Author response to comments of Reviewer 1:

We would like to thank the reviewer for their efforts in helping us to improve our manuscript. Referee remarks are shown in red and our responses are given in black font. Changes to the manuscript text are given in blue font.

General Comments:
This paper provides relevant scientific information on the differences in biogeochemical tracers between running a biogeochemical model in conjunction with an online OGCM versus running the same biogeochemical model offline using the Transport Matrix Method (TTM), when the transport matrices are based on exactly the same OGCM. Comparisons are made for key biogeochemical tracers (nitrate, phosphate, oxygen) and for phytoplankton biomass (diazothophs, other phytoplankton). The authors conclude that the differences are relatively minor compared to the differences between observations and the modeled values.

For the OGCM used, the University of Victoria Earth System Climate Model, an additional benefit of using the TTM approach is the feasibility of parallelism in the implementation. This resulted in two orders of magnitude gain in the wall-clock time needed to run the biogeochemical model.

The paper is well organized, and quite readable. The use of Taylor diagrams to contrast the results is helpful in summarizing the differences. There is sufficient detail about how this was implemented to facilitate reproducing the results.

It is unfortunate that basic changes to the advection scheme of the online OGCM and a compensating parameter change in a diffusion parameter were required in order to facilitate generation of the transport matrices. That means that the transport matrices are derived from a different version of the OGCM than the version with which the biogeochemical model is usually run. Comparisons of the results using the two OGCM versions are given; and comparisons of results using the modified OGCM and the TTM model are given. What is missing, as a practical matter, is a comparison of the biogeochemical results using the standard OGCM with the results using the TTM model.

We would not regard changing the scheme to suit our purpose “unfortunate”. (More of a “nuisance”.) There is no single, perfect advection scheme (or, more generally, numerical method). And while it is true that historically UVic ESCM has been run with a particular nonlinear advection scheme (which the reviewer calls “standard”), that is a choice dictated by a variety of factors and the model parameters had to be tuned appropriately. There is no reason to not use a different scheme if circumstances call for it. Practical choices like this are made all the time in scientific computing.

Regarding the reviewer’s suggestion to include a comparison of biogeochemical results from the “standard” OGCM with the TMM, we emphasize that the primary purpose of our manuscript is to provide a comparison between identical online and offline biogeochemical model runs in order to assess the offline method. This is why we did not spend much time re-tuning the physical model beyond increasing the vertical diffusion parameter to stabilise the overturning circulation. The reason we included a comparison between online runs of the OGCM performed with the two different advection schemes is that we recognize that other users of the UVic ESCM who may be interested in using the offline method, but currently use the default nonlinear advection scheme, will want to know about the consequences of switching advection schemes. But, again, it is not the main purpose of the paper. Moreover, we also note that one of the other referees thought that it wasn’t even necessary to

present this comparison, and that it be either removed or, at the very least, moved to an appendix. We agree that it does distract from the main goal of our paper and have followed that reviewer's advice and moved it to an appendix.

Specific Comments:

Page 1, Line 9, forward: The paper abstract needs to inform the reader that modifications to the OGCM were necessary in order to make it feasible to extract the TMs.

The following sentence has been added to Line 9:

The default, non-linear advection scheme was first replaced with a linear, third order upwind-biased advection scheme to satisfy the linearity requirement of the TMM.

Pg. 2, line 22, forward: This line implies that the ocean component of the ESCM was always run along with the atmosphere-biosphere-cryosphere-geosphere, since there is no statement to the contrary. Is this correct? If not, state which components were run in conjunction with the ocean model. This is particularly important in order to set the context for the great gain in the computer time that was made using the TTM.

In addition to the ocean GCM and biogeochemistry, the UVic ESCM has a prognostic sea ice model, atmospheric energy balance model (EBM) and land biosphere component. All of these were switched on for the online runs. Switching off the land model make very little difference to the computational time (although in our prescribed CO_2-experiments it was not necessary to turn it on, an oversight on our part). The EBM takes up about 20% of the computational time but, as currently implemented, it is not possible to switch it off and drive the OGCM with prescribed surface fluxes. (Of course, with appropriate code modifications it should be possible to do this.) To clarify, Sec. 3.2 has been modified to read:

The UVic ESCM is a serial code and thus unable to exploit more than one computational core. With the biogeochemical component switched on the model throughput on a typical Linux machine is about ~250 model years per day. (It should be noted that the model was run with the atmospheric energy balance model (EBM) switched on. This adds roughly 20% to the computational cost of running the model. However, as currently implemented, it is not possible to switch off the EBM in UVic ESCM and drive the ocean GCM with prescribed fluxes.) A 5000-year spin-up of the online biogeochemical model thus takes ~3 weeks. On the other hand, the PETSc-based TMM version can run in parallel, even though the underlying biogeochemical code is identical. While we have not carried out a detailed scaling analysis of the TMM version's performance, a similar 5000-year spin-up can be accomplished in 3.8 hours with 256 cores on NCAR's Yellowstone IBM iDataPlex cluster, and in 5.2 hours with 160 cores on GEOMAR/Kiel University's NEC HPC machine.

Were the forcings being used representative of the current era, without increased warming? Please state what forcing scenario was used.

As stated on Page 4, line 34, all simulations were carried out with a fixed, pre-industrial atmospheric CO_2 concentration of 277.4 ppm. The forcings are thus representative of that period.

Please give a little more in the description of the biogeochemical model; such as NZPD (declared later), and what phytoplankton groups, grazers, and nutrients are being tracked.

Page 2, line 24 has been amended and expanded to read:

The marine nutrients-phytoplankton-zooplankton-detritus (NPZD) biogeochemistry has increased in complexity since Eby et al. (2009), with the addition of iron limitation and revisions to zooplankton grazing (Keller et al., 2012), and subsequent minor updates. The NPZD model contains two phytoplankton types (a general type and diazotrophs) and a single zooplankton type, DIC, alkalinity, nitrate, phosphate (the base unit), and oxygen. Iron limitation is prescribed using a seasonally varying iron mask. Full model details can be found in Keller et al. (2012) and associated references.

Page 6, oxygen: It is mentioned that diazotrophs are disproportionately sensitive to low oxygen levels, as denitrification can be triggered. Small differences spatially in suboxic conditions can have significant impact. The extent of these differences between the OGCM and TTM in the modeled low oxygen, and where the low oxygen regions occur, needs to be shown.

Figure 9 (now Figure 4) has been changed to include suboxic regions (less than 5 mmol m3), and the depth of the map plots is lowered to 300 m.

Technical Corrections:
Pg. 1, Line 8: Insert “course-resolution” after “widely used”.

Done

Pg. 1, Line 9: Replace “for” with “from”.

Done

Pg. 1, Line 23: Replace “GCM” with “OGCM”.

Done

Pg. 2, Line 24: Delete “bug fixes and”.

Done

Pg. 4, Line 4: If possible, explain the effect of this (such as making the TMM model 5 times slower), so that the trade-off can be better understood.

We have modified the relevant paragraph in Sec. 2.1, which now reads as:

Lastly, UVic ESCM applies Fourier filtering in the zonal direction at high latitudes to remove grid-scale noise. The efficiency of the TMM arises from the fact that the discretized advection-diffusion operator has a limited stencil, i.e., only couples nearby points, giving rise to a sparse matrix. Fourier filtering on the other hand couples all points in the zonal direction, greatly reducing the sparsity of the transport matrix and hence the computational efficiency of the sparse matrix-vector products at the heart of the TMM. While the cost of a sparse matrix-vector product is implementation- and hardware-dependent and non-trivial to analyze (e.g., Gropp et al., 2000), it roughly scales with the number of non-zero elements per row. With a 3rd order upwind scheme, there are a maximum of 5 x 5 x 5 = 125 non-zero elements per row. With Fourier filtering that becomes nx x 5 x 5, where nx is the number of zonal grid points. In UVic ESCM, nx=100, implying that the TMM would be roughly

nx/5=20 times slower with Fourier filtering turned on. We therefore turn off polar filtering for the passive tracers used to extract the TMs. The numerical treatment of temperature and salinity by the model is not altered.

Pg. 5, Line 21: Modify “the impact of the advection scheme” to “the impact of the differences in the advection schemes” (or similar).

Changed to:

"highlight the impact of different advection schemes"

Pg. 6, Line 33: Insert before “offline and online models”, “ the biogeochemical component of the”.
Done

Pg. 7, Line 1: Insert after “monthly in these experiments”, “compared to [what time period] in the online simulation”.

We’re not sure what the reviewer is getting at here. There is no averaging necessary in the online model.

Pg. 7, Line 16: “it’s” should be “its”.

Done

Pg. 7, Line 24: End the sentence “with respect to observations”. Period. The hundredfold improvement in the time is a separate benefit of the particular implementation because of the added parallelization.

Done

Pg. 7, Line 28: Insert “biogeochemistry” at the beginning of the line, before “ in a few hours”.

Done

Pg. 7, Lines 19 forward: It might enhance the clarity of the conclusions to split this into two paragraphs, one on the biogeochemical results comparison, and another on the large improvement in the computer time required.

Done

Pg. 13, Fig. 2, forward: State that the top panels are zonal averages (assuming that they are).

Done