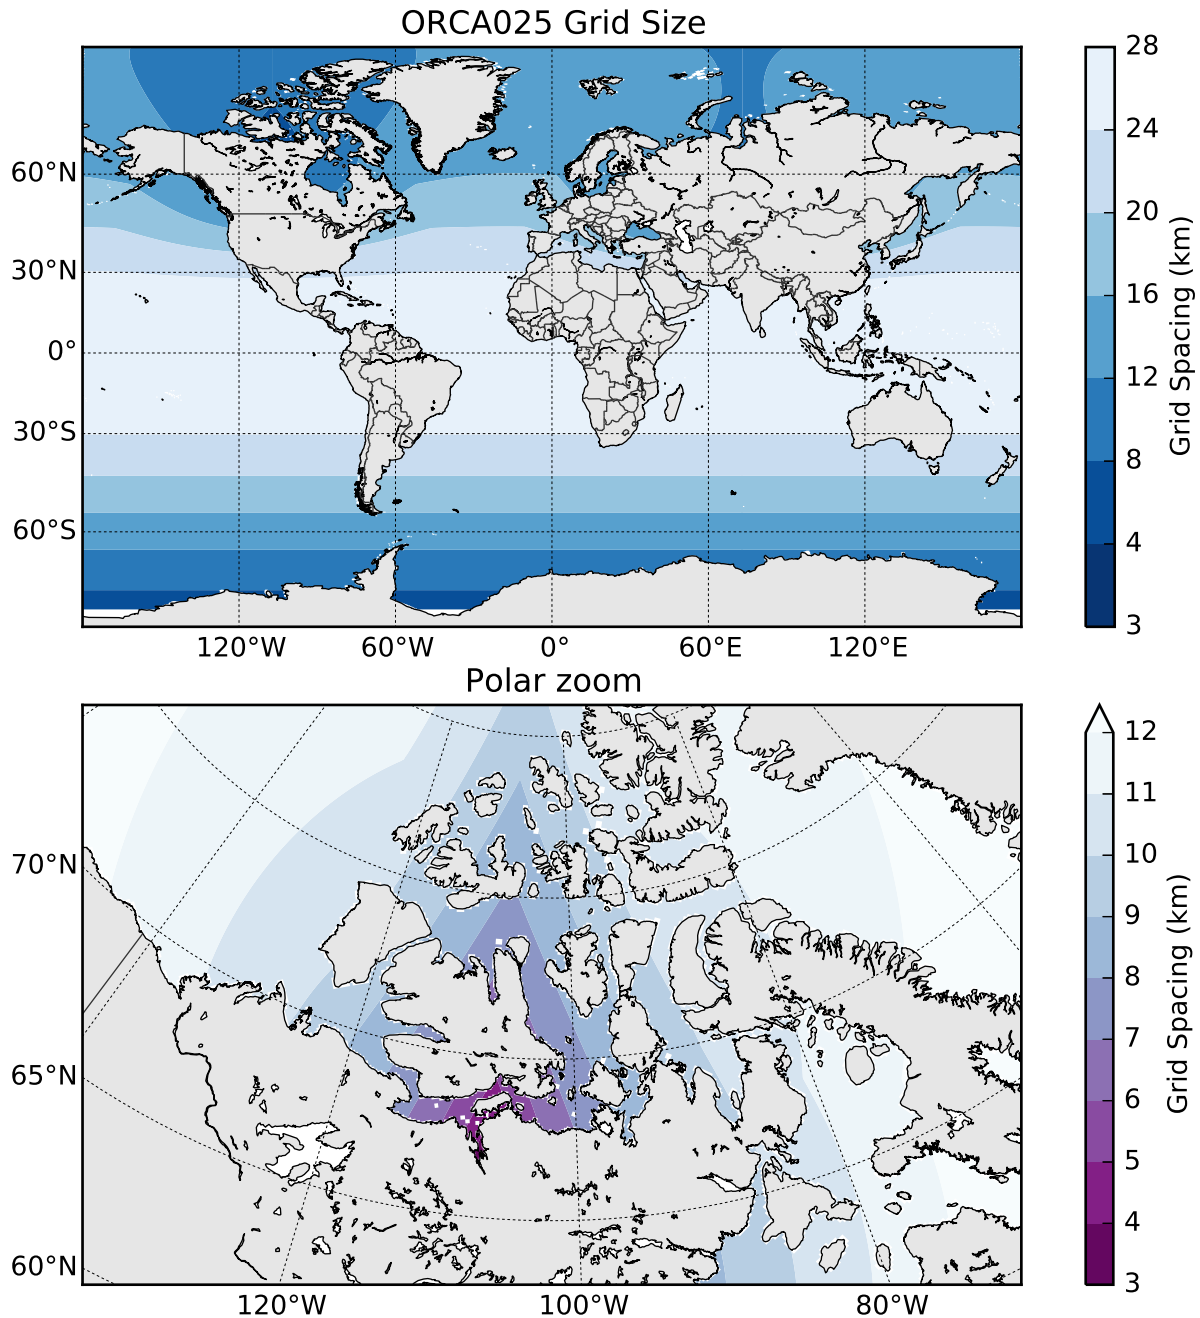


Figure 1. Grid size (defined as $\min(e1t, e2t)$) on the ORCA025 grid. At top: in the global view, gridpoint spacing gradually decreases from the equator towards the north and south poles. At bottom: in a detail view of the north polar region, the grid is especially high-resolution in the southern portion of the Canadian Arctic Archipelago, with gridpoint spacing as low as 3km.

~~The on the prescribed computational grid, and the~~ properties of the fluid parcels (in the ocean setting, temperature, salinity, and horizontal velocity) at the ~~beginning of the time step~~ end of the timestep (on the computational grid) are found by interpolating ~~these values from their gridded locations~~ the previous-step, gridded values to the origin point of ~~the each parcel's~~ trajectory.

90 This method provides an implicit treatment of advection, allowing timesteps with CFL-advective Courant numbers greater than those ~~allowable under wholly explicit methods~~ usually permitted by explicit, Eulerian-form models.

In this work, we describe the initial implementation of a semi-Lagrangian advection routine in NEMO-OPA, based on the configuration of Smith et al. (2018). This configuration uses a linear free surface where the vertical coordinate does not move in time, but we believe that the described method can be generalized.

95 1.2 Existing work

~~This technique is standard in atmospheric models (Robert, 1982), but it is not widely applied to the ocean.~~

In the atmosphere, semi-Lagrangian advection is a standard technique (Robert, 1982) for the implicit treatment of advection, but especially at large scales ~~the effects of topography are relatively gentle, and~~ In particular, trajectory calculations can proceed under the assumption that the fluid parcel ~~experiences relatively little acceleration. This assumption is more significantly~~ broken by the irregular bathymetry and coastline of the ocean ~~does not experience strong boundary-related acceleration. In the ocean domain this assumption is strongly violated, particularly for z-level vertical grids where the bathymetry changes abruptly at lateral cell boundaries.~~

100 does not experience strong boundary-related acceleration. In the ocean domain this assumption is strongly violated, particularly for z-level vertical grids where the bathymetry changes abruptly at lateral cell boundaries.

Some attempts have been made previously to incorporate semi-Lagrangian advection into the ocean context. The work of Casulli and Cheng (1992), which is used as part of the ELCOM lake and estuary model (Hodges and Dallimore, 2006), calculates parcel trajectories via a substepping approach, where fluid parcel trajectories are integrated via an explicit Euler method over many short steps per model timestep. The two-dimensional, unstructured shallow water model of Walters et al. (2007) takes a similar approach, where it also must take at least one substep per element boundary traversed by a fluid parcel.

105 calculates parcel trajectories via a substepping approach, where fluid parcel trajectories are integrated via an explicit Euler method over many short steps per model timestep. The two-dimensional, unstructured shallow water model of Walters et al. (2007) takes a similar approach, where it also must take at least one substep per element boundary traversed by a fluid parcel.

In this work, we overcome this difficulty with an iterative trajectory calculation which reduces in the limit to an implicit trapezoidal rule. In addition, we also derive the semi-Lagrangian advection scheme in a form which calculates effective advective tendencies, such that the advection routine can operate as a “plug-in” scheme for models which traditionally use Eulerian fluxes. We apply this to the NEMO-OPA model, and we believe this algorithm may be useful when applied to other ocean models with a structured grid.

110 advective tendencies, such that the advection routine can operate as a “plug-in” scheme for models which traditionally use Eulerian fluxes. We apply this to the NEMO-OPA model, and we believe this algorithm may be useful when applied to other ocean models with a structured grid.

In exchange, however, the semi-Lagrangian advection formulation departs from NEMO's finite-volume interpretation of its tracer and velocity components. By tracing infinitesimal fluid parcels, semi-Lagrangian advection treats gridpoint values analogously to a finite-difference method, and as a consequence the scheme does not naturally offer conservation guarantees. This is not a primary concern for the short to medium-term forecasting applications that form the direct target for this work, but extensions of the semi-Lagrangian scheme to ensure conservation (Lauritzen, 2007) may be needed before the technique is applicable to longer-term climate simulations.

115 analogously to a finite-difference method, and as a consequence the scheme does not naturally offer conservation guarantees. This is not a primary concern for the short to medium-term forecasting applications that form the direct target for this work, but extensions of the semi-Lagrangian scheme to ensure conservation (Lauritzen, 2007) may be needed before the technique is applicable to longer-term climate simulations.

Additionally, Leclair and Madec (2011) has developed an “arbitrary Lagrangian-Eulerian” vertical coordinate scheme, implemented in recent versions of NEMO. This scheme splits vertical motions into fast (high temporal frequency) and slow motions, and the

120 in recent versions of NEMO. This scheme splits vertical motions into fast (high temporal frequency) and slow motions, and the

former are treated by co-moving vertical coordinate surfaces with a regriding step. This coordinate system reduces spurious diapycnal mixing caused by the high-frequency vertical motions, but its Lagrangian treatment of these motions is likely to allow for a larger maximum Courant number in the vertical.

1.3 Organization

125 We first introduce the time discretization of the semi-Lagrangian scheme in section 2, in order to develop a formulation that remains compatible with the common leapfrog scheme. In section 3, we begin to spatially discretize the semi-Lagrangian scheme by specifying the horizontal and vertical interpolation operators, and in section 4 we complete the discretization by defining the trajectory calculations. We present preliminary numerical examples in section 5, demonstrating the stability of the advection scheme.

130 2 Time discretization

The first requirement of a semi-Lagrangian advection scheme for the NEMO-OPA model is that it be consistent with the model's overall timestepping approach: the advection scheme is but one component of the full model.

In version 3 of NEMO-OPA, non-diffusive, non-damping processes such as advection are implemented via the leapfrog scheme (Mesinger and Arakawa, 1976), where at each timestep a field f receives its new value at $f^{t_0+\Delta t}$ (f "after" the timestep) based on its value at the previous timestep and forcing terms, which are all evaluated on the reference grid \mathbf{x}_{ref} . This gives a schematic of:

$$f^{t_0+\Delta t}(\mathbf{x}_{ref}) = f^{t_0-\Delta t}(\mathbf{x}_{ref}) + 2\Delta t \text{RHSE}_f \mathbf{F}(\mathbf{x}_{ref}), \quad (1)$$

where RHSE_f (Eulerian right hand side) generally stands in for both true forcing terms (such as heating or evaporation) as well as time-tendency terms from the transport and momentum equations. For advective processes f^A is the field calculated at time $t_0 + \Delta t$, f^B is the field evaluated at the known prior time $t_0 - \Delta t$ ("before"), f^N is the field at the provided time t_0 ("now"), and \mathbf{F} is the forcing operator. The forcing operator includes advective processes at the RHSE terms arising from tracer and momentum flux are evaluated at the current "now" time-level, whereas RHSE terms arising from but diffusive, damping, and hydrostatic pressure may terms might be evaluated at either the previous or new "before" or "after" time-levels.

This is an Eulerian approach to fluid motion, where tracer and momentum values are tracked at specific locations (namely the grid points) over time along the fixed reference grid at all times, and fluid flows through this grid.

2.1 Semi-Lagrangian advection

In contrast, the semi-Lagrangian advection scheme considers Lagrangian advection schemes consider the fluid parcel to be the fundamental unit of discretization. In this perspective, if f is a property of a fluid parcel that is conserved along a trajectory¹,

¹This is true for temperature, salinity, and momentum provided the ocean is treated as an incompressible fluid. This assumption is satisfied by NEMO-OPA's adoption of the Boussinesq approximation.

it satisfies the continuous equations:

$$150 \quad \frac{D}{Dt} f(\mathbf{x}(t)) = \text{RHSL}_f \mathbf{F}_L(\mathbf{x}(t)), \quad (2)$$

where $\frac{D}{Dt} = \partial_t + \mathbf{u} \cdot \nabla$ is the material derivative and $\text{RHSL}_f \mathbf{F}_L$ (Lagrangian right-hand side) contains all the same forcing terms as $\text{RHSE} \mathbf{F}$ *except* those arising from tracer and momentum flux, which are included inside the material derivative.

Ordinarily, (2) is discretized so that $\text{RHSL}_f \mathbf{F}_L$ is evaluated following the Lagrangian particles. ~~One such approach is the semi-implicit in the moving coordinate frame $\mathbf{x}(t)$, satisfying the trajectory equation:~~

$$155 \quad \frac{D}{Dt} \mathbf{x}(t) = \mathbf{u}(\mathbf{x}(t)). \quad (3)$$

~~From an Eulerian point of view, (3) is a trivial identity based on the definition of the material derivative, but from the Lagrangian point of view (3) must be solved to define \mathbf{x} over time.~~

~~One technique for solving (2) and (3) is the two time-level implicit semi-Lagrangian (SISL) method, used in the GEM atmospheric model (Girard et al., 2014) among others, where the RHSL_f . Here, the \mathbf{F}_L terms are evaluated using with a trapezoidal rule to give:-~~

$$f^{t_0+\Delta t}(\mathbf{x}(t_0 + \Delta t)) = f^{t_0}(\mathbf{x}(t)) + \frac{\Delta t}{2} (\text{RHSL}_f^{t_0+\Delta t}(\mathbf{x}(t_0 + \Delta t)) + \text{RHSL}_f^{t_0}(\mathbf{x}(t))),$$

~~discretizing (2) and (3) as:~~

$$f^A(\mathbf{x}_{ref}) = f^N(\mathbf{x}^D) + \frac{\Delta t}{2} (\mathbf{F}_L^A(\mathbf{x}_{ref}) + \mathbf{F}_L^N(\mathbf{x}^D)) \quad (4a)$$

$$\mathbf{x}_{ref} = \mathbf{x}^D + \frac{\Delta t}{2} (\mathbf{u}^A(\mathbf{x}_{ref}) + \mathbf{u}^N(\mathbf{x}^D)). \quad (4b)$$

165 ~~where $\mathbf{x}(t_0 + \Delta t)$ is constrained to be coincident with the computational grid (abbreviated as the unadorned \mathbf{x} in the subsequent) and $\mathbf{x}(t)$ is solved for iteratively. The trajectory equation (4b) acts to implicitly define the paths of the traced fluid parcels, where each location on \mathbf{x}_{ref} is associated with a corresponding departure-point location \mathbf{x}^D . Over the single timestep, fluid parcels depart from \mathbf{x}^D (which in general is not aligned with the grid) and arrive on the reference grid.~~

~~The difference between and is more than just cosmetic. The forcing terms in are evaluated. This off-grid either directly or via interpolation, and doing so inside, departure point evaluation of \mathbf{u} and \mathbf{F}_L is fundamental to Lagrangian and semi-Lagrangian methods, and $f^N(\mathbf{x}^D)$ ($\mathbf{F}_L^N(\mathbf{x}^D)$) can be written more simply as f^D (\mathbf{F}_L^D) for “departure-point f (\mathbf{F}).” Neither the time-implicit evaluations (generally) nor the off-grid evaluations (of non-advective forcing) are compatible with the core structure of NEMO-OPA would be incompatible with its core structure, which considers advection to be just one of many operators that influences the RHSE term in independent operators influencing the \mathbf{F} term of (1).~~

175 ~~To reconcile these differences, the-~~

2.2 Reconciliation

180 ~~Implementing semi-Lagrangian method described in this work adopts advection in NEMO-OPA requires adopting as much of the framing framework of (1) as possible, including evaluating the forcing terms strictly on-grid. To adapt without changing the evaluation of non-advective forcing terms. Effectively, the semi-Lagrangian method to this framework advection routine must ultimately supply a time-trend that, from the perspective of the leapfrog timestep algorithm, is indistinguishable from a conventional flux-form advection operator.~~

~~To effect this, consider (2) without forcing terms ($\text{RHS}_L = 0$) over the interval $t_0 - \Delta t$ to $t_0 + \Delta t$ ($\mathbf{F}_L = 0$). The function f is preserved following the flow, so this gives ~~the simply-written:~~~~

$$\frac{D}{Dt} f^A = \frac{0^{t_0+\Delta t}((t_0 + \Delta t)) - 0^{t_0-\Delta t}((t_0 - \Delta t))}{\Delta t} f^D. \quad (5)$$

185 ~~Equating This is approximated by taking one timestep of (1) (with only advective forcing \mathbf{F}_{adv}), but the latter involves integrating over the whole interval from $t_0 - \Delta t$ to $t_0 + \Delta t$. Thus, we should identify f^D (and the departure points generally) not with the “now” time-level in the leapfrog scheme, but with the “before” time-level. Doing so and equating (5) and ~~noting that \mathbf{x} in corresponds $\mathbf{x}(t_0 + \Delta t)$ in, gives:~~(1) gives:~~

$$f_{\sim}^{t_0+\Delta t}((t_0 + \Delta t))^A = f_{\sim}^{t_0-\Delta t}((t_0^B + \Delta t)) + 2\Delta t \cdot \text{RHS}_{f,adv}((t_0 + \Delta t)) \mathbf{F}_{adv} = f_{\sim}^{t_0-\Delta t B}(\mathbf{x}(t_0 - \Delta t^D)), \text{ or} \quad (6)$$

$$190 \quad \text{RHS}_{f,adv}() \mathbf{F}_{adv} = \frac{1}{2\Delta t} f_{\sim}^{t_0-\Delta t}((t_0^B - \Delta t)) - f_{\sim}^{t_0-\Delta t B}(\mathbf{x}(t_0 + \Delta t^D)), \quad (7)$$

~~which defines the advective forcing term $\text{RHS}_{f,adv}$ that, in the context of the leapfrog method, emulates the behaviour of semi-Lagrangian advection.~~

Equation ~~forms the fundamental equation for this work, and the subsequent focuses on the interpolations necessary to approximate $\mathbf{x}(t_0 - \Delta t)$ and $f(\mathbf{x}(t_0 - \Delta t))$. Importantly, the effects of the Earth’s curvature and the curvilinear grid used~~
 195 ~~are not included in this equation, even when (7) is prescriptive, and it gives the necessary trend for the leapfrog algorithm. Evaluating it requires f assumes the role of the u or v components of velocity. only at the already-known “before” time-level and calculation of the departure points \mathbf{x}^D . This calculation is further simplified by basing the departure points on the time-centered velocities \mathbf{u}^N , and the exact algorithm for this calculation will be discussed in more detail in section 4.~~

2.3 ~~Effects of the Asselin filter~~

200 ~~To prevent decoupling of odd and even timesteps (damping the computational mode), NEMO-OPA separately includes the effects of planetary vorticity and of the coordinate metric in the governing equations, is typically configured to use the Asselin time filter (Asselin, 1972), which adds a small time-damping proportional to $\frac{\partial^2}{\partial t^2} f$. Using the notation of Shchepetkin and McWilliams (200~~

adapted to (1), the filter extends the time-marching scheme to the sequence:

$$f^{A*} \leftarrow f^B + 2\Delta t \mathbf{F}^{N*} \quad (8a)$$

$$205 \quad f^N \leftarrow \epsilon f^{A*} + (1 - 2\epsilon) f^{N*} + \epsilon f^B \quad (8b)$$

$$f^{N*} \leftarrow f^{A*} \quad (8c)$$

$$f^B \leftarrow f^N \quad (8d)$$

Equation (8a) is the direct equivalent of (1), creating a provisional “after” value f^{A*} . Equation (8b) applies the filter (with a strength parameter ϵ) with this value and the previous step’s provisional field to define a final “now” field, and finally equations (8c) and ~~including those pseudo-forces inside would double-count these effects~~(8d) are “bookkeeping” steps to shift field labels to become ready for the next timestep. The forcing operator \mathbf{F}^{N*} is evaluated based on the provisionally-defined fields.

210 In applying this filter with the semi-Lagrangian forcing, equation (7) is oblivious to the presence of the filter or the difference between f^N and f^{N*} . Substituting (7) into (8a) and applying (8c) and (8d) to (8a) and (8b) gives the update equation:

$$\begin{pmatrix} f^{N*} \\ f^B \end{pmatrix} \leftarrow \begin{pmatrix} f^B(\mathbf{x}^D) \\ \epsilon(f^B(\mathbf{x}^D) + f^B(\mathbf{x}_{ref})) + (1 - 2\epsilon)f^{N*}(\mathbf{x}_{ref}) \end{pmatrix}. \quad (9)$$

215 In the case of one-dimensional advection by a constant velocity u_0 , the trajectory calculation is trivial and:

$$\mathbf{x}^D = \mathbf{x}^D = \mathbf{x}_{ref} - 2\Delta t u_0. \quad (10)$$

Since (9) is linear, we can also take its Fourier decomposition in space and consider only a single, arbitrary wave mode, giving $f = \hat{f}(t) \exp i k x$ for a time-varying coefficient \hat{f} . Applying this to (9) casts the update in a matrix form as:

$$\begin{pmatrix} \hat{f}^{N*} \\ \hat{f}^B \end{pmatrix} \leftarrow \begin{pmatrix} 0 & \exp(-2i k u_0 \Delta t) \\ 1 - 2\epsilon & \epsilon(1 + \exp(-2i k u_0 \Delta t)) \end{pmatrix} \begin{pmatrix} \hat{f}^{N*} \\ \hat{f}^B \end{pmatrix}. \quad (11)$$

220 The time-stability of this filter is then governed by the eigenvalues of this matrix. Using the shorthand $\omega = -k u_0 \Delta t$, these eigenvalues are:

$$\lambda_{1,2} = \frac{1}{2} \left(\epsilon(1 + \exp(2i\omega)) \pm \sqrt{\epsilon^2(1 + \exp(2i\omega))^2 + (4 - 8\epsilon) \exp(2i\omega)} \right), \quad (12)$$

and to leading order in ϵ these eigenvalues have squared magnitudes of:

$$|\lambda_{1,2}|^2 = 1 - 2\epsilon \pm 2\epsilon \cos(\omega) + O(\epsilon^2), \quad (13)$$

225 signifying stability ($|\lambda| \leq 1$) for all values of ω and thus all Courant numbers.

3 Interpolation

To perform the off-grid interpolations in ~~to find $f(\mathbf{x}(t_0 - \Delta t))$ (abbreviated $f(\mathbf{x}_d)$)~~(7) to find $f^B(\mathbf{x}^D)$, this method fits a cubic polynomial to the underlying function. If the single departure point $\mathbf{x}_d = (x_d, y_d, z_d)$ lies within² $(x_a, x_{a+1}) \times (y_b, y_{b+1}) \times (z_c, z_{c+1})$ for integer values of a , b , and c coinciding with grid-point locations, the full interpolation stencil consists of the
 230 grid-index cube $i \in [a - 1, a + 2]$, $j \in [b - 1, b + 2]$ and $k \in [c - 1, c + 2]$.

This grid-cube contains up to 64 grid points where $f(\mathbf{x})$ might be defined (subject to boundary conditions), and building a complete interpolation stencil would be cumbersome and inefficient. Instead, the interpolation procedure takes advantage of the tensor-product nature of the grid to separate interpolation along each dimension:

Algorithm 1. *Three-dimensional interpolation. To find $f(x_d, y_d, z_d)$ for some off-grid point (x_d, y_d, z_d) :*

- 235
1. Interpolate $f(\mathbf{x})$ in the vertical to the location $[x_i, y_j, z_d]$, for $i \in [a - 1, a + 2]$ and $j \in [b - 1, b + 2]$
 2. Interpolate along the first dimension in this two-dimensional grid to give $f(x_d, y_j, z_d)$, for $j \in [b - 1, b + 2]$.
 3. Finally, interpolate along the second dimension to give $f(x_d, y_d, z_d)$.

To effect the one-dimensional interpolations in algorithm 1, we make use of the cubic Hermite polynomials (Hildebrand, 1974). On the interval $0 \leq \chi \leq 1$, these polynomials are:

$$\begin{aligned}
 h_{00}(\chi) &= 2\chi^3 - 3\chi^2 + 1, \\
 h_{01}(\chi) &= -2\chi^3 + 3\chi^2, \\
 h_{10}(\chi) &= \chi^3 - 2\chi^2 + \chi, \text{ and} \\
 h_{01}(\chi) &= \chi^3 - \chi^2,
 \end{aligned}
 \tag{14}$$

240

and a function $f(\chi)$ defined on this interval is interpolated via:

$$f(\chi) \approx f(0)h_{00}(\chi) + f'(0)h_{10}(\chi) + f(1)h_{01}(\chi) + f'(1)h_{11}(\chi).
 \tag{15}$$

Here, we prefer to use the cubic Hermite polynomials over simple Lagrange polynomial interpolation because the former choice allows greater freedom (via (15)) in implementation. If f' is approximated by a four-point finite difference stencil,
 245 then (15) reduces to Lagrange interpolation. However, we can also make other choices for f' to impose desirable properties: restricting f' to have the same sign as the discrete difference imposes a type of slope limiting, and calculating f' through a three-point stencil provides for continuous derivatives. These approaches are discussed in more detail in the following sections.

Interpolation using the above algorithm involves appropriately defining the interval to be scaled to $[0, 1]$ and approximating f' at the endpoints. Because of the high aspect ratio of oceanic flows and the special character of vertical motion in a stratified
 250 ocean, these approximations differ between the horizontal and vertical interpolations.

²If \mathbf{x}_d lies along an edge or corner of this interval, then at least one of the resulting interpolations will be trivial. In that case, the choice of which neighbouring interval \mathbf{x}_d lies “within” is arbitrary.

3.1 Horizontal interpolation

In the horizontal, the interpolation in (15) can be directly conducted in grid-index space. Even when the underlying grid is mapped to the sphere, such as in the ORCA global grid (Madec, 2008, Ch. 16), the grid generally transitions smoothly and slowly from point to point³. The physical trajectory departure location x_d (and y_d) can be translated into a fractional grid index
255 offset by dividing by the appropriate grid scale factor, available inside the NEMO-OPA source code as one of `e[12]` [`tuv`].

Achieving third-order accuracy inside (15) is possible, but doing so requires an equally-accurate estimate for f' . Unfortunately, interpolating successively in [1D-one dimension](#) using the above algorithm does not allow for precomputation of these derivatives: after the vertical interpolation step, all of the function values need to be taken off-grid, so any precomputed derivatives would themselves require interpolation. Instead, sufficiently-accurate estimates of the derivative are available by applying
260 a finite-difference formula to the function values themselves.

Derivative estimates

For notational simplicity, begin with the last step of the above algorithm where we have $f(x_d, y_j, z_d)$ and would like to estimate $f(x_d, y_d, z_d)$. If y_d lies between y_0 and y_1 , then the four-point interpolation stencil implies that we have computed $f(x_d, y_j, z_d)$ for $j = -1, 0, 1, 2$. To emphasize that this is now a one-dimensional interpolation problem, let $g(j) = f(x_d, y_j, z_d)$, such that
265 $f(x_d, y_d, z_d) = g(j')$ for some $j' \in [0, 1]$. In this domain, $g'(0)$ and $g'(1)$ can be approximated by the finite differences:

$$g'(0) \approx -\frac{1}{3}g(-1) - \frac{1}{2}g(0) + g(1) - \frac{1}{6}g(2) \text{ and} \quad (16a)$$

$$g'(1) \approx \frac{1}{6}g(-1) - g(0) + \frac{1}{2}g(1) + \frac{1}{3}g(2), \quad (16b)$$

which then substitute for the appropriate derivatives in (15).

These finite differences are exact expressions for the first derivative for polynomials up to third order in j , and their use
270 essentially converts (15) to interpolation via Lagrange polynomials. The Hermite polynomial form, however, allows for an easier imposition of boundary conditions.

Boundary conditions

On the NEMO-OPA z -level grid, the lateral boundaries coincide with u - and v -points (velocity points), which are spaced halfway between t -points (tracer points). Tracer points that lie inside the land region are masked (`tmask = 0`) as are velocity
275 points that are at the edge of or within the masked region. This arrangement is illustrated for a sample region in figure 2.

The physical interpretation of the boundary varies with respect to the field being interpolated. For tracers, lateral boundaries imply no-flux conditions for the purposes of advection, which in turn implies a zero derivative at the boundary. The normal velocity (u with respect to a boundary along the first grid dimension, v with respect to a grid boundary along the second) is obviously constrained to zero by geometry to give a Dirichlet boundary condition, whereas the tangential velocity can be set as

³This is not necessarily the case, however, for grids that have manually-specified, non-smooth regions of enhanced resolution. In such cases a more nuanced treatment of interpolation would be advisable.




Figure 2. Grid point locations (letters) and land region (grey region) for a sample horizontal plane in the Gulf of Saint Lawrence, between Nova Scotia and New Brunswick. The horizontal velocities (u and v) are staggered with respect to temperature and salinity (t), and the edge of the land area is coincident with the lines between velocity-point locations.

280 free-slip, no-slip, or some combination via a namelist entry. In the subsequent, we assume that velocity has a free-slip boundary condition, with boundary friction left for future work.

If a boundary occurs in the left portion of this interpolation stencil, there are a total of seven possible cases:

Algorithm 2. *Lateral boundary conditions*

To find $g(j')$ for $0 \leq j' \leq 1$ in the vicinity of a lateral boundary:

- 285
1. If $g(-1)$ corresponds to a point at the boundary and the boundary is of the Dirichlet-type, then $g(-1) = 0$ and (16a) and (15) apply normally.
 2. If $g(-1)$ is inside the boundary, $g(0)$ is inside the fluid domain (that is, the boundary is between these two points), and the boundary is of the Neumann-type, then $g(-1)$ is taken to be equal to $g(0)$, essentially making it a ghost point.

- 290 3. If $g(-1)$ is inside the boundary, $g(0)$ is inside the fluid domain, and the boundary is of the Dirichlet-type, then $g(-1)$ is taken to be $-g(0)$.
4. If $g(0)$ is at the boundary and the boundary is of the Dirichlet-type, then $g(0) = 0$ and $g(-1) = -g(1)$.
5. If $g(0)$ is inside the boundary and $j' < 0.5$, then the interpolated point is itself inside the boundary and should be masked.
6. If $g(0)$ is inside the boundary, $j' > 0.5$, and the boundary is of the Neumann-type, then $g(0) = g(1)$ and $g(-1) = g(2)$.
- 295 7. If $g(0)$ is inside the boundary, $j' > 0.5$, and the boundary is of the Dirichlet-type, then $g(0) = -g(1)$ and $g(-1) = -g(2)$.

For boundaries that occur in the right portion of the interpolation stencil, the values taken for ghost points are given symmetrically.

The combination of “the grid point is at the boundary” and “the boundary is of Neumann-type” is missing from algorithm 2. This construction is forbidden by the grid structure of NEMO-OPA, where tangential velocity is located one half-cell away from a boundary.

300

For two dimensional interpolation, algorithm 2 applies independently to each dimension. When interpolating along x , the points $f(x_d, y_j, z_d)$ will each individually be either in the ocean domain and valid or in the land domain and masked, which provides the values necessary to compute $f(x_d, y_d, z_d)$. This off-grid point itself must lie in water, which imposes a strong requirement on the trajectory calculations to be discussed in section 4.

305 **Slope limiting**

As a final step, once values for the function and its derivative at the interval endpoints are specified, the derivative values are limited to help prevent new maxima in the interpolated function. In particular, if $g(0)$ is a local minimum (maximum) among itself, $g(-1)$, and $g(1)$, then $g'(0)$ is set to zero if the above procedure finds that it would be negative (positive). A similar procedure applies symmetrically for $g'(1)$ if $g(1)$ is a relative extremum.

310 This limiting is milder than methods derived from Bermejo and Staniforth (1992), which would strictly preserve positivity for any j' , but it effectively limits excursions when j' is close to 0 or 1. Without such limiting, numerical testing showed that semi-Lagrangian advection of temperature and salinity could cause weak instabilities near the coastline, where a locally extreme temperature or salinity could become “trapped” near the coast and slowly amplified.

Two-dimensional application

315 ~~With the one-dimensional algorithm completely specified above, the interplay between steps 2 and 3 of Algorithm 1 is now clear. After step 2, interpolation along the first dimension, we have now specified $f(x_d, y_j, z_d)$ such that each point contains either an interpolated value or is found (by step 3 of Algorithm 1) to be inside the land region and hence masked. That provides the requisite four gridded function values to compute $f(x_d, y_d, z_d)$.~~

320 Generally, this value should *not* be inside land, as by this is the value that forms the advective trend of a presumably inside-water grid point. This imposes an obligation for the trajectory calculations in section 4 to return only valid departure points, and this requirement will be discussed in more length there.

3.2 Vertical interpolation

Vertical motion in the NEMO-OPA model differs from horizontal motion in a number of respects:

- 325 – Vertical gradients of temperature and density are much stronger than typical horizontal gradients, especially near the surface.
- Typical vertical grids used with NEMO-OPA are strongly stretched, with a higher resolution near the surface and a lower resolution in the deep ocean.
- Vertical flow is often oscillatory, where vertical motion is driven by barotropic and baroclinic waves.

330 ~~Although the~~ The horizontal interpolation described in section 3.1 is third-order accurate, ~~it is ill-suited for interpolation in the vertical. The worst characteristic of the horizontal interpolation is that it is~~; with the provided one-sided formulas for calculating the endpoint derivatives it reduces to a four-point (cubic) Lagrangian interpolation process. However, the smooth field implied by this interpolation process is only C^0 continuous, ~~where the interpolating function generally has a discontinuous derivative at the gridpoints. This is reasonable for horizontal flow that does not often change direction, but with the more-oscillatory vertical flow C^0 continuity causes~~: $f(x_j - \epsilon)$ “sees” f'_j calculated from $f(x_{j-2})$ to $f(x_{j+1})$, whereas
335 $f(x_j + \epsilon)$ sees f'_j from $f(x_{j-1})$ to $f(x_{j+2})$.

We do not find this to be a practical concern for horizontal interpolation, since horizontal currents in most of the ocean tend to be dominated by relatively steady quasi-geostrophic motions. In the vertical, however, we found that even low-amplitude oscillations caused by high-frequency gravity waves would cause the temperature and salinity fields to drift. ~~In particular, The mechanism is that~~ a fluid parcel ~~that is~~ displaced upwards by ϵ in one timestep and downwards by ϵ ~~on the next would see a net change of temperature and salinity proportional to ϵ times~~ in the next timestep would see an effective diffusion proportional to the jump between the upward and ~~downward-directed derivatives. This drift does not just remain localized; it causes larger-scale impacts through the generation of baroclinic waves.~~ downward-looking vertical derivatives.

~~Maintaining global accuracy requires that the interpolation be~~ To maintain global accuracy, we impose C^1 continuous continuity in the vertical. ~~Fortunately, the Hermite polynomial interpolation form is already equipped to handle this scenario.~~
345 ~~For the purposes of the first step of algorithm 1, the relevant coordinate system is physical depth, with $z = 0$ at the surface and $z = z_{max}$ at the local ocean bottom.~~ direction through an alternative treatment of the vertical derivative. Instead of applying equations (16), we treat the physical depth (rather than grid index) as the relevant coordinate and construct a centered estimate of the derivative.

350 ~~Unlike for horizontal interpolation, where derivatives at the interval endpoints are defined in grid-coordinate space, vertical derivatives at the gridpoints are defined in the physical space using a three-point stencil that reduces to the classic centered~~

~~difference when the grid spacing is uniform.~~ For a function $f(z_n)$ defined at the z_n levels, define $\Delta f^+ = f(z_{n+1}) - f(z_n)$, $\Delta f^- = f(z_n) - f(z_{n-1})$, $\Delta z^+ = z_{n+1} - z_n$, and $\Delta z^- = z_n - z_{n-1}$. These differences combine to give ~~an estimate for the the~~ estimated derivative:

$$f_z(z_n) \approx \frac{1}{\Delta z^- + \Delta z^+} \left(\frac{\Delta z^-}{\Delta z^+} \Delta f^+ + \frac{\Delta z^+}{\Delta z^-} \Delta f^- \right), \quad (17)$$

355 which is accurate to $O(\Delta z^2)$ for the derivative and accurately reproduces quadratic functions of z . In the limiting case of a constant Δz (equispaced vertical levels), this formula reduces to the classic centered difference.

Because vertical interpolation comes first in algorithm 1, (17) need be evaluated only at grid points, and in fact it may be precomputed for the entire grid for a given function and timestep. This is a key advantage of placing vertical interpolation first in the interpolation sequence, and it avoids duplication of work.

360 Whereas interpolation near the horizontal boundaries is complicated by the many combinations of grid staggering and physical boundary conditions, interpolation near the vertical boundaries is much simpler. On the NEMO grid, tracers and horizontal velocities lie along the same vertical level, and these levels are staggered one half-cell away from the boundaries. Likewise, the natural vertical boundary condition for both tracer and horizontal velocity fields is a no-flux boundary condition; NEMO-OPA models boundary-layer friction in another module. Interpolation near the boundaries then proceeds in two steps.

365 The first step is to define f_z at the top and bottom points in the water column, for which the central difference formula of (17) is not directly valid. Here, we approximate the physical no-flux condition through a fictitious ghost point such that $\Delta f^- = 0$ at the top boundary and $\Delta f^+ = 0$ at the bottom boundary, with the respective Δz matching the layer thickness ($e3t$).

The second step is to define how (15) applies to the interval between the grid level and the physical boundary. Here, the no-flux boundary conditions reduce to even symmetry, and the derivative at the ghost points is the negative of the vertical
370 derivative calculated for the in-boundary point. Near the free surface, if the interpolation point is above the level of the free surface (above $z = 0$) then it is clamped to the surface itself. Near the ocean bottom, if the interpolation point is below the level of the ocean bottom (below $z = z_{max}$) then the point is masked and is treated as an “inside the boundary” point for the purposes of horizontal interpolation above.

Treatment of partial cells

375 Over most of the domain, this interpolation works well. Although there is no guarantee of positivity in the derivative formulation of (17), overshoots and the consequent generation of spurious maxima are limited. For the tests presented in section 5, there was no need for slope-limiting for vertical interpolation over most of the domain.

One exception to this rule is at the bottom boundary. Here, vertical levels are spaced far apart, but to better-represent the ocean bottom the z-level grid of NEMO-OPA uses a partial cell configuration (Madec, 2008, sec. 5.9). For water columns
380 where the bottommost bottom-most cell is much deeper than its neighbours, a local (small) upwelling can cause an overshoot of temperature or salinity that spuriously increases the local density but does not diminish the upwelling. Over time, the maxima-increasing trend can accumulate and cause some points at the bottom boundary to reach implausibly cold temperatures (below -10°C , for example) or high salinities. In the absence of explicit horizontal diffusion (which would mix this maximum into

more dynamically-active regions), these spurious maxima do not generally corrupt the flow, although they obviously would
385 corrupt whole-ocean (or whole-level) statistics such as average or extreme temperatures.

Near these boundary cells, vertical limiting is implemented in the simplest possible way: the interpolation of (15) is replaced with a constant, such that $f(z) = f(z_k)$ over the ~~whole~~ interval from z_k downwards to the physical boundary.

Implementing this limiting over the whole bottom level is possible, but that is far stronger than necessary and leads to erroneous diffusion along gentle slopes. ~~Imagine a horizontally nondivergent flow, such that the vertical velocity $w = 0$ everywhere, with salinity and temperature being given by a function of z alone such that the underlying stratification is stable. Here, horizontal advection should nowhere result in a trend in salinity or temperature, and effecting this (even approximately) in the framework of semi-Lagrangian advection requires extrapolating tracer values downwards towards the boundary. When the bottom layer is composed of partial cells of varying thickness, even interpolation along a horizontal plane (that is, without changing physical depth) requires vertical interpolation in grid space to find that constant level in adjoining columns. Imposing~~
390 ~~vertical limiting along the whole bottom level effects undesired horizontal diffusion, even though the problem solved by limiting is observed when adjoining cells have large relative thickness variations.~~

As a compromise, ~~for this work we only limit vertical interpolation as described for cells that are~~ between these two errors, we only apply the described limiting to vertical interpolation for cells at the bottom boundary ~~and which~~ have a layer thickness greater than 1.75 times that of the neighbouring cell with the smallest local layer thickness ~~their “thinnest” neighbour.~~

400 This exact threshold is empirical, and other grids might require a re-tuning of this parameter. Ideally, the grid generation itself would avoid abrupt transitions in cell-layer thicknesses, but adding such a restriction would make this advection scheme useless as a drop-in replacement for the standard advection routines of NEMO-OPA.

3.3 A numerical example

As a simple numerical example, consider the case of a tracer being advected in a rectangular, two-dimensional domain by an
405 internal wave and a background current. This tracer satisfies the advection equation:

$$\frac{\partial \sigma}{\partial t} - u(x, z, t) \frac{\partial \sigma}{\partial x} - w(x, z, t) \frac{\partial \sigma}{\partial z} = 0, \quad (18)$$

for some prescribed velocity field (u, w) .

If this tracer field $\sigma(x, z, t)$ would be a function of z alone ($\bar{\sigma}(z)$) if not for the wave motion, then its motion is analytically given by:

$$410 \quad \sigma(x, z, t) = \bar{\sigma}(z - \eta(x - (c + u_0)t, z)), \quad (19)$$

where $\eta(x, z)$ is the isopycnal displacement, u_0 is the x -directed background current (uniform in z), and c is the phase speed of the wave. ~~A~~ Following Turkington et al. (1991), a streamfunction defined as:

$$\psi(x, z, t) = c\eta(x, z, t) - u_0 z \quad (20)$$

gives velocities:

$$415 \quad u = -\psi_z \tag{21a}$$

$$w = \psi_x, \tag{21b}$$

which are exact solutions of (18) for $\sigma(x, z, t)$.

To give an internal wave that respects no-flux conditions at the top and bottom of the domain, we set:

$$\eta(x, z, t) = A \cos(k(x - (c + u_0)t)) \sin(mz) \tag{22}$$

420 where k and m are horizontal and vertical wavenumbers respectively and A is the wave amplitude. For a domain of size L_x in the horizontal (periodic) and L_z in the vertical, $k = 2\pi/L_x$ and $m = \pi/L_z$ give the lowest internal wave mode, used here.

In dimensional units, we take the model domain to be a channel $L_x = 1\text{km}$ long and $L_z = 100\text{m}$ deep with a background current of $u_0 = 1\text{ms}^{-1}$, and we set $c = N/\sqrt{k^2 + m^2}$ based on a mean buoyancy frequency of $N = .03\text{s}^{-1}$, which corresponds to a 1% density change from the surface to the bottom of the channel. With a wave amplitude of $A = 10\text{m}$, the maximum wave-
425 induced current is about 10% of u_0 , and the phase speed is $c \approx 0.94\text{ms}^{-1}$.

In order to represent the pycnocline found in many ocean waters, we choose⁴ $\bar{\sigma}(z) = \tanh\left(\frac{1}{2} - \frac{1}{10}zL_y^{-1}\right)$.

The domain is discretized by $N_x \times N_y$ points, defined as:

$$x_i = -\frac{L_x}{2} + L_x \frac{i - 0.5}{N_x}, \text{ and} \tag{23a}$$

$$z_j = \frac{L_z}{2} \left(1 + \frac{\alpha_j + \alpha_j^3}{2}\right), \tag{23b}$$

430 where $i = 1, 2, \dots, N_x$, $j = 1, 2, \dots, N_z$, and:

$$\alpha_j = 2^{\frac{j - 0.5}{N_z}} - 1. \tag{23c}$$

This implements a stretched vertical coordinate that increases the vertical resolution in the vicinity of the pycnocline.

Semi-Lagrangian advection

In integrating this system with semi-Lagrangian advection, the leapfrog method reduces to an Euler method of twice the
435 timestep because there is no external forcing. The time-discrete equation is:

$$\sigma(x_i, z_j, t + 2\Delta t) = \sigma(x_{d(ij)}, z_{d(ij)}, t), \tag{24}$$

where $(x_{d(ij)}, z_{d(ij)})$ is the departure point of the trajectory that arrives at the gridpoint (x_i, z_j) , and the off-grid evaluation of σ proceeds via the interpolation processes described earlier without slope-limiting.

⁴Since this section tests advection alone, the scaling of σ is not dynamically relevant. In fact, the wave structure of (22) corresponds to an exact internal mode of the incompressible Navier-Stokes equations for a linear stratification.

The departure points are given by the trapezoidal rule⁵ with a time-centered evaluation of velocity:

$$440 \quad x_i - x_{d(ij)} = \Delta t (u(x_i, z_j, t + \Delta t) + u(x_{d(ij)}, z_{d(ij)}, t + \Delta t)) \quad (25a)$$

$$z_j - z_{d(ij)} = \Delta t (w(x_i, z_j, t + \Delta t) + w(x_{d(ij)}, z_{d(ij)}, t + \Delta t)), \quad (25b)$$

where the velocities are evaluated exactly via (21). The overall system (25) is solved via simple iteration, with an initial guess given by setting $(x, z)_{d(ij)} = (x, z)_{ij}$.

This algorithm is stable for large timesteps, so we tested this system for timesteps corresponding to CFL-Courant numbers
445 of 0.2 and 2.1, with spatial grid resolutions between 40×4 and 2560×256 . The final integration time was chosen to be $t_{fin} = 5L_x/u_0$, which allowed the wave to propagate through the domain several times. Since the exact solution is analytically known, we recorded the maximum error experienced over the integration, and error convergence rates are shown in figure 3.

Flux-form advection

As a control, we also integrate this system in flux form ($\sigma_t - \nabla \cdot (\mathbf{u}\sigma) = 0$) via centered differences, with σ evaluated at the
450 midpoints between grid cells via a simple average, matching the central difference tracer advection scheme in NEMO-OPA.

The velocity field given by (20) is divergence free, so this form of the equation is pointwise equivalent to (18). However, this no longer holds after discretization. In order to eliminate the divergence error, the velocity field is defined by creating the streamfunction at the staggered points $(x_{i+1/2}, z_{j+1/2})$ and defining discrete velocities u and w via the discrete equivalents to (21). With this modification, the discrete flux-form operator is equivalent to a discrete advection equation.

455 After leapfrog discretization in time, the discretized equation is:

$$\begin{aligned} \sigma(x_i, z_j, t + \Delta t) = & \sigma(x_i, z_j, t - \Delta t) + 2 \frac{\Delta t}{\Delta x \Delta z_j} \left(\frac{\Delta z_j}{2} (u(x_{i-1/2}, z_j, t)(\sigma(x_{i-1}, z_j, t) + \sigma(x_i, z_j, t)) - \right. \\ & \left. u(x_{i+1/2}, z_j, t)(\sigma(x_i, z_j, t) + \sigma(x_{i+1}, z_j, t))) + \right. \\ & \left. \frac{\Delta x}{2} (w(x_i, z_{j-1/2}, t)(\sigma(x_i, z_{j-1}, t) + \sigma(x_i, z_j, t)) - \right. \\ & \left. w(x_i, z_{j+1/2}, t)(\sigma(x_i, z_j, t) + \sigma(x_i, z_{j+1}, t))) \right), \end{aligned} \quad (26)$$

where $\Delta z_j = z_{j+1/2} - z_{j-1/2} = \frac{1}{2}(z_{j+1} - z_{j-1})$. For the first timestep, a single Euler step is taken of size Δt with time-centered velocities ($t = \Delta t/2$).

As usual, this leapfrog timestepping algorithm is only stable to a CFL-number of $\Delta t \max(u)/\Delta x < 1$, so we tested maximum
460 Courant number of 1. With this staggered grid and vertical grid stretching, the Courant number can be defined by:

$$C_{ij} = \Delta t \left(\frac{\max(u_{i+1/2, j}, 0) - \min(u_{i-1/2, j}, 0)}{\Delta x} + \frac{\max(w_{i, j+1/2}, 0) - \min(w_{i, j-1/2}, 0)}{\Delta z_j} \right). \quad (27)$$

For the mode-one internal wave with background current used in this section, the maximum Courant number is reached at the top and bottom of the domain (where $w = 0$), so $\max(C) = \max(u)/\Delta x$.

⁵The trajectory calculation scheme of section 4 could be used instead, but since the overall trajectory lengths are small compared to the length scales of the velocity field (L_x and L_z), that method would give equivalent results to the trapezoidal rule.




Figure 3. Top: Maximum L2 error ($(\int (\sigma - \sigma_{e,x})^2 dA / (L_x L_z))^{1/2}$) over the integration period for the test case of section 3.3 versus resolution for flux-form Eulerian advection with a CFL-Courant number of 0.2 (blue, solid), semi-Lagrangian advection with a CFL-Courant number of 2.1 (red, dashed), and semi-Lagrangian advection with a CFL-Courant number of 0.2 (green, dot-dashed), showing second-order convergence (line). Bottom: L2 error over time for these algorithms, on the 2560×256 grid.

465 We present results for (26) only at a CFL at a maximum Courant number of 0.2, chosen to give a “small timestep” for later comparison with semi-Lagrangian results. The results are insensitive to the timestep within the stable range, with less than 5% change in maximum-norm error over the range $0.2 \leq \max(C) \leq 0.99$.

Results

The error over time of this test case is shown in figure 3. As expected, each method achieves second-order convergence. For the Eulerian advection control case, this is governed by its two-point central difference scheme. For the semi-Lagrangian cases, the dominant contribution to the error field comes from the lower-order vertical interpolation. While the semi-Lagrangian method has a higher order of accuracy for horizontal motion, here the problem is constructed such that horizontal and vertical motions are of equal importance.

470

As is often observed with semi-Lagrangian methods, the overall error of the scheme is somewhat lower for the high-CFL case than for the low-CFL case. The interpolation used to evaluate σ off the grid necessarily introduces error with each interpolation,

475 and the overall contribution of this error necessarily scales in proportion to the number of interpolations ~~-, and hence-~~ and
inversely with the timestep size.

Overall, this simplified test case supports the conclusion that the semi-Lagrangian treatment of advection is a viable re-
placement for flux-form advection. The semi-Lagrangian method achieves similar (for low-CFL flows) or better (for high-CFL
flows) error, and it remains stable for CFL values substantially larger than unity.

480 4 Trajectory calculation

With the mechanism for evaluating the ~~$f^{t_0-\Delta t}(\mathbf{x}(t_0-\Delta t))$ term in~~ $f^B(\mathbf{x}^D)$ term in (7) established in section 3, the remaining
half of the semi-Lagrangian advection algorithm is the estimation of the ~~$\mathbf{x}(t_0-\Delta t)$ departure points~~ (again abbreviated \mathbf{x}_d) \mathbf{x}^D
departure points. This corresponds to the positions at the “before” timestep time-level ($t_0 - \Delta t$) of those fluid parcels that will
arrive ~~at the grid points at on the reference grid \mathbf{x}_{ref}~~ the “after” timestep time-level ($t_0 + \Delta t$). One such upstream location
485 exists for each valid grid location, so in general \mathbf{x}_d needs to be estimated for each τ , u , and v point on the NEMO-OPA grid to
provide for (respectively) the tracer and velocity advective forcings.

In general, calculation of the departure points is an implicit and nonlinear problem, requiring knowledge of the flow velocity
at every sub-grid place and time between the “before” and “after” time-levels, before the flow at the latter has been computed.
To make this problem tractable, we make a series of simplifying assumptions.

490 The first such assumption is to freeze the flow, such that trajectories are computed based on strictly the “now” velocities
(that is, u , v , and w at the intermediate time-level). This is consistent with the underlying leapfrog timestepping algorithm
and the other advection schemes in NEMO-OPA, where most fluxes are computed instantaneously with respect to the same
“now” velocities. In physical terms, this constrains fluid parcels to follow paths based on estimated, instantaneous streamlines.
In exchange, this decouples the trajectory computation from the “after” velocities and makes the process time-explicit, which
495 eliminates what would otherwise be a need to iterate the entire timestepping process.

4.1 Exponential integration

Ordinarily, the next assumption in the trajectory calculation is to approximate the particle paths, either by a straight line or by
a low-degree polynomial. In this case, the Lagrangian equation:

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}(\mathbf{x}) \tag{28}$$

500 is integrated with an approximate quadrature. Using the trapezoidal rule gives the approximation:

$$\mathbf{x}_a - \mathbf{x}_d = \Delta t(\mathbf{v}(\mathbf{x}_d) + \mathbf{v}(\mathbf{x}_a)), \tag{29}$$

where $\mathbf{x}_a = \mathbf{x}(t_0 + \Delta t)$ is the on-grid “arrival” point of the fluid parcel. This approximation is second-order in time, and it
results in an iterative method where $\mathbf{v}(\mathbf{x}_d)$ is interpolated, leading to a revised estimate of \mathbf{x}_d .

Unfortunately, this approximation is not suitable for trajectory calculations in the general ocean because it does not ap-
505 propriately handle flow near a solid boundary. Consider the case of two-dimensional flow in the positive half-plane, where

fluid velocities are prescribed as $(u, v) = (x, -y)$, ~~and examine the departure point corresponding to a.~~ This forms an analytic continuation of flow near a boundary along the line $x = 0$.

Now, apply equation (29) to the fluid parcel that arrives at $x_a = (1, 0)$. ~~Here, the y -directed velocity along the streamline is zero by inspection (and so y_d is also zero), so reduces to a one-dimensional equation with~~ Along this streamline, $v = 0$ by
510 inspection so this equation reduces to one dimension and has the solution:

$$x_d = \frac{1 - \Delta t}{1 + \Delta t}. \quad (30)$$

For small values of Δt , this solution is ~~very reasonable. Unfortunately, as Δt becomes large, gives unphysical results with respect to the boundary. Once~~ reasonable. For $\Delta t > 1$, the solved-for departure point however, this solution leads an unphysical trajectory, where the departure point is found to lie in the left-half plane (and thus lies inside the boundary, where in reality
515 there is no fluid to advect!).

The ~~exact form of depends on the modeling of the boundary, but there is ultimately no universally-valid answer. If the departure point is to lie within the physical domain, then failure here is a specific example of (29) either must depend only on physically-realized velocities (independent of any particular extension into the boundary) or the departure point cannot be an equilibrium solution, failing the Lipschitz trajectory-crossing criterion (Smolarkiewicz and Pudykiewicz, 1992), which~~
520 requires $u_a \Delta t < C \approx 1$. The trajectory implied by (30) crosses the trajectories of fluid parcels that arrive at $x_a = (1 \pm \epsilon, 0)$, and the resulting advection loses its physical meaning.

~~A~~ This trajectory-crossing criterion is a physical limit for solutions which develop discontinuous shocks, such as those that can arise in simulations of the non-dispersive, nonlinear shallow water equations. However, these shocks are not typical of three-dimensional hydrostatic flows in the ocean, and they are certainly not universally seen at solid boundaries. The true
525 trajectories of fluid parcels, if evaluated exactly, do not cross (and do not have origins inside the land domain), so a better approach is to directly integrate (28) without approximating the time derivative. Here, this one-dimensional system reduces to the ordinary differential equation:

$$x_t = x \quad (31)$$

with the boundary condition $x(t_0 + \Delta t) = x_a = 1$. The solution to this equation is obviously of the form $x(t) = C \exp(t)$
530 for some constant C , and applying the boundary condition gives $x(t) = \exp(t - (t_0 + \Delta t))$ and a departure point of $x_d = \exp(-2\Delta t)$.

This solution is very well-behaved, lying exclusively in the right half-plane and asymptotically approaching the wall at $x = 0$ as $\Delta t \rightarrow \infty$. This approach works when that of (29) fails because the direct integration properly captures the exponential-in-time path of the fluid parcel.

535 A generalization of this approach forms the basis for trajectory calculation in this work. Since the solution of (28) is not analytically possible with an arbitrary velocity field, we exactly solve (28) based on an approximate, linearly-varying velocity field. This is similar to an approach discussed by Walters et al. (2007), where within a single, two-dimensional finite-element cell the linear velocity form is exactly-given by the underlying discretization rather than an approximation.

Assume that an arbitrary fluid parcel arrives at \mathbf{x}_a , and that we know the velocity there (\mathbf{v}_a) and at another point $\mathbf{v}(\mathbf{x}_c) = \mathbf{v}_c$.
 540 We know that the fluid parcel must arrive at \mathbf{x}_a travelling in the direction of $\hat{\mathbf{v}}_a = \mathbf{v}_a/\alpha_a$, with $\alpha_a = \|\mathbf{v}_a\|$. Then \mathbf{v}_c can be written in terms of this direction as $\mathbf{v}_c = \alpha_c \hat{\mathbf{v}}_a + \beta_c \hat{\mathbf{n}}_a$, for scalar α_c and β_c and some $\hat{\mathbf{n}}_a$ normal to $\hat{\mathbf{v}}_a$.

This forms a two-dimensional system spanned by vectors $\hat{\mathbf{v}}_a$ and $\hat{\mathbf{n}}_a$. If we additionally make the assumption that $\mathbf{v}(\mathbf{x})$ varies linearly in this plane, we can construct a simplified, two-dimensional coordinate system to solve (28). Here, the origin of the coordinate system corresponds to \mathbf{x}_a , and the rotated coordinates \hat{x} and \hat{y} align with $\hat{\mathbf{v}}_a$ and $\hat{\mathbf{n}}_a$ respectively. This implies
 545 that \mathbf{x}_c projects onto the point $(\mathbf{x}_c \cdot \hat{\mathbf{v}}_a, \mathbf{x}_c \cdot \hat{\mathbf{n}}_a) = (x_c, y_c)$. The linearly-interpolated velocities lie strictly in this plane, so the equations of motion for a fluid parcel are:

$$x_t = \alpha_a + (\alpha_c - \alpha_a) \frac{x}{x_c}, \text{ and} \quad (32a)$$

$$y_t = \beta_c \frac{y}{x_c}, \quad (32b)$$

subject to the boundary condition that $x(t_0 + \Delta t) = y(t_0 + \Delta t) = 0$. (32a) can be solved first, and applying the boundary
 550 condition $x(t_0 + \Delta t) = 0$ gives:

$$x(t) = \frac{\alpha_a x_c}{\alpha_c - \alpha_a} \left(\exp\left(\frac{\alpha_c - \alpha_a}{x_c} (t - (t_0 + \Delta t))\right) - 1 \right). \quad (33a)$$

Applying this to (32b) along with its boundary condition $y(t_0 + \Delta t) = 0$ gives:

$$y(t) = \frac{\beta_c \alpha_a}{\alpha_c - \alpha_a} \left(\frac{x_c}{\alpha_c - \alpha_a} \left(\exp\left(\frac{\alpha_c - \alpha_a}{x_c} (t - (t_0 + \Delta t))\right) - 1 \right) - (t - (t_0 + \Delta t)) \right). \quad (33b)$$

When the along-trajectory acceleration is small ($|(\alpha_c - \alpha_a)\Delta t/x_c| \ll 1$), (33) reduces to a trapezoidal rule with second-order
 555 accuracy in time.

Trajectory iteration

Evaluating (33) at $t = t_0 - \Delta t$ and re-projecting the coordinates to the grid forms the basis of an iterative algorithm for trajectories:

Algorithm 3. Trajectory iteration overview

560 At each grid point:

1. Begin with a candidate departure point $\mathbf{x}_c = \mathbf{x}_a - 2\Delta t \mathbf{v}_a$
2. Interpolate the “now” velocities off-grid to this point
3. Evaluate (33) at $t = t_0 - \Delta t$ to give a revised candidate departure point \mathbf{x}'_c
4. Set $\mathbf{x}_c \leftarrow \mathbf{x}'_c$
5. Repeat from step 2 until the change is smaller than a tolerance of 10^{-3} grid cells

565

This algorithm is ideally suited to cases that look like flow away from a stagnation point, where a fluid parcel is accelerating as it reaches the grid point at $t_0 + \Delta t$. In those cases, the $(\alpha_c - \alpha_a)/x_c$ terms will be positive, and the exponential terms will limit the size of the trajectory for finite Δt . In the opposite case, however, the exponential terms will tend to lengthen the trajectory. For large Δt or large deceleration, this effectively demands that (32)–(33) extrapolate beyond the velocity sample at \mathbf{x}_c , a potential source of instability.

To remedy this, a limiter is added to step 3 of algorithm 3, whereby $x(t_0 - 2\Delta t)$ is constrained to the greater⁶ of that from (33a) and $-2\Delta t \max(\alpha_a, \alpha_c)$. When limiting is necessary it effectively reduces the timestep used for the trajectory iteration, so for consistency a revised $\Delta t'$ is computed by inverting (33a) with the limited x'_c , which is then used to evaluate (33b).

4.2 Underrelaxation and land boundaries

While the construction of algorithm 3 guarantees that trajectories cannot converge to an out-of-boundary point, there are no guarantees that the algorithm remains in-boundary during the iteration process or that the iteration converges. The problem of a divergent or oscillatory iteration is more likely when the underlying velocity field does not resemble the linearly-approximated velocity field integrated by (32), as then each iteration might result in very different approximations.

Addressing the latter point first, this work heuristically applies underrelaxation when algorithm 3 is slow to converge. After 10 local iterations, step 3 is replaced by $\mathbf{x}_c \leftarrow \frac{1}{2}(\mathbf{x}_c + \mathbf{x}'_c)$, after 20 iterations the right-hand side becomes $\frac{1}{4}(3\mathbf{x}_c + \mathbf{x}'_c)$, and after 30 iterations the right-hand side becomes $\frac{1}{8}(7\mathbf{x}_c + \mathbf{x}'_c)$. At 40 iterations, the trajectory is truncated by ending the iteration with the first in-domain point returned from the process; this ensures some sort of advection even if the iterative process enters a limit cycle.

This underrelaxation also addresses the possibility that \mathbf{x}_c might lie outside of the ocean domain. If \mathbf{x}_c is masked, then there is no valid velocity to provide via off-grid interpolation, so instead of evaluating (33) \mathbf{x}'_c is set to \mathbf{x}_a in step 3 of algorithm 3. This combines with the underrelaxation after 10 iterations to reduce the trajectory length until an in-boundary point is found, whereupon iteration resumes normally.

These values for iteration counts and underrelaxation parameters are conservatively specified. In the numerical tests discussed in this work, the vast majority of trajectories converge after one or two iterations, without needing to resort to underrelaxation or trajectory truncation.

4.3 Velocity interpolation

The trajectory algorithm requires the off-grid interpolation of velocities at each iteration. In principle, these velocities can be interpolated using the interpolation process of section 3. Doing so would be ideal for ultimate consistency with the final off-grid interpolation, but this process is also computationally expensive. In practice, it is more efficient to evaluate the off-grid velocity field in step 2 of algorithm 3 using trilinear interpolation; doing so causes little change in the numerical test cases in this work.

Trilinear interpolation proceeds with the same order of operations as algorithm (1): velocities are first interpolated in depth to the (x, y) corners of the grid-box at the off-grid level, then along the x -direction, and finally along the y -direction. Each individ-

⁶Since the rotated x -axis is aligned with the fluid velocity at \mathbf{x}_a , x_c is generally negative in the rotated frame.




Figure 4. Illustration of modified linear interpolation near corners. At top, linear interpolation results in an interpolated velocity field that does not respect the boundary conditions along $0 < y < 0.5$ and a discontinuous interpolated field at $0.5 < y < 1$. At bottom, modifying the linear interpolation with a corner solution results in a field that respects the boundary condition.

ual interpolation respects the relevant boundary condition, so for example the u -velocity is considered to reflect symmetrically around a boundary in y and z but is constrained to zero at a boundary in x .

600 One complication of linear interpolation, however, is that the velocity points are staggered by half a cell with respect to the physical boundary. In two dimensions, if the tracer point $T(0,0)$ (to use grid-cell coordinates for the tracer grid denoted T) is a land point but $T(0,1)$, $T(1,0)$, and $T(1,1)$ are all ocean points, then u -velocity point $U(0,0)$ (denoting the u -velocity grid as U), halfway between $T(0,0)$ and $T(1,0)$, lies along the boundary. The boundary continues to $U(0,0.5)$, whereupon $U(0,0.5)–U(0,1)$ lies inside the ocean. This violates a basic assumption of linear interpolation, that the velocity should vary
605 smoothly (and approximately linearly) within the u -cell.

This causes two problems for trajectory computation. The first problem is that after repeated one-dimensional interpolation, the boundary condition is no longer necessarily respected by the interpolated velocity, which can result in a trajectory iteration that “pushes” the departure point into the wall, causing non-convergence. The second problem is that while the interpolation process guarantees continuity of the interpolated field at the cell corners, the boundary conditions can cause large discontinuities
610 along the cell edges, again resulting in a convergence failure. In the above example, the interpolated velocity at $U(+\epsilon, 0.6)$

would be influenced by both $U(0, 1)$ and the zero velocity at the physical boundary of $U(0, 0)$, but the interpolated velocity at $U(-\epsilon, 0.6)$ would be influenced by $U(0, 1)$ and its reflection at a ghost point. These problems are illustrated in the top panel of figure 4.

The solution to both of these problems is to blend the linearly-interpolated function with a corner singularity solution. A bilinear function is a solution to Laplace's equation ($\nabla^2 f = 0$), so it is reasonable to consider corner solutions that are also solutions to Laplace's equation.

Without loss of generality, consider a grid cell defined by $(x, y) \in [0, 1]^2$, such that there is a solid boundary along $(x = 0, y < 0.5)$ as depicted in figure 4. Treating the boundary as an *infinite* half-plane, with $f(0, y) = 0$ for $y < 0.5$ and $f(0, y) = f(0, 1)$ for $y > 0.5$, the ‘‘corner’’ solution to Laplace's equation is:

$$f_{corner}(x, y) = \frac{f(0, 1)}{2} \left(1 + \cos \left(\tan^{-1} \left(\frac{y - 0.5}{x} \right) \right) \right), \quad (34)$$

while bilinear interpolation would give:

$$f_{bilinear}(x, y) = (1 - x)yf(1, 0) + x(1 - y)f(0, 1) + xyf(1, 1). \quad (35)$$

These two solutions are blended together, with (34) taking precedence along the solid boundary ($x = 0$ and $0 \leq y \leq 0.5$) and (35) taking precedence along the $x = 1$ and $y = 1$ boundaries of the cell. This gives:

$$f_{blend}(x, y) = \sigma(x, y)f_{bilinear}(x, y) + (1 - \sigma(x, y))f_{corner}(x, y), \quad (36)$$

where $\sigma = \max(1 - x, 2(y - 0.5))$.

The blended function exactly respects the solid boundary condition, and the discontinuity at the cell edges is significantly reduced. Blended functions for other configurations of the solid wall are given by applying the appropriate reflections and rotations to (36).

630 5 Results

5.1 Flow past an island

To demonstrate the impacts of semi-Lagrangian advection on a simple test case with a lengthened timestep, we first present the quasi-two dimensional test case of isothermal flow past an interposed island.

This test case consists of a $280 \times 70 \times 3$ point grid, with grid resolution $\Delta x = \Delta y = 5\text{m}$ and $\Delta z = 10\text{m}$. A $50\text{m} \times 50\text{m}$ region (10×10 points) is masked as land in the middle of the domain. The inflow boundary condition is set to $u = 0.03\text{m/s}$, $v = 0$; this was also imposed throughout the domain as an initial condition. The reference frame was also irrotational, with a Coriolis parameter of 0.

Relevant namelist parameters are given in table 1, with parameters that differ between the control and semi-Lagrangian runs highlighted. The control run used flux-form velocity advection⁷ via the QUICKEST scheme (Leonard, 1979, 1991), whereas

⁷This choice of velocity advection provided the best results for the control run, of the advection models supported in NEMO version 3.1.

Parameter	Value	Comments
rdt	<i>Varies</i>	Varied from 5s – 160s
nitend	<i>Varies</i>	Set so that $\text{rdt} * \text{nitend} = 8000\text{s}$
ln_zps	.TRUE.	Enables the z-level coordinate; no partial steps were necessary
atfp	0.1	Asselin time filter parameter
ln_dynvor_een	.FALSE.	Flux-form advection
ln_dynvol_qck	.TRUE.	QUICKEST velocity advection (for control run)
ahm0	0	Horizontal eddy viscosity for momentum
avm0	$1.2e-4$	Vertical eddy viscosity
ln_zdfevd	.TRUE.	Enhanced vertical diffusion
avevd	100	Vertical coefficient of enhanced diffusion
n_evdm	1	Apply enhanced vertical diffusion to momentum
nn_botfr	3	Free slip bottom boundary condition

Table 1. Selected namelist parameters for the test case of section 5.1.

640 the semi-Lagrangian run used semi-Lagrangian advection of momentum in flux form as described in sections 3 and 4. To
emphasize the dynamical differences between the advection schemes, both test cases were run with no explicit horizontal
diffusion of momentum. Vertical mixing terms, largely irrelevant for this quasi-two dimensional case, were set consistently
with the ORCA025 simulations in section 5.2.

Both series of runs used the implicit free surface formulation (enabled with the compile-time key `key_dynspg_flt`),
645 which damped the large initial surface gravity waves caused by the imposition of the blocking island on the steady-state flow.

After the initial gravity-wave adjustment, this test case quickly develops a set of recirculating vortices in the lee of the island.
Over time these vortices grow in extent and would begin detaching to form a vortex street, but this does not happen before the
8000s end of the simulation. Although there is no explicit horizontal diffusion of momentum in these runs, the flow regime is
much more laminar than would be implied by the physical Reynolds number of $1.5 \cdot 10^6$, based on the free-stream velocity,
650 edge-length of the island, and molecular viscosity of water.

In moving around the box, the flow locally accelerates to a maximum steady velocity of about 0.05m s^{-1} , and this maximum
velocity is reached in the vicinity of the leading-edge corners of the box. The exact value of this maximum depends on both the
simulation time and the timestep, but our expected pattern holds: the control simulation is stable with a timestep of 64 seconds,
which corresponded to a maximum steady ~~CFL-number~~ Courant number ($\max(u_{\text{steady}})/\Delta x$) above 0.6 (and a maximum
655 transient CFL of 0.95), but it is unstable with a timestep of 80 seconds.

Semi-Lagrangian advection maintains stability for much longer timesteps. Figure 5 shows the free surface height and flow
streamlines for Δt between 5 and 160 seconds, and only the semi-Lagrangian method remains stable for 80 and 160-second
timesteps. For both advection schemes, the longer timestep is associated with a more diffuse flow pattern, with lengthening
(and less intense) recirculating vortices in the lee of the island.

660 This effect is stronger with semi-Lagrangian advection than with Eulerian advection. We attribute this to the nature of the
flow at the leading edge of the island. Here, the dominant flow balance is cyclostrophic, where the pressure gradient at these
corners balances the local vorticity. The operator splitting method used here treats the advective terms in a frame following the

flow, but it can only apply the pressure force at the destination cell; ~~there is~~. This results in an inconsistency that grows with Δt . ~~This is~~, related to the forces in equation (2) being available only at the endpoint of the Lagrangian trajectory – an $O(\Delta t)$ approximation to the integral.

This inconsistency is most evident in the 160-second timestep case (bottom right panel of figure 5), where the maximum steady CFL-Courant number of 1.6 implies that fluid parcels are advected by about three grid cells over the $2\Delta t$ leapfrog step. There, the lowest pressure region at the leading edge of the flow has moved slightly further downstream.

In the full ocean, the geostrophic effect predominates, with a leading-order balance between the pressure gradient and the Coriolis force (planetary vorticity), so we expect this issue to be less pronounced.

5.2 Global forced runs

To evaluate semi-Lagrangian advection in a more realistic forecasting setting, we conducted a preliminary series of free runs of the NEMO-OPA model. The runs consisted of:

- A control run, based on the configuration of Environment and Climate Change Canada’s $1/4^\circ$ nominal-resolution Global Ice/Ocean Prediction System (GIOPS) (Smith et al., 2016) with a 10-minute model timestep⁸. Tracers were advected with the model’s ~~total variation diminishing (TVD) scheme~~ tracer variance dissipation scheme (Lévy et al., 2001), and momentums were advected in vector form with the model’s energy and enstrophy conserving scheme⁹ (Arakawa and Lamb, 1981),
- A “semi-Lagrangian tracer” run, where momentum was advected as in the control scheme and the semi-Lagrangian advection described in this work was used for advection of salinity and temperature. Additionally, this run disabled horizontal diffusion of salinity and temperature, and
- A “semi-Lagrangian momentum and tracer” run, where momentum as well is advected with the semi-Lagrangian scheme. The configuration was otherwise the same as the semi-Lagrangian tracer run, save for a 15-minute model timestep.

The runs were all initialized at October 1, 2001 on the ORCA025 grid. The ocean was at rest, and temperature and salinity were given by the 2011 World Ocean Atlas climatology (Locarnini et al., 2013; Zweng et al., 2013). Atmospheric forcing was provided at one-hour intervals from Environment and Climate Change Canada’s $1/4^\circ$ global atmospheric reforecast, and sea ice was modeled via coupling with version 4.0 of the CICE model (Hunke and Dukowicz, 1997), with dynamically active (moving) ice. Selected namelist parameters are provided in table 2.

⁸This timestep is shorter than other commonly-used ORCA025 configurations, such as in the ocean reanalysis of Ferry et al. (2016). This shorter timestep is required to stabilize the coupling of ocean/ice stress with the CICE model, where following Roy et al. (2015) the ice/ocean drag coefficient is larger than typically considered. We chose to maintain this configuration and coupling approach to provide for the cleanest like-for-like comparisons against the operational configuration

⁹For compatibility with the operational model, as run in this work the scheme did not include the “fix” for the Hollingsworth instability (Hollingsworth et al., 1983) reported in Ducouso et al. (2017). This instability is more prominent at higher resolutions, and we do not believe it meaningfully impacted the results as presented in this section.

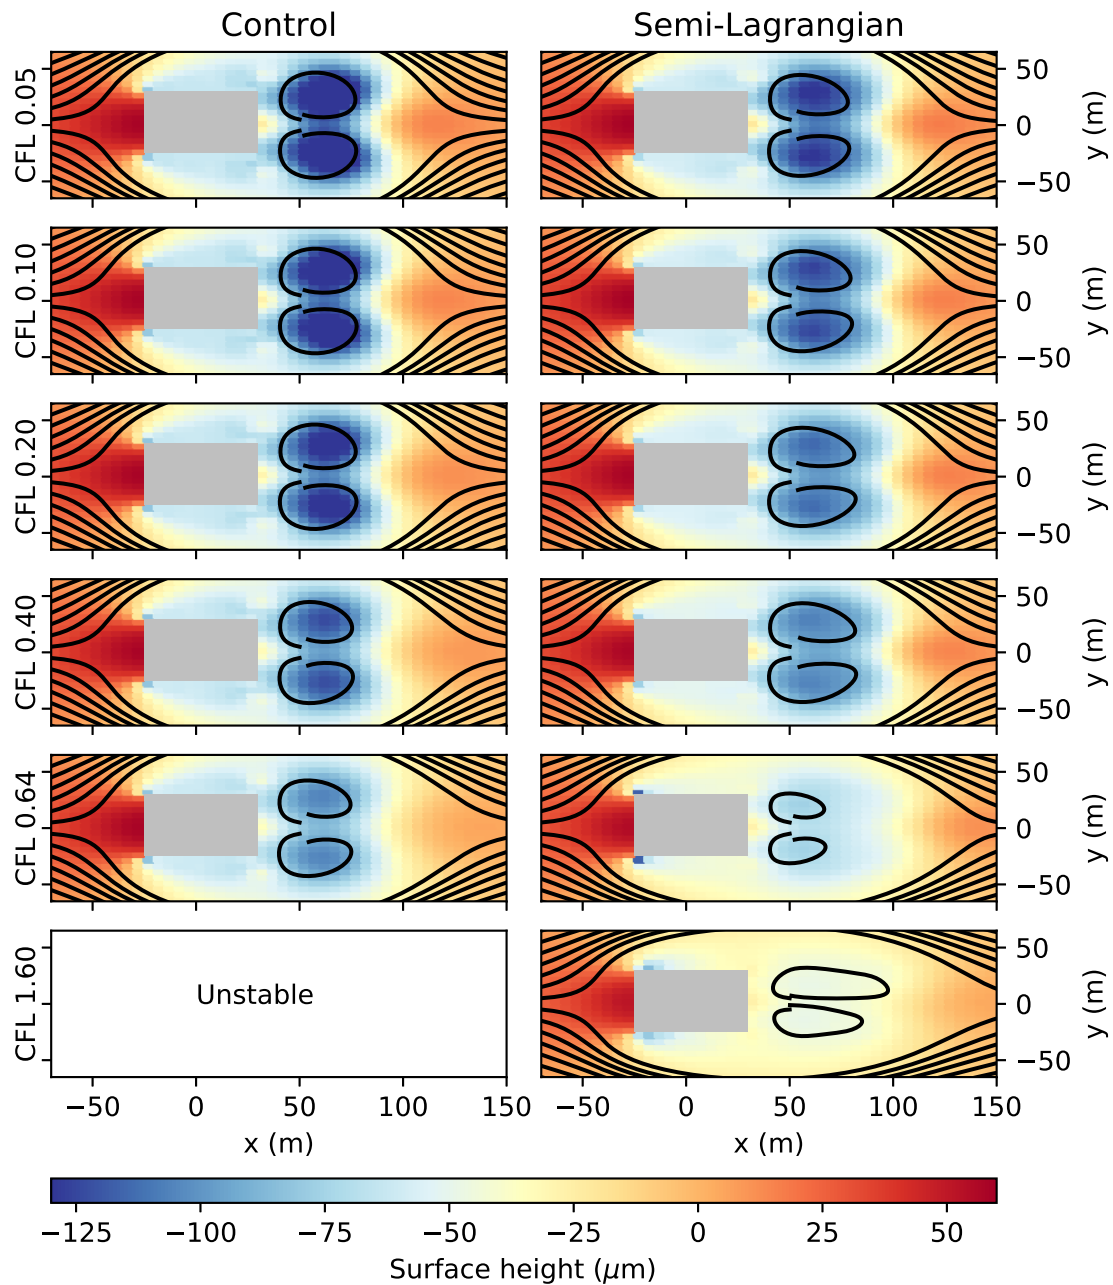


Figure 5. Free surface height and streamlines for the test case of section 5.1, after 8000s for $\Delta t = 5, 10, 20, 40, 64,$ and 160 seconds (top to bottom, with the approximate CFL-Courant number listed). Results for the Eulerian advection scheme are at left, and results for the semi-Lagrangian advection of momentum are at right. As the timestep increases both advection schemes show more diffuse behaviour, however the semi-Lagrangian advection scheme remains stable to $\Delta t = 160$ s whereas the Eulerian scheme becomes unstable after $\Delta t = 64$ s.

690 As with section 5.1, the test cases used NEMO-OPA’s linear free surface with a time-implicit solver, and tidal forcing was not present in these configurations. In a typical timestep, the vast majority of semi-Lagrangian trajectories converged in one iteration (mean 1.004 over the “semi-Lagrangian tracer” run). A very small minority of cells required an extended number of iterations or underrelaxation as described in section 4.2, but this did not affect the overall trajectory-calculation performance because convergence was measured (and iterations limited) on a per-cell basis.

695 Each run continued through late 2009. For reasons of space efficiency, we recorded the two-dimensional sea surface height, temperature, and salinity fields for each model day, and we preserved every fifth daily-mean, three-dimensional output of temperature, salinity, and horizontal ocean velocity.

For short and medium-term forecasts, the operational coupled forecasting systems at CMC are constrained by observations and periodic re-initialization. With a focus on this forecasting horizon the objective with these long free-runs was:

- 700
- To provide a test of model stability with semi-Lagrangian advection, in terms of both avoiding crashes and providing plausible ocean fields;
 - To check for any large-scale conservation errors, which might be difficult to correct given the sparsity of observation data for the deep ocean, and
 - To note any qualitative improvement or deterioration in the effective resolution of the model.

705 This first goal of model stability was met in part by the successful completion of these runs. Use of semi-Lagrangian advection for both tracers and momentum allowed us to increase the effective timestep from 10 minutes (with typical maximum ~~CFL number~~ Courant number¹⁰ of 0.2, found in the vertical direction) to 15 minutes (~~CFL~~ Courant number 0.3). Further increases led to instability and model crashes not from the advection component, but from the ice model. In this version of the model, the ocean/ice stress is coupled in a time-explicit way between the water and ice components. Concurrent work towards

710 a time-implicit coupling has given encouraging preliminary results on further timestep increases.

The use of semi-Lagrangian advection also gives global flows qualitatively similar to the control run, and average transports in the Atlantic overturning circulation and Circumpolar current are comparable between the control and semi-Lagrangian runs (figure 6). ~~However, the use of~~ The semi-Lagrangian runs appear to result in a slightly weaker overturning circulation and a slightly stronger circumpolar current than the control run, but these results may not be robust to re-tuned physical

715 parameterizations. Using semi-Lagrangian advection for the velocity components results in a significant decrease to overall ocean kinetic energy (figure 7), both during and after the spin-up period.

The cause of this energy disparity is under investigation, but we believe the most likely cause is the application of slope-limiting to the u and v fields independently. Future work will focus on taking a more nuanced approach to filtering, but this effect may not be very significant in a shorter-term forecast setting with frequent re-initializations from an analysis.

720 The second goal of global conservation was met. Although semi-Lagrangian advection does not guarantee conservation of tracers, the impact on the global balance of temperature and salinity was small. Figure 8 shows the evolution of ocean-

¹⁰Defining the Courant numbers in each direction as $\max(|u|)/e1u$, $\max(|v|)/e2v$, and $\max(|w|)/e3w$ respectively.

Parameter	Value	Comments
<i>Parameters common to all runs</i>		
atfp	0.1	Asselin time filter parameter
ln_zps	.TRUE.	Z-level vertical coordinate with partial (cut) cells
e3zps_min	25	Absolute minimum thickness of a cut cell
e3zps_rat	0.2	Relative minimum thickness of a cut cell
shlat	0	Free-slip lateral momentum boundary condition
nn_botfr	2	Nonlinear bottom friction
nn_bfro2	1e-3	Nonlinear bottom friction coefficient
nn_bfeb2	2.5e-3	Background turbulent kinetic energy coefficient
ngeo_flux	0	No bottom temperature geothermal heat flux
ln_dynhpg_imp	.TRUE.	Semi-implicit computation of the hydrostatic pressure gradient
ln_dynldf_bilap	.TRUE.	Bi-Laplacian hyperdiffusion of momentum
ln_dynldf_hor	.TRUE.	... acting in the horizontal direction
ahm0	-3e11	Momentum hyperviscosity coefficients
nsolv	2	Use the successive over-relaxation (SOR) free-surface solver
nsol_arp	0	... with an absolute-tolerance stopping condition
nn_sstr	0	No sea surface temperature damping
nn_sssr	0	No sea surface salinity damping
ndmp	0	No temperature or salinity damping in the water column
<i>Parameters for the control run</i>		
rdt	600	Model timestep
ln_traadv_tvd	.TRUE.	Total variation diminishing <u>Tracer variance dissipation</u> (TVD) tracer advection scheme
ln_traldf_lap	.TRUE.	Laplacian diffusion for the tracer
ln_traldf_iso	.TRUE.	... acting in the iso-neutral direction
aht0	300	Horizontal tracer diffusion coefficient
ln_dynadv_vec	.TRUE.	Vector form of the momentum advection operator
ln_dynvor_eeen	.TRUE.	... using the energy and enstrophy conserving scheme
resmax	1e-10	Absolute residual tolerance for the SOR free-surface solver
<i>Parameters for the semi-Lagrangian tracer run</i>		
rdt	600	Model timestep
ln_traldf_lap	.FALSE.	No explicit horizontal tracer diffusion
ln_dynadv_vec	.TRUE.	Vector form of the momentum advection operator
ln_dynvor_eeen	.TRUE.	... using the energy and enstrophy conserving scheme
resmax	1e-11	Absolute residual tolerance for the SOR free-surface solver
<i>Parameters for the semi-Lagrangian momentum and tracer run</i>		
rdt	900	Model timestep
ln_traldf_lap	.FALSE.	No explicit horizontal tracer diffusion
ln_dynadv_vec	.FALSE.	Flux form of the momentum advection operator
resmax	1e-11	Absolute residual tolerance for the SOR free-surface solver

Table 2. Selected dynamical and numerical namelist parameters for the test cases of section 5.2.




Figure 6. 61-day mean transports for the Atlantic overturning circulation (top; net northward flux above 1000m depth at 27.25° north latitude in the Atlantic Ocean) and Antarctic circumpolar current (bottom; net eastward flux at 67.75° west longitude in the Drake Passage) over time for the test cases of section 5.2

average temperature and salinity over time in these runs, and the effect of non-conservation attributable to the semi-Lagrangian advection of tracers is comparable to the ~~mangitude~~ magnitude of uncertainty in the global balance of atmospheric forcing – the imbalance seen in the control run.

725 Each case saw an overall temperature drift of about 0.04 K over the simulated period, with the semi-Lagrangian cases having a slight warming trend against the control run’s slight cooling trend, and all three runs showed a very small increase in ocean average salinity, by about 0.01 PSU.

The temperature change versus depth over the simulated period is shown in figure 9. Both the control and semi-Lagrangian runs showed a warming trend in the surface layers, but the semi-Lagrangian runs showed temperature stability in fluid layers below 1000 meters depth whereas the control run showed a cooling trend in these waters.

730

Despite the energy shortfall with semi-Lagrangian advection of momentum, we see tentative signs that the method increases the model’s effective resolution. Figure 10 shows one particular sea surface temperature realization, from the 31 December 2005

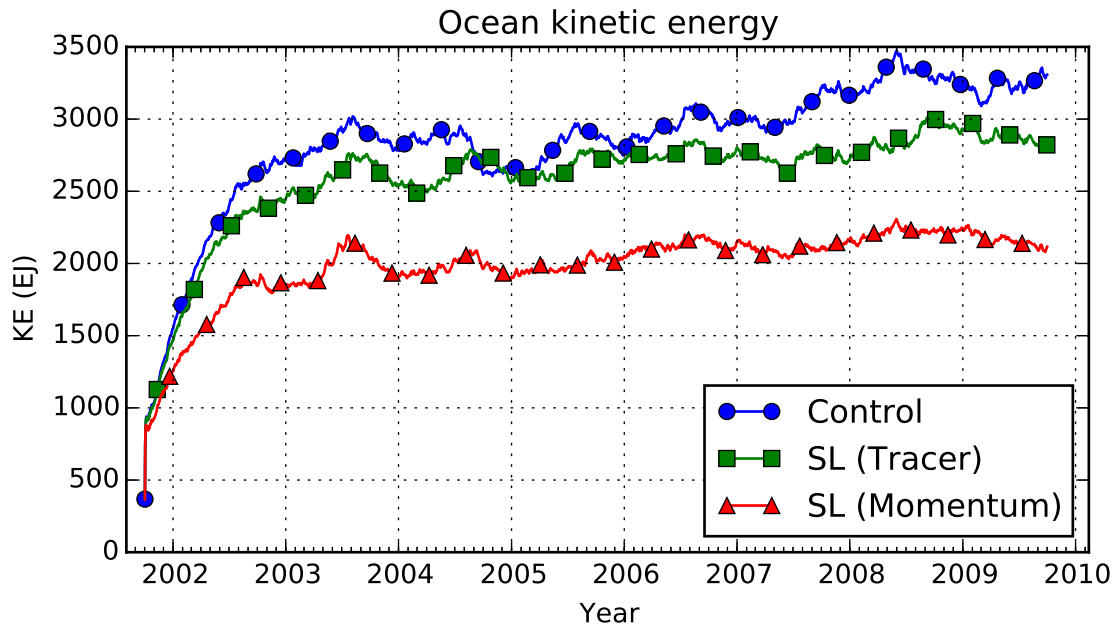


Figure 7. Total ocean horizontal kinetic energy (EJ) over time for the test cases of section 5.2. All of the test cases generally reproduce the monthly to yearly variability of kinetic energy, but the use of semi-Lagrangian momentum advection results in significantly lower total kinetic energy.

of each test case, along with the magnitude of the temperature gradient. The large-scale flows are similar between the control and semi-Lagrangian runs (and most similar between the control and semi-Lagrangian tracer run), but the semi-Lagrangian runs have noticeably stronger gradients in the sea surface temperature, in patterns that resemble smaller-scale eddies.

6 Conclusions and further work

This work has derived a semi-Lagrangian advection scheme for the NEMO-OPA model. After advecting a tracer or momentum field along estimated fluid parcel trajectories, it calculates a time-trend to provide to the remainder of the model; in this way the semi-Lagrangian scheme serves as a drop-in replacement for other tracer and (flux-form) momentum schemes.

[The development of this advection module relied on several new or newly-applied algorithms that might be relevant to other ocean models or other domains:](#)

- [The “semi-Lagrangian trend” form of equation \(7\) might be useful in other models when researchers wish to implement semi-Lagrangian advection after the fact, without disrupting the calculation of other forcing terms.](#)




Figure 8. Ocean average temperature (top, °C) and salinity (bottom, PSU) over time for the test cases of section 5.2. Although conservation is not guaranteed by semi-Lagrangian advection, long-term trends are similar between the semi-Lagrangian runs and the control run.

745

– The Hermite interpolation form in section 3, especially combined with the C^1 -continuous estimate of the vertical derivative in section 3.2 might find application in other domains where, as in the ocean, one dimension (the vertical) is more oscillatory than others.

– The exponential integration of trajectories in (4) may be useful in other applications that feature strong accelerations over trajectories. In particular, it forbids trajectory-crossing in one dimensional flows, and here that property ensures that trajectories remain inside the ocean domain.

750

– The “corner solution” treatment of velocity for trajectory calculations near corners might find use in other applications with staggered velocity components.

Overall, we find that the semi-Lagrangian method is effective at extending the realizable timestep in the NEMO-OPA model. In the simple domain of section 5.1, this method resulted in a stable simulation with advective CFL-Courant numbers in excess of 1. Although we only extended the timestep from 10 to 15 minutes for the semi-Lagrangian momentum run in section 5.2,




Figure 9. [Initial ocean-average temperature profile \(left, °C\) versus depth and change at the end of the simulated period \(right, 4 October 2009\) for the test cases of section 5.2. Both semi-Lagrangian runs show temperature stability in deeper waters, whereas the control run shows a small cooling.](#)

755 this limitation was imposed by the ice model. Disabling ice dynamics allowed us to increase the timestep to 30 minutes, but this would have made the results incomparable with those of the control and semi-Lagrangian tracer runs. Preliminary work with the CICE sea model and implicit coupling of the ice-ocean stress seems to allow us to relax the ice-related timestep restriction.

[Performance and implementation](#)

In spite of this increased timestep, the semi-Lagrangian method by itself does not yet improve overall computational performance. The semi-Lagrangian momentum and tracer run of section 5.2 took approximately one hour of computational time per five days of simulated time, using 128 Intel Xeon E5530 processors at 2.4GHz. With a 10-minute timestep, the semi-Lagrangian tracer run took approximately 50 minutes for the same five days of simulated time, whereas the control run took just 30 minutes. We expect these results to improve with further numerical optimization work. In particular, we did not take great care to ensure that loops were vectorized where possible, and it is much more difficult for compilers to automatically vectorize the point-by-point semi-Lagrangian computations compared to volume flux calculations in the traditional advection schemes.

[About one-third of the additional computational cost comes from trajectory iterations, and the remainder comes from the cubic interpolation. This suggests that the relative cost of semi-Lagrangian advection will be lower than presented here if](#)




Figure 10. Sea surface temperature (left) and the magnitude of its gradient (right) for the control (top), semi-Lagrangian tracer (middle), and semi-Lagrangian momentum and tracer (bottom) test cases of section 5.2, for 31 December 2005 in the Labrador Sea. Although the large-scale flows are similar, the runs with semi-Lagrangian advection of tracers have noticeably more fine-scale variability.

trajectories can be reused for multiple tracer species (such as biogeochemical constituents). Additionally, it suggests that a further optimization may be to re-use tracer trajectories for momentum advection, at least away from the boundaries where interpolating (staggering) trajectories might be reasonable. It seems unlikely that optimization will reduce the per-timestep penalty to the 20% value seen by Ritchie et al. (1995) for an atmospheric – model owing to the lack of three-dimensional implicit equations and expensive physical parameterizations elsewhere in NEMO-OPA – but we are hopeful that semi-Lagrangian advection will nonetheless improve overall system performance.

The parallel (MPI) implementation of this algorithm was straightforward. With the relatively modest increase in Courant number for the cases in this work, we simply needed to increase the inter-processor lateral halo (parameters `jpreci` and `jprecj`) to three points, which was sufficient to allow a fluid parcel arriving at a processor's edge to apply the full interpolating stencil for the Courant numbers reached in the presented simulations. This increase in halo size was small compared to the processor tile size of about 50×260 grid points for the runs in section 5.2. Extending this to support very large horizontal Courant numbers, however (if another solution could be found to stabilize baroclinic waves) would require either prohibitively large halo sizes or additional interprocessor communication to track fluid parcels that cross MPI boundaries.

Qualitative comments on results

Although the semi-Lagrangian method does not guarantee tracer conservation, we see no evidence that its implementation here leads to a degradation relevant in a weekly to seasonal forecast setting. ~~However, even small perturbations may~~ In particular, the deep-water temperature stability shown in figure 9 is an encouraging sign that semi-Lagrangian advection will preserve the deep-water structure that is weakly constrained by data. Even small imbalances, however, might become significant over ~~decadal to century-long climate simulations, so further~~ the decade-to-century timescales of climate simulations. Further work will be necessary to characterize this method before we can safely recommend ~~the use of~~ semi-Lagrangian advection in such settings.

For the test cases in section 5.2, semi-Lagrangian advection of tracers appears to slightly increase the effective resolution of the model. However, this effect is much more mixed when momentum is also advected with the semi-Lagrangian method, in part because the underlying currents differ. Both of these differences will be the subject of future study, with the specific intention of assessing these effects in the setting of shorter-term forecasts. We speculate that the overall loss of kinetic energy with semi-Lagrangian advection of momentum is attributable to the use of the slope limiter: limiting each component of velocity separately may be causing unrealistic diffusion of smaller-scale structures in the presence of background vorticity. We hope to address this issue with more selective limiting.

Future development

Finally, the development in this paper implicitly assumes that the coordinate system is static with time. This is not the case in NEMO-OPA when using its nonlinear free surface option, which necessarily implies time-varying vertical levels. Adapting the semi-Lagrangian method to this more general coordinate system will be a focus of future work, which will be required to apply this advection scheme to higher-resolution domains that require tide-permitting simulations.

805 Additionally, future versions of NEMO intend to move to a third-order Runge-Kutta time-stepping algorithm (Shu and Osher, 1997)
, which constructs a full timestep as a linear combination of forward Euler steps. We expect that the semi-Lagrangian “advective
trend” of (7) can be adapted to this framework in a straightforward manner by basing the calculated trend on the current-step
values of tracers and velocities, but the adaptation may require care to preserve the higher-order temporal accuracy of the
overall scheme.

Code availability. The modified NEMO (CeCILL license, version 2.0) code along with scripts and data used in this paper are available under Subich et al. (2020). The modified CICE 4.0 used in section 5.2 is not redistributable under a free license, but it has been made available for the topical editors and anonymous reviewers.

810 *Author contributions.* All authors were responsible for the concept. CS, PP, and FD were responsible for the necessary software development. CS contributed to the original draft preparation, and all authors contributed to review and editing.

Competing interests. The authors declare that they have no competing interests.

815 *Acknowledgements.* The authors would like to acknowledge Francois Roy of Environment Canada, who provided considerable technical support, especially with the forcing runs of section 5.2, and Jerome Chanut of Mercator Ocean, who provided helpful comments on a draft of this paper. The authors also are grateful for the peer review provided by Florian Lemarie and Mike Bell, which greatly improved this paper from its original submission.

The authors also owe a debt of inspiration to the semi-Lagrangian dynamical core of MC2 (Thomas et al., 1998).

References

- Arakawa, A. and Lamb, V. R.: A Potential Enstrophy and Energy Conserving Scheme for the Shallow Water Equations, *Monthly Weather Review*, 109, 18–36, [https://doi.org/10.1175/1520-0493\(1981\)109<0018:APEAEC>2.0.CO;2](https://doi.org/10.1175/1520-0493(1981)109<0018:APEAEC>2.0.CO;2), [https://journals.ametsoc.org/doi/abs/10.1175/1520-0493\(1981\)109%3C0018:APEAEC%3E2.0.CO%3B2, 1981](https://journals.ametsoc.org/doi/abs/10.1175/1520-0493(1981)109%3C0018:APEAEC%3E2.0.CO%3B2, 1981).
- Asselin, R.: Frequency Filter for Time Integrations, *Monthly Weather Review*, 100, 487–490, [https://doi.org/10.1175/1520-0493\(1972\)100<0487:FFFTI>2.3.CO;2](https://doi.org/10.1175/1520-0493(1972)100<0487:FFFTI>2.3.CO;2), <https://journals.ametsoc.org/doi/abs/10.1175/1520-0493%281972%29100%3C0487%3AFFFTI%3E2.3.CO%3B2>, publisher: American Meteorological Society, 1972.
- Bermejo, R. and Staniforth, A.: The Conversion of Semi-Lagrangian Advection Schemes to Quasi-Monotone Schemes, *Monthly Weather Review*, 120, 2622–2632, [https://doi.org/10.1175/1520-0493\(1992\)120<2622:TCOSLA>2.0.CO;2](https://doi.org/10.1175/1520-0493(1992)120<2622:TCOSLA>2.0.CO;2), [https://journals.ametsoc.org/doi/abs/10.1175/1520-0493\(1992\)120%3C2622%3ATCOSLA%3E2.0.CO%3B2, 1992](https://journals.ametsoc.org/doi/abs/10.1175/1520-0493(1992)120%3C2622%3ATCOSLA%3E2.0.CO%3B2, 1992).
- Casulli, V. and Cheng, R. T.: Semi-implicit finite difference methods for three-dimensional shallow water flow, *International Journal for Numerical Methods in Fluids*, 15, 629–648, <https://doi.org/10.1002/fld.1650150602>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/fld.1650150602, 1992>.
- Ducouso, N., Le Sommer, J., Molines, J. M., and Bell, M.: Impact of the “Symmetric Instability of the Computational Kind” at mesoscale- and submesoscale-permitting resolutions, *Ocean Modelling*, 120, 18–26, <https://doi.org/10.1016/j.ocemod.2017.10.006>, <http://www.sciencedirect.com/science/article/pii/S1463500317301518, 2017>.
- Ferry, N., Parent, L., Masina, S., Storto, A., Zuo, H., and Balmaseda, M.: PRODUCT USER MANUAL For Global Ocean Reanalysis Products GLOBAL-REANALYSIS-PHYS-001-009 GLOBAL-REANALYSIS-PHYS-001-011 GLOBAL-REANALYSIS-PHYS-001-017, Tech. Rep. 2.9, Copernicus Marine Service, <https://resources.marine.copernicus.eu/documents/PUM/CMEMS-GLO-PUM-001-009-011-017.pdf, 2016>.
- Girard, C., Plante, A., Desgagné, M., McTaggart-Cowan, R., Côté, J., Charron, M., Gravel, S., Lee, V., Patoine, A., Qaddouri, A., Roch, M., Spacek, L., Tanguay, M., Vaillancourt, P. A., and Zadra, A.: Staggered vertical discretization of the Canadian Environmental Multiscale (GEM) model using a coordinate of the log-hydrostatic-pressure type, *Monthly Weather Review*, 142, 1183–1196, 2014.
- Hildebrand, F.: *Introduction to Numerical Analysis: Second Edition*, McGraw-Hill Inc., 1974.
- Hodges, B. R. and Dallimore, C.: *Estuary, Lake and Coastal Ocean Model: ELCOM v2.2 science manual*, Tech. rep., Center for Water Research, University of Western Australia, 2006.
- Hollingsworth, A., Källberg, P., Renner, V., and Burridge, D.: An internal symmetric computational instability, *Quarterly Journal of the Royal Meteorological Society*, 109, 417–428, 1983.
- Hunke, E. C. and Dukowicz, J. K.: An elastic–viscous–plastic model for sea ice dynamics, *Journal of Physical Oceanography*, 27, 1849–1867, [https://doi.org/10.1175/1520-0485\(1997\)027<1849:AEVPMF>2.0.CO;2](https://doi.org/10.1175/1520-0485(1997)027<1849:AEVPMF>2.0.CO;2), <https://journals.ametsoc.org/doi/full/10.1175/1520-0485%281997%29027%3C1849%3AAEVPMF%3E2.0.CO%3B2, 1997>.
- Lauritzen, P. H.: A Stability Analysis of Finite-Volume Advection Schemes Permitting Long Time Steps, *Monthly Weather Review*, 135, 2658–2673, <https://doi.org/10.1175/MWR3425.1>, <https://journals.ametsoc.org/doi/10.1175/MWR3425.1>, publisher: American Meteorological Society, 2007.
- Leclair, M. and Madec, G.: z -Coordinate, an Arbitrary Lagrangian–Eulerian coordinate separating high and low frequency motions, *Ocean Modelling*, 37, 139–152, <https://doi.org/10.1016/j.ocemod.2011.02.001>, <http://www.sciencedirect.com/science/article/pii/S1463500311000126, 2011>.

- 855 Lemarié, F., Debreu, L., Madec, G., Demange, J., Molines, J. M., and Honnorat, M.: Stability constraints for oceanic numerical models: implications for the formulation of time and space discretizations, *Ocean Modelling*, 92, 124–148, <https://doi.org/10.1016/j.ocemod.2015.06.006>, <http://www.sciencedirect.com/science/article/pii/S1463500315001055>, 2015.
- Lemieux, J.-F., Beaudoin, C., Dupont, F., Roy, F., Smith, G. C., Shlyaeva, A., Buehner, M., Caya, A., Chen, J., Carrieres, T., Pogson, L., DeRepentigny, P., Plante, A., Pestieau, P., Pellerin, P., Ritchie, H., Garric, G., and Ferry, N.: The Regional Ice Prediction System (RIPS): verification of forecast sea ice concentration, *Quarterly Journal of the Royal Meteorological Society*, 142, 632–643, 860 <https://doi.org/10.1002/qj.2526>, <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.2526>, 2016.
- Leonard, B. P.: A stable and accurate convective modelling procedure based on quadratic upstream interpolation, *Computer Methods in Applied Mechanics and Engineering*, 19, 59–98, [https://doi.org/10.1016/0045-7825\(79\)90034-3](https://doi.org/10.1016/0045-7825(79)90034-3), <http://www.sciencedirect.com/science/article/pii/0045782579900343>, 1979.
- 865 Leonard, B. P.: The ULTIMATE conservative difference scheme applied to unsteady one-dimensional advection, *Computer Methods in Applied Mechanics and Engineering*, 88, 17–74, [https://doi.org/10.1016/0045-7825\(91\)90232-U](https://doi.org/10.1016/0045-7825(91)90232-U), <http://www.sciencedirect.com/science/article/pii/004578259190232U>, 1991.
- Locarnini, R. A., Mishonov, A. V., Antonov, J. I., Boyer, T. P., Garcia, H. E., Baranova, O. K., Zweng, M. M., Paver, C. R., Reagan, J. R., Johnson, D. R., Hamilton, M., and Seido, D.: *World Ocean Atlas 2013, Volume 1: Temperature*, 2013.
- 870 Lévy, M., Estublier, A., and Madec, G.: Choice of an advection scheme for biogeochemical models, *Geophysical Research Letters*, 28, 3725–3728, <https://doi.org/10.1029/2001GL012947>, <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2001GL012947>, 2001.
- Madec, G.: NEMO ocean engine, Note du Pôle de modélisation, Institut Pierre-Simon Laplace (IPSL), France, No 27, 2008.
- Madec, G. and Imbard, M.: A global ocean mesh to overcome the North Pole singularity, *Climate Dynamics*, 12, 381–388, <https://doi.org/10.1007/BF00211684>, <https://doi.org/10.1007/BF00211684>, 1996.
- 875 Mesinger, F. and Arakawa, A.: Numerical methods used in atmospheric models, Global Atmospheric Research Programme (GARP), 1976.
- Murray, R. J.: Explicit Generation of Orthogonal Grids for Ocean Models, *Journal of Computational Physics*, 126, 251–273, <https://doi.org/10.1006/jcph.1996.0136>, <http://www.sciencedirect.com/science/article/pii/S0021999196901369>, 1996.
- Ritchie, H., Temperton, C., Simmons, A., Hortal, M., Davies, T., Dent, D., and Hamrud, M.: Implementation of the Semi-Lagrangian Method in a High-Resolution Version of the ECMWF Forecast Model, *Monthly Weather Review*, 123, 489–880 514, [https://doi.org/10.1175/1520-0493\(1995\)123<0489:IOTSLM>2.0.CO;2](https://doi.org/10.1175/1520-0493(1995)123<0489:IOTSLM>2.0.CO;2), <https://journals.ametsoc.org/mwr/article/123/2/489/65285/Implementation-of-the-Semi-Lagrangian-Method-in-a>, publisher: American Meteorological Society, 1995.
- Robert, A.: A semi-Lagrangian and semi-implicit numerical integration scheme for the primitive meteorological equations, *Journal of the Meteorological Society of Japan. Ser. II*, 60, 319–325, 1982.
- Roy, F., Chevallier, M., Smith, G. C., Dupont, F., Garric, G., Lemieux, J.-F., Lu, Y., and Davidson, F.: Arctic sea ice and fresh- 885 water sensitivity to the treatment of the atmosphere-ice-ocean surface layer, *Journal of Geophysical Research: Oceans*, 120, 4392–4417, <https://doi.org/10.1002/2014JC010677>, <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014JC010677>, [_eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2014JC010677](https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2014JC010677), 2015.
- Shchepetkin, A. F. and McWilliams, J. C.: The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model, *Ocean Modelling*, 9, 347–404, <https://doi.org/10.1016/j.ocemod.2004.08.002>, <http://www.sciencedirect.com/science/article/pii/S1463500304000484>, 2005.
- 890

- Shu, C.-W. and Osher, S.: Efficient Implementation of Essentially Non-oscillatory Shock-Capturing Schemes, II, in: *Upwind and High-Resolution Schemes*, edited by Hussaini, M. Y., van Leer, B., and Van Rosendale, J., pp. 328–374, Springer Berlin Heidelberg, Berlin, Heidelberg, https://doi.org/10.1007/978-3-642-60543-7_14, https://doi.org/10.1007/978-3-642-60543-7_14, 1997.
- 895 Smith, G., Liu, Y., Colan, D. S., Reszka, M., Roy, F., Dupont, F., Lemieux, J.-F., Beaudoin, C., Lellouche, J.-M., Garric, G., Testut, C.-E., Pellerin, P., Ritchie, H., Lu, Y., and Davidson, F.: The CMC Global Ice Ocean Prediction System v2.1, Tech. rep., Canadian Centre for Meteorological and Environmental Prediction, collaboration.cmc.ec.gc.ca/cmc/cmoi/product_guide/docs/tech_notes/technote_giops-210_e.pdf, 2016.
- 900 Smith, G. C., Bélanger, J.-M., Roy, F., Pellerin, P., Ritchie, H., Onu, K., Roch, M., Zadra, A., Colan, D. S., Winter, B., Fontecilla, J.-S., and Deacu, D.: Impact of Coupling with an Ice–Ocean Model on Global Medium-Range NWP Forecast Skill, *Monthly Weather Review*, 146, 1157–1180, <https://doi.org/10.1175/MWR-D-17-0157.1>, 2018.
- Smolarkiewicz, P. K. and Pudykiewicz, J. A.: A Class of Semi-Lagrangian Approximations for Fluids, *Journal of the Atmospheric Sciences*, 49, 2082–2096, [https://doi.org/10.1175/1520-0469\(1992\)049<2082:ACOSLA>2.0.CO;2](https://doi.org/10.1175/1520-0469(1992)049<2082:ACOSLA>2.0.CO;2), <https://journals.ametsoc.org/jas/article/49/22/2082/23130/A-Class-of-Semi-Lagrangian-Approximations-for>, publisher: American Meteorological Society, 1992.
- 905 Subich, C., Pellerin, P., Smith, G., and Dupont, F.: Code repository for 'Semi-Lagrangian advection in the NEMO ocean model', a GMD discussion paper, <https://doi.org/10.5281/zenodo.3724713>, <https://zenodo.org/record/3724713>, 2020.
- Thomas, S. J., Girard, C., Benoit, R., Desgagné, M., and Pellerin, P.: A new adiabatic kernel for the MC2 model, *Atmosphere-Ocean*, 36, 241–270, <https://doi.org/10.1080/07055900.1998.9649613>, <https://doi.org/10.1080/07055900.1998.9649613>, 1998.
- Turkington, B., Eydeland, A., and Wang, S.: A Computational Method for Solitary Internal Waves in a Continuously Stratified Fluid, *Studies in Applied Mathematics*, 85, 93–127, <https://doi.org/10.1002/sapm199185293>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/sapm199185293>, [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/sapm199185293](https://onlinelibrary.wiley.com/doi/pdf/10.1002/sapm199185293), 1991.
- 910 Walters, R. A., Lane, E. M., and Henry, R. F.: Semi-Lagrangian methods for a finite element coastal ocean model, *Ocean Modelling*, 19, 112–124, <https://doi.org/10.1016/j.ocemod.2007.06.008>, <http://www.sciencedirect.com/science/article/pii/S1463500307000947>, 2007.
- Zweng, M., Reagan, J., Antonov, J., Locarnini, R., Mishonov, A., Boyer, T., Garcia, H., Baranova, O., Johnson, D., Seidov, D., and Biddle, M.: *World Ocean Atlas Volume 2: Salinity*, 2013.