

# Numerically Relevant Timescales in the MG2 Microphysics Model

Sean Patrick Santos<sup>1</sup>, Christopher S. Bretherton<sup>1,2</sup>, and Peter M. Caldwell<sup>3</sup>

<sup>1</sup>Department of Applied Mathematics, University of Washington, Seattle, WA, USA

<sup>2</sup>Department of Atmospheric Sciences, University of Washington, Seattle, WA, USA

<sup>3</sup>Lawrence Livermore National Laboratory, Livermore, CA, USA

## Key Points:

- MG2 contains several processes that are poorly resolved at a time step of 5 minutes.
- Substepping MG2's rain-related processes significantly changes their rates.
- However, reducing MG2's time step within E3SM has little to no impact on the climate.

## Abstract

Climate models rely on parameterizations of a variety of processes in the atmospheric physics, but a common concern is that the temporal resolution is too coarse to consistently resolve the behavior that individual parameterizations are designed to capture. This study examines timescales numerically derived from the Morrison-Gettelman (MG2) microphysics as implemented within the Energy Exascale Earth System Model, version 1 (E3SMv1). Numerically-relevant timescales in MG2 are derived by computing the eigen-spectrum of its Jacobian. These timescales are found to often be smaller than the default 5 min timestep used for MG2. The fast timescales are then heuristically connected to individual microphysics processes. By substepping a few particular rain processes within MG2, the time discretization error for those processes was significantly reduced with minimal additional expense to the overall microphysics. While this improvement has a substantial effect on the target processes and on the vertical distribution of stratiform-derived rain within E3SMv1, the overall model climate is found to not be sensitive to MG2 time step.

## Plain Language Summary

The atmospheric components of climate models contain a number of physics parameterizations, subcomponents that are designed to capture particular aspects of the atmospheric physics. Cloud microphysics models are parameterizations designed to represent very small-scale cloud processes, including phase changes and the formation of precipitation. The accuracy of these parameterizations depends on the model time step: a shorter time step typically requires more computational resources, but also improves the model's accuracy. This paper examines a particular microphysics model, MG2, used with a time step of five minutes in the E3SMv1 climate model. By linearizing MG2, we can find characteristic time scales associated with this model, which are often much shorter than five minutes. This suggests that the usual time step is too large to fully capture the physics that MG2 represents. We also experiment with using a shorter time step for parts of the rain physics, and find that in many cases the rain mass is significantly affected. However, reducing the MG2 time step does not have much effect on the overall climate of E3SMv1.

## 1 Introduction

Atmospheric general circulation models (GCMs) consist of resolved-scale fluid “dynamics” and parameterized “physics” components. Dynamics is based on the Navier-Stokes equations, which makes its discretization and numerical solution straightforward (though computationally taxing). In particular, the dynamics timestep is typically controlled by the CFL condition (cite?) and related stability concerns. Parameterized physics, on the other hand, handles the grid-scale effects of the remaining collection of sub-grid scale atmospheric processes. Because the proper equations for these processes are often not known, and because the processes themselves often do not behave smoothly (e.g. across the clear-sky/condensate boundary), numerical treatment of atmospheric physics has received much less attention. In addition, use of limiters and implicit numerical schemes often prevents outright instability in physical parameterizations even at arbitrarily long time steps. As a result, physics timesteps are often set based on a tradeoff between throughput requirements and model sensitivity to further timestep reductions.

This situation leads to a lack of confidence in the numerical accuracy of parameterized physics schemes. As a result, a more formal numerical analysis of atmospheric physics routines would be of great interest. Stability analysis is one way to evaluate the time steps that should be used for various parameterizations. When solving a linear problem, we can readily determine whether any given method is absolutely stable (i.e. when the error does not grow exponentially in time), by checking the eigenvalues associated

with this problem. In the nonlinear case, there is no simple procedure that rigorously bounds the error growth in this way, but if the equations governing process rates are smooth, it can still be informative to analyze stability using the linearization of the problem, i.e. by analyzing the eigenvalues of the Jacobian used to linearize the problem (LeVeque, 2007). Using a time step size small enough to solve the linearized system will not guarantee any particular level of accuracy, but we can regard it as a necessary condition. If the numerical method being used would ordinarily be unstable at a given timestep, and model stability is only being maintained by limiters to prevent this instability, it is likely that the model's time step is too large to accurately approximate the solution of the underlying mathematical equations for the parameterized physics.

This study is focused on version 2 of the Morrison-Gottelman microphysics (MG2, (Gottelman & Morrison, 2015)), the component of the Energy Exascale Earth System Model version 1 (E3SMv1) responsible for microphysics in stratiform clouds. Previous work on the Community Atmosphere Model version 5 (CAM5), a predecessor to E3SMv1, shows that cloud cover and cloud ice distribution are significantly affected by the model timestep (Wan et al., 2014), and also that the stratiform microphysics considerably affects the temporal convergence rate of the overall model (Wan et al., 2015). While these results were based on an earlier version of the Morrison-Gottelman microphysics (MG1), it has also been noted that MG2 relies significantly on limiters to avoid negative cloud liquid mass, even when running at the default 5 minute timestep (Gottelman et al., 2015). While the limiters within MG2 specifically are designed to avoid introducing errors in the conservation of water mass, it is known that other limiters implemented to avoid negative mass in E3SMv1 can produce significant water mass conservation violations in the atmosphere at coarse resolution (Zhang et al., 2018). All this strongly suggests a need for a more detailed investigation of the timestep sensitivity of E3SMv1's cloud processes, since it appears that many of these processes are not adequately resolved, and therefore that the model physics is excessively reliant on conservation limiters and resolution-dependent tuning to produce a reasonable climate.

Section 2 will describe the relevant features of MG2 and its usage within E3SMv1. Then we will address questions about the numerically-relevant timescales in MG2 in several sections. First, we examine the eigenvalues of the Jacobian in Section 3, which are derived numerically based on a broad sample of conditions from the global model. Then, we associate these eigenvalues with specific processes in Section 4. This allows us to connect the timescales associated with MG2 to subsets of the specific physical processes it implements. Next, in Section 5 we examine the time scales of MG2 in different regimes where specific processes are dominant (e.g. warm versus cool grid cells, cloudy grid cells versus precipitation in an otherwise clear sky). Finally, in Section 6, we check these results by substepping the processes identified as having fast timescales. Each of these four sections is divided between one or more methodology sections, followed by a section presenting the results. We will conclude by discussing the relevance of these results to the numerics of MG2 generally, and introduce some preliminary information about the impact of MG2's timestep on the global model.

## 2 Model Description

E3SM is an ongoing U.S. Department of Energy (DOE) project designed to produce a state-of-the-art earth system model that can leverage the DOE's largest supercomputers to produce high-resolution simulations. The scientific goals of this project relate to three main topics: (1) the water cycle, (2) the cryosphere, and (3) biogeochemistry (Golaz et al., 2019). While most of this paper is concerned with its stratiform microphysics (MG2) running in isolation using a stand-alone driver, we do use the global model to generate input data and do some preliminary investigation of the impact of MG2's timestep on global climate. For this purpose, we use version one of the E3SM Atmosphere

Model (EAMv1), which is coupled to a data ocean and run at an approximately 100km (1) horizontal resolution.(Rasch et al., 2019)(Xie et al., 2018)

Within EAMv1 at this resolution, much of the physics, including the deep convection scheme, runs at a 30 minute time step, which is also the frequency of physics-dynamics coupling. However, both MG2 and the Cloud Layers Unified By Binormals parameterization (known as CLUBB, which handles cloud macrophysics, shallow convection and turbulence) use a shorter time step. These parameterizations, as well as a handful of aerosol- and ice-related microphysical processes calculated outside of MG2, are substepped in a loop within the physics driver, and use a 5 minute (300s) timestep. Although the time step for much of the physics is reduced to 15 minutes when running EAMv1 at high resolution, the MG2 and CLUBB time step is not changed. We therefore regard this as the default MG2 time step, and will concern ourselves with the behavior of MG2 only at 5 minute and smaller time step sizes.

The MG2 microphysics uses a two-moment representation of four hydrometeor types. For cloud liquid, the two prognostic variables are the grid-cell-average mass mixing ratio ( $q_c$ ) and number concentration ( $n_c$ ), and process rates are calculated assuming that the particle diameter follows a gamma distribution, with the mean size dictated by the ratio  $q_c/n_c$ , and the other shape parameter diagnosed from the number concentration. Similarly, an average mass mixing ratio and number concentration are used to describe the model’s cloud ice ( $q_i$ ,  $n_i$ ), rain ( $q_r$ ,  $n_r$ ), and snow ( $q_s$ ,  $n_s$ ), though these other hydrometeors are assumed to have sizes that follow a simpler exponential distribution. The temperature ( $T$ ), the humidity ( $q$ ), and these hydrometeor variables together make up the ten prognostic state variables for MG2, and for the purposes of this paper we define the “state” used by MG2 to consist of only these variables. Other inputs are also used by MG2, such as the cloud fraction, but these are diagnosed by other parameterizations, and we assume that they are roughly constant over the course of the 300s time step at which MG2 runs.

Parallel splitting is used for most of the processes in MG2, i.e. the inputs to these processes are generally the same as the inputs to MG2 itself. Once the process rates are calculated, the state is updated by adding contributions from all processes and taking a single forward Euler method step. As part of this process, a series of conservation checks are performed. For each hydrometeor, MG2 checks whether the forward Euler step would produce a negative mass mixing ratio or number concentration. If so, all process rates are scaled down to avoid producing negative values. All processes that are applied at this stage are listed in Table 1. Note that the droplet activation is a special case; its rate is calculated outside of MG2, and it is applied sequentially using a single forward Euler step before all the other processes in this table. There are also three other processes controlled by “external” schemes. While these are applied as if they had been calculated within MG2 itself, within E3SMv1 they are prescribed by a different scheme.

After this, MG2’s sedimentation is run on the updated state. The sedimentation calculates the mean fall speed of each hydrometeor in each grid cell, and uses a first-order explicit upwind scheme to update the profile of the entire column. Hydrometeors that reach the surface are considered precipitation and reported as a surface flux to the host model.

Some limiters and instantaneous adjustments (e.g. forcing all precipitation to freeze/melt when introduced to very cold/warm grid cells) are also applied at two stages: at the beginning of MG2, and after the sedimentation.

### 3 Timescales in MG2

In this section we investigate the inherent timescales present in the MG2 code by calculating the eigenvalues of a numerically derived Jacobian.

Short name	Variables affected	Description
Rain Evap.	$T, q, q_r, n_r$	Evaporation of rain droplets
Snow Subl.	$T, q, q_s$	Sublimation of snow
Vapor/Ice Transfer	$T, q, q_i, n_i$	Vapor deposition onto cloud ice minus ice sublimation
Berg. (Snow)	$T, q_c, q_s$	Bergeron process on snow
Liq. Accr. Snow	$T, q_c, n_c, q_s$	Collection of cloud water by snow
Sec. Ice Prod.	$T, q_c, q_i, n_i$	Secondary ice production via the Hallett-Mossop process
Het. Rain Frz.	$T, q_i, n_i, q_r, n_r, q_s, n_s$	Heterogeneous rain freezing
Rain Accr. Snow	$T, q_r, n_r, q_s$	Collection of rain by snow
Berg. (Cloud)	$T, q_c, q_i$	Bergeron process on cloud ice
Autoconversion	$q_c, n_c, q_r, n_r$	Autoconversion of cloud droplets to rain
Accretion	$q_c, n_c, q_r$	Accretion of cloud water by rain
Ice Auto.	$q_i, n_i, q_s, n_s$	Autoconversion of cloud ice to snow
Ice Accretion	$q_i, n_i, q_s$	Accretion of cloud ice by snow
Rain Self-col.	$n_r$	Self-collection of rain
Snow Self-col.	$n_s$	Self-aggregation of snow
Drop. Activ.	$n_c$	Droplet activation from external aerosol scheme
Nucleation Dep.	$T, q, q_i, n_i$	External classical nucleation scheme.
Immersion Frz.	$T, q_c, n_c, q_i, n_i$	External heterogeneous freezing scheme.
Contact Frz.	$T, q_c, n_c, q_i, n_i$	External heterogeneous freezing scheme.
Size Limiters	$n_c, n_i, n_r, n_s$	Limiters constraining hydrometeor particle sizes to remain in relevant ranges

**Table 1.** MG2 process descriptions and short names used in figures.

### 3.1 Methodology - Running MG2 as a Standalone Process

Within E3SM, MG2 and CLUBB are substepped together using a 300 second time step, which is both the default and maximum recommended time step for CLUBB. We can view the main role of MG2 as providing the average rates of change of MG2's ten state variables due to microphysical processes over a given timestep. We can evaluate how well MG2 answers this question without use of the full E3SM model, as long as we have a representative sample of atmospheric columns and a method for running MG2 outside of E3SM.

To obtain a sample of input columns, we ran a standard pre-industrial global E3SM simulation with  $\sim 100$  km atmospheric resolution (ne30\_ne30 grid) and prescribed sea surface temperature (compset F1850C5AV1C-04P2). The code used for this integration was a beta version (git hash 7a17edbe) which is structurally similar to the E3SMv1 release. Code modifications between this version and the release are not expected to make a difference for the results here. At the end of a five day run, we saved a “snapshot” of every column on the planet at the final timestep, including every input used in the final call to MG2.

This provided us with a representative sampling of columns for a particular day in January, which is the basis for the analysis in the remainder of this paper. A larger or more varied sample of columns, taken over the course of a run lasting a year or more, would have some statistical differences, especially since such a sample would capture seasonal variation. However, since this sample causes MG2's most important processes to operate under a wide variety of conditions, we do not believe that broader sampling would affect our overall conclusions.

In order to run MG2 in isolation from the rest of the model, we used F2PY to create a Python interface and compiled MG2 and the interface into a stand-alone library. This library was used by a set of drivers to produce the MG2 stand-alone results in this paper.

We narrowed down the number of active processes to a subset of those in MG2 by implementing switches to disable two types of process. First, sedimentation was disabled in our runs. MG2's sedimentation runs sequentially after the state is updated by all other processes, and it uses an adaptive timestep, smaller than that of MG2 as a whole, to satisfy its CFL condition. The numerical problems posed by the sedimentation are of a different character from those posed by the rest of MG2, and so we believe that it would be best to examine the sedimentation in a separate study.

Since sedimentation is the (only) process in the MG2 microphysics that transports mass between levels, disabling sedimentation allowed us to view MG2 as a tool for solving an ODE in a collection of uncoupled grid cells, rather than as a tool for solving a 1D PDE on an atmospheric column. This provides two benefits:

1. All remaining processes have a well-defined and easily controlled time step, since only the sedimentation uses an adaptive time step.
2. If sedimentation was enabled with 72 vertical layers per column, then the state vector input to MG2 would be 720-dimensional (given the 10-dimensional state for each grid cell). With sedimentation disabled, MG2 operated on individual grid cells with a 10-dimensional state in each one.

Second, many “instantaneous” processes were disabled. Instantaneous processes are those processes which MG2 implements not by calculating an expected process rate, but by making an immediate adjustment from an state that is considered unstable or a violation of model assumptions, to a more stable or valid state. Usually such processes

involve rapid changes of state, e.g. quickly melting snow that has fallen into a very warm grid cell.

We disabled these processes because we were interested in how the microphysical processes in MG2 are resolved, but instantaneous processes are by definition so fast that the model is not designed to resolve them in the first place. Furthermore, disabling the instantaneous processes removed any explicit dependence of the calculation on the time step size, which allowed the time step chosen during the Jacobian calculation to be chosen arbitrarily without significantly impacting the results.

One final change that was made to MG2 was to extract the calculation of precipitation area fraction from the main code, so that it was only calculated once per driver call. In conjunction with disabling sedimentation, this change was needed to allow MG2 to be run on individual grid cells rather than full columns, which greatly reduces computational cost. Furthermore, MG2's default precipitation area fraction calculation depends mostly on the cloud fraction calculated by CLUBB and is insensitive to changes in the state vector. As a result, forcing this variable to remain constant for short MG2 runs does not produce significantly different results from allowing it to vary over time.

### 3.2 Methodology - Measuring Differences in MG2 States

We often wanted to compare MG2 states in order to (a) establish whether MG2 is active enough to produce a final state that is significantly different from its input state, or (b) measure the magnitude of a change in MG2's output given a change in MG2's parameters. However, MG2's ten state variables include temperature, mass mixing ratios, and number concentrations, which have different units and orders of magnitude, so an isotropic norm such as the  $L^2$  norm is ill-suited to represent the distance between MG2 states. One way to remedy this would be to use a set of constant weights to convert all variables to a common set of units; for instance, a difference in temperature could be converted to an approximate change in mass mixing ratio of liquid water that would be necessary to produce that change in temperature via evaporation/condensation.

A simpler approach is to focus solely on the mass mixing ratio terms in the state vector, from which we define a quantity called the total water mass difference ( $D_w$ ). This is simply half of the  $L^1$ -norm of the water mass, or for two states labeled as  $\mathbf{s}_1$  and  $\mathbf{s}_2$ :

$$D_w(\mathbf{s}_1, \mathbf{s}_2) = \frac{|q_1 - q_2| + |q_{c1} - q_{c2}| + |q_{i1} - q_{i2}| + |q_{r1} - q_{r2}| + |q_{s1} - q_{s2}|}{2} \quad (1)$$

This value can be interpreted as the total amount of water that is in a different category between two states. For instance, if the difference between the two states is equivalent to the evaporation of  $10^{-3}$  g/kg of water, then  $D_w = 10^{-3}$  g/kg.

The total water mass difference is arguably the most straightforward method for measuring the differences between MG2 outputs. Technically it does not account directly for differences in temperature or number concentration. However, MG2 can only create changes in temperature via phase change, which will affect the mass mixing ratios, and in practice changes in number concentration will only occur in circumstances where processes that affect mass mixing ratio are also present. Differences in the overall MG2 state can therefore be distinguished by looking at mass mixing ratio alone.

### 3.3 Methodology - Jacobian Eigenvalue Calculation

A standard method for analyzing the stability of a numerical method on a nonlinear problem is to linearize the system about a given state by calculating the Jacobian.



The stability of the method on the linearized problem depends on the location of the eigenvalues of the Jacobian (multiplied by the time step) in the complex plane.

Alternatively, one can think of the inverse of these eigenvalues as a set of time scales associated with the linearized system. In order to keep the system stable without depending on limiters or other artificial corrections to a given method, our time step must be less than or equal to the smallest time scales derived from the Jacobian.

However, calculating this Jacobian presents some difficulty. The system of equations represented by MG2 is quite complex, and contains a large number of thresholds that adversely affect the smoothness of the process rates (e.g. the process rates are generally continuous, but not  $C^1$ , at the freezing point of water).

Rather than attempt to analytically differentiate MG2 around each point, we numerically calculated the Jacobian using the *numdifftools* Python package. The specific method we used from this package is a second-order forward difference method, supplemented with one stage of Richardson extrapolation (so the resulting method is 3rd order accurate).

The use of a forward (rather than central) difference method was important. MG2 produces floating point exceptions when given negative concentrations, so we needed to use a one-sided method in order to numerically linearize about a state where any constituent has zero mass. Furthermore, by using a small linear change of variables, we could also ensure that this one-sided method does not use perturbed states that violate MG2's size limiters. Since MG2 enforces these limiters at several points throughout the code, any perturbation that leads to a state outside of the acceptable size ranges will result in an instantaneous adjustment back to a valid state, and these instantaneous adjustments are not of interest.


Since the Jacobian is a  $10 \times 10$  matrix, finding its eigenvalues and eigenvectors was a negligible cost compared to the MG2 calls necessary to calculate the Jacobian in the first place. Thus we simply used SciPy's *linalg.eig* (which wraps LAPACK calls that use the QR algorithm) for this calculation. The MG2 Jacobian almost always has a full set of eigenvectors when MG2 is active, though the Jacobian becomes defective when no hydrometeors are present. This appears to be because the process rates are completely insensitive to the temperature and humidity variables (rates are always zero in the absence of hydrometeors), but perturbations to the hydrometeor masses will result in temperature and humidity changes. This introduces asymmetric off-diagonal elements to the Jacobian that make it non-diagonalizable. We had no interest, however, in states where MG2 is inactive or nearly-inactive. We therefore only included cases where the process rates are sufficiently large, as defined by a total water mass difference threshold. Specifically, for an initial state  $\mathbf{s}_1$  and a final state  $\mathbf{s}_2$  we were only interested in columns where  $D_w(\mathbf{s}_1, \mathbf{s}_2)/\Delta t \geq 10^{-7} \text{ g kg}^{-1} \text{ s}^{-1}$ . Since most grid cells in the model have little to no condensate, this eliminated most of the total grid cells from consideration (20.3%), which also significantly sped up calculations on the full sample.

The eigenvalues of the Jacobian are furthermore almost always real or have small imaginary parts. We therefore solely focused on the real parts of these eigenvalues.

### 3.4 MG2 Timescale Results

Figures 2 through 3 contain histograms of MG2's eigenvalues for grid cells above the  $10^{-7} \text{ g kg}^{-1} \text{ s}^{-1}$  cutoff. These eigenvalues are further categorized based on whether they are associated with active or inactive processes, where an active process is one that either affects water mass above a rate of  $10^{-7} \text{ g kg}^{-1} \text{ s}^{-1}$ , or affects a hydrometeor number above  $2.98 \text{ kg}^{-1} \text{ s}^{-1}$ , which is equivalent to the rate of  $400 \mu\text{m}$  rain particles that would






**Figure 1.** Graphs of complex eigenvalues, using a log scale. Eigenvalues with magnitude less than  $10^{-5}$ s are placed at the origin. Left panel shows distribution of eigenvalues in the complex plane, with a circle at  $(300\text{s})^{-1}$  for comparison, as well as the region of absolute stability for the Euler method. Note that the number of eigenvalues on the real axis is far larger than what can be represented on this color scale, implying that real eigenvalues dominate. Left panel compares real zeros (taken to be those with  $\leq 10^{-5}\text{s}^{-1}$  imaginary part) with all zeros, showing that almost all eigenvalues are real.

have to be introduced to produce a mass change of  $10^{-7}\text{gkg}^{-1}\text{s}^{-1}$ . We can roughly divide these eigenvalues into five categories:

1. Eigenvalues associated with inactive processes (shown in blue). These eigenvalues are typically related to physics that is not active in a given regime, e.g. ice physics in a warm grid cell, so they are not as relevant to the numerics in practice. We will discuss how we associate eigenvalues with particular processes in a later section.
2. Negative eigenvalues of magnitude greater than  $(300\text{s})^{-1}$ . These eigenvalues correspond to short-timescale processes, which have rates that may decay too rapidly for the default MG2 timescale to handle. In general, MG2 requires limiters to avoid instability caused by these processes.
3. Negative eigenvalues of magnitude less than  $(300\text{s})^{-1}$ . These eigenvalues correspond to processes which MG2 is able to resolve, assuming roughly linear behavior of MG2.
4. Near-zero eigenvalues. These eigenvalues correspond either to extremely slow feedbacks within MG2, or to forbidden directions of motion in the phase space (particularly eigenvectors that are perpendicular to surfaces of constant energy or mass).
5. Positive eigenvalues. These eigenvalues correspond to eigenvalues of processes that are temporarily in a state of positive feedback (e.g. accretion produces larger raindrops, and those larger raindrops will be effective at accreting even more cloud water). Generally speaking, MG2 avoids instability from these processes due to its nonlinearity, since all MG2 processes “use up” some form of mass or number, and therefore will eventually slow down over time.

In general we were most interested in dealing with the second case, the large negative eigenvalues, since these are common and correspond to cases where MG2’s large time scale is likely to cause problems. Figure 2 shows that there are a great number of




**Figure 2.** Histogram of the real parts of the eigenvalues of MG2’s Jacobian, focusing on negative eigenvalues. Red bars represent eigenvalues associated primarily with an active process. Green bars represent eigenvalues associated with at least one active process, but no “primary” process. Blue bars represent eigenvalues associated with inactive processes. Black line placed at  $(300\text{s})^{-1}$  for comparison with the MG2 timestep.

timescales that are not adequately resolved at MG’s current time step (by this measure, even more than the number that do seem adequately resolved). We also wanted to understand the nature of the positive eigenvalues.

In order to better interpret this result, however, we needed to associate the eigenvalues more closely with both specific sets of processes within MG2, and the physical conditions under which these eigenvalues arise. For instance, there may be cases where the model is formally unstable without limiters at a given timestep, but in practice there is little difference between the resolved and limited behaviors.

#### 4 Connecting Timescales to Specific Processes

Documenting the timescales of MG2 processes was of inherent interest, but we were particularly interested in the finding that MG2 is often integrated using a  $\Delta t$  which is too large to represent the processes it represents. Substepping all of MG2 in order to capture these timescales is not computationally feasible, but numerically accurate solutions



**Figure 3.** Histogram of the real parts of the eigenvalues of MG2’s Jacobian, focusing on positive eigenvalues. Red bars represent eigenvalues associated primarily with an active process. Green bars represent eigenvalues associated with at least one active process, but no “primary” process. Blue bars represent eigenvalues associated with inactive processes.

may still be affordable if the need for substepping could be isolated to just a few processes.

#### 4.1 Methodology - Measuring Eigenvalue-to-Process Associations

We sought to assign each eigenvalue from Sect. 3 to one or more of the MG2 processes listed in Table 1. Let us label the grid-cell output tendencies from applying MG2 to state  $\mathbf{s}$  as  $\mathbf{r} = \partial \mathbf{s} / \partial t$ . The Jacobian  $J_{\mathbf{r}}$  has  $i, j$ th entry  $\partial r_i / \partial s_j$  where  $r_i$  is the  $i$ th entry in  $\mathbf{r}$  and  $s_j$  is the  $j$ th entry in  $\mathbf{s}$ . As noted above,  $J_{\mathbf{r}}$  is diagonalizable for the states that were of interest for us, so matrices of eigenvalues  $\Lambda$  and eigenvectors  $V$  exist such that  $V \Lambda V^{-1} = J_{\mathbf{r}}$ .

If we label the tendencies due to a particular process  $p$  as  $\mathbf{r}_p$ , then  $\sum_{p=1}^P \mathbf{r}_p = \mathbf{r}$ . This leads directly to the identity

$$\Lambda = V^{-1}(J_{\mathbf{r}})V = \sum_{p=1}^P V^{-1}(J_{\mathbf{r}_p})V \quad (2)$$

which is the heart of our association method. Now construct a matrix  $\tilde{C}$  of dimensions  $N$  (number of eigenvalues) by  $P$  (number of processes to consider) whose  $n, p$ th entry

$\tilde{C}_{np}$  is the  $n$ -th diagonal element of  $V^{-1}(J_{\mathbf{r}_p})V$ . Since it depends only on these diagonal elements,  $\tilde{C}$  is independent of the scaling of the columns of  $V$ . Additionally,

$$\sum_{p=1}^P \tilde{C}_{np} = \lambda_n \quad (3)$$

so the  $n$ th row of  $\tilde{C}$  partitions  $\lambda_n$  into contributions from each process.

As a final step, normalize  $\tilde{C}$  by taking the absolute value of each element, and by dividing each row by its 1-norm, producing a new matrix  $C$ . Then each element of  $C$  describes the fractional contribution of its column's process to its row's eigenvalue. If any element of  $C$  is greater than 0.5 (our heuristic for whether at least 50% of an eigenvalue's magnitude can be attributed to a particular process), then we say that the eigenvalue for that row is *primarily associated* with the process for that column. This is how the colors in Figures 2-3 are derived: the red bars count eigenvalues primarily associated with active processes, blue bars correspond to inactive processes, and green bars represent eigenvalues without a primary association. We are mainly concerned with large-magnitude eigenvalues, which almost all have a primary association with a particular process.

For purposes of this study, we treated the effects of MG2's size limiters as if they were a separate physical process. This was largely because MG2's state can become unrealistic or invalid if these limiters are disabled, so we were not able to disable them as we had with other instantaneous processes, and it was therefore necessary to account for their effects on MG2's state. Note also that these limiters are not applied purely for artificial, numerical reasons, since the maximum size limiters are also the means by which MG2 accounts for the spontaneous breakup of large particles, which is important especially for a reasonable treatment of precipitation.

## 4.2 Methodology - Correlation of Multiple Processes


In addition to identifying processes associated with particular timescales, we wanted to identify tightly coupled processes. Such processes must be solved simultaneously or substepped together in order to obtain numerically accurate solutions. In addition, accounting for such coupling is needed in order to create useful conceptual models of microphysical behavior.

To identify tightly coupled processes, we first tallied the number of eigenvalues with primary associations to each process. We then looked at the average value of  $C$  across all modes primarily associated with a given process. If, for instance, modes that are primarily associated with autoconversion also typically have strong association with accretion, then we would focus on autoconversion/accretion coupling for further study and optimization.

## 4.3 Process Association Results

We will first examine the association between positive eigenvalues and specific processes, shown in Figure 4.

First, we found positive eigenvalues that are associated mainly with accretion-related processes. These eigenvalues are positive when the accumulating particles are small and few in number, since in this case the initial accretion increases the size of the particles, making them better at accumulating additional mass. In the long run, this is not a threat to the stability of the model, since eventually the accretion will begin to deplete the cloud, which will cause the rate of accretion to slow (see Figure 5). We will not examine this issue further here, but it should be noted that long timesteps may delay the onset of heavy precipitation in these cases.




**Figure 4.** PDF of positive MG2 Jacobian eigenvalues, based on process primarily associated with each eigenvalue. Each row sums to 1, except for processes with no associated positive eigenvalues. The right-hand y-axis label shows the number of positive eigenvalues associated with each process.

Second, we found eigenvalues that are associated mainly with evaporation or sublimation. These eigenvalues are again positive mainly when the affected particles are small and few in number, as such particles rapidly evaporate when out-of-cloud. In such cases, the particle mass rapidly drops to zero, and the temporal resolution should not be relevant to the final state of the system.

Turning back to the negative eigenvalues, we wanted to examine the processes associated with the smallest timescales. As seen in Figure 6, these processes are:

1. Accretion of cloud water by rain.
2. Rain evaporation.
3. Rain self-collection.
4. Snow sublimation.
5. Vapor/Ice transfer.

Some other processes were active with relatively short time scales (particularly heterogeneous rain freezing and the Bergeron process), but these were less active. The par-



**Figure 5.** Diagram of accretion regimes. In case (a), the cloud mass is still much larger than the rain mass. As the rain accretes more liquid mass, it becomes more effective at accreting additional mass, so accretion experiences positive feedback. In case (c), the rain has already accreted most of the cloud, and further depletion of cloud mass slows the accretion rate, producing negative feedback. In case (b), cloud and rain mass are comparable, and these effects are in balance, producing only weak positive or negative feedback.


ticular processes that we examined in more detail were rain evaporation, self-collection, and accretion of cloud water in more detail. Partially this was because the purely liquid processes were easier to examine in isolation from the other physics, and partly this was because in practice the vapor/ice transfer was usually associated with small timescales only when the effect of this process was small anyway, so the error due to finite time resolution was small. In particular, sublimation of a small ice mass in relatively dry air can be quite rapid, but leads to the same behavior as given by the limited behavior, namely that the ice sublimates completely and rapidly.

We can also note that almost no eigenvalues are associated with the external processes, which is what we expect since these process' rates don't directly depend on the MG2 inputs, and only interact with MG2's physics due to being affected by the conservation limiters. In our data set, we found only three eigenvalues associated with contact freezing, a few thousand associated with immersion freezing, and none with nucleation deposition or droplet activation.

Figure 7 shows the degree to which different processes were associated with the same eigenvalues (and hence timescales). Looking at timescales associated primarily with the processes listed above, we found the following:

1. Accretion of cloud water by rain is associated with autoconversion.
2. Rain self-collection is strongly associated with rain evaporation.
3. The degree of association between ice processes is quite complex. In particular, vapor/ice transfer is associated with several processes, possibly because it is mildly active in a very large number of grid cells.

As Figure 6 shows, the timescales associated with both types of autoconversion were relatively long. We therefore expected that the regime in which, say, liquid autoconversion and accretion interact most heavily would be different from the regime in which accretion by rain is associated with short timescales, and therefore we hypothesized that autoconversion should not require a short time step to adequately resolve, despite its association with accretion.




**Figure 6.** PDF of negative MG2 Jacobian eigenvalues, based on process primarily associated with each eigenvalue. Each row sums to 1, except for processes with no associated negative eigenvalues. The right-hand y-axis label shows the number of negative eigenvalues associated with each process.

On the other hand, rain self-collection and evaporation were both primarily associated with the same short time scales, and so it appeared less likely that we could resolve rain-related processes without using a relatively short time scale for both.

## 5 Decomposition by Weather Regime

Microphysics operates differently in different meteorological conditions. In particular, only a few microphysical processes are typically active at any particular point in space and time. Thus breaking cases down by weather regime can simplify the task of understanding model behavior. Such decomposition is also important because processes are likely to have different timescales depending on the weather regime they're in. For instance, the rate of accretion is fairly steady when rain and cloud mass are similar, but the accretion rate decreases rapidly when cloud becomes depleted, and the associated timescale is therefore much smaller in the latter case.





**Figure 7.** Association between pairs of processes in MG2. Each row shows the average association index ( $C$ ) for eigenvalues associated with a given primary process. Each entry on the diagonal has had its value reduced by 0.5, to fit in the same color range.

## 5.1 Methodology

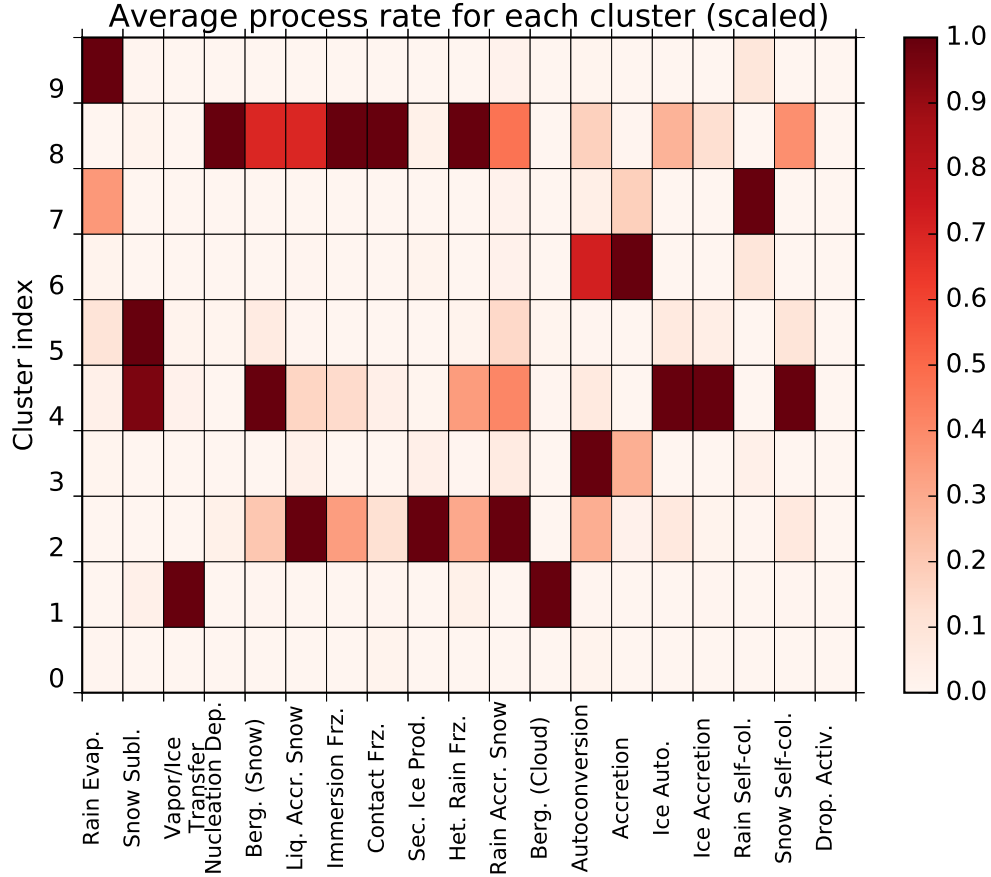
We accomplished this decomposition by normalizing all process rates by typical values (so that they were at the same order of magnitude) and using a simple  $k$ -means algorithm to generate clusters that we could treat as separate regimes.

We were interested in finding clusters containing qualitatively different, common types of behavior to examine in further detail, not in an objective categorization of all types of grid cells in the model. For this purpose we were content to use a degree of hand-tuning on both the relative weights of different process rates and the total number of clusters, until we found 10 clusters that had reasonably distinct process rates from one another.

Once this was accomplished, we could then look at the eigenvalues associated with the Jacobian for each cluster, in order to determine which regimes were likely to have short time scale behavior based on process rates alone.

## 5.2 Results


Figure 8 shows a representation of each cluster showing its most active processes. Cluster 0, which consists of generally clear-sky grid cells where no processes is particularly active, comprises the majority of grid cells.



**Figure 8.** Relative activity of each process in each cluster found via k-means.

Note that this provides another, simpler way to look at associations between processes. Looking again at liquid autoconversion and accretion, we see that these two processes are active together in clusters 3 and 6, but that the rate of accretion is much larger in cluster 6. We can then look at the eigenvalues of the Jacobian based on cluster, which is shown in Figure 9. Unsurprisingly, the negative eigenvalues are much larger in cluster 6, implying that the short timescales are associated with the heavy accretion present in cells with a relatively large in-cloud rain mass.

Similarly, we can use other clusters to examine other sets of short timescale processes. Cluster 1 is the best test case for short timescales associated with the Bergeron process and ice deposition. Rain evaporation and self-collection are both active and associated with short timescales in clusters 7 and 9.



**Figure 9.** PDF of negative MG2 Jacobian eigenvalues, based on cluster of grid cell where the eigenvalue was calculated.

## 6 Impact of Shorter Timesteps

In the previous sections we identified processes and combinations of processes that evolve much more rapidly than the default model step. In this section we test whether accurately resolving those processes has a big impact on model behavior.

### 6.1 Methodology

In this section, we focused on two clusters, each associated with a set of processes that we believed would change significantly if substepped so as to resolve their behavior:

1. Grid cells with large rain self-collection and evaporation rates, corresponding to heavy out-of-cloud rain. (Cluster 9)
2. Grid cells with large accretion rates and moderate autoconversion rates, corresponding to heavy in-cloud rain. (Cluster 6, filtered to remove grid cells with any sublimation/evaporation process tendency above  $1.e - 7\text{g/kg}$ )

The filtering of cluster 6 was necessary to remove cases where accretion and autoconversion were not the only processes with significant activity, since the clustering

algorithm alone did not separate those grid cells with only rain production from grid cells that combined rain production with evaporation/sublimation of precipitation falling from above.

For each of these cases, we produced modified versions of MG2 that allowed these processes to be run at a smaller time step using the forward Euler method. For instance, for heavy in-cloud rain, autoconversion and accretion were substepped inside a nested series of loops, so that it was possible to adjust the timestep of each independently of MG2 as a whole. If both processes were run at a finer timestep, it was also possible to independently adjust the coupling frequency. By adjusting these timesteps, it was possible to determine which processes needed to be better resolved to improve MG2’s accuracy, and which were less relevant.

In order to assess the accuracy of these substepped simulations, it was necessary to decide upon both a measure of error, and a “converged” result for comparison. To measure the error in each grid cell, we used the total water mass difference defined above ( $D_w$ ), which for Cluster 9 was effectively identical to the evaporation rate due to the absence of other influences. The converged result was chosen by using MG2 run at a very fine timestep of 75/512s, which is 1/2048 of the normal MG2 timestep of 300s.

## 6.2 Results

We hypothesized that for the grid cells with large rain evaporation, there would be a significant benefit from substepping rain self-collection as well, since our Jacobian-based analysis associates these two processes with the same short timescales.

Since these timescales are typically around 20s and we are still using the forward Euler method in the substepping loop, we also expect to see roughly first-order convergence for time steps shorter than this, but not for larger time step sizes, because for larger time steps this method (or rather, its linearization around a typical state) is not stable on this problem. The model therefore becomes increasingly dependent on limiters for longer time steps.


The actual evaporation/self-collection results are shown in Figure 10. Substepping evaporation by itself is indeed much less effective than substepping both processes together. The self-collection of raindrops is a very fast process in MG2, and accounting for this considerably reduces the rate of evaporation.

Furthermore, the rate of convergence of the model becomes first order at a time scale of roughly 30–60s, both when all of MG2 is substepped and when only only these two processes are substepped together. At very short time steps, the error levels off when substepping only evaporation and self-collection due to small contributions from other processes in the grid cell.

For grid cells dominated by accretion and autoconversion, we expected the effect of substepping accretion to matter much more than that of autoconversion, since autoconversion is associated almost exclusively with timescales longer than the maximum MG2 step size of 300s. Similar to the evaporation case, we expected that first-order convergence would occur only for timescales shorter than about 5s.

Figure 11 shows the results of this substepping of accretion and autoconversion in Cluster 6. As expected, substepping autoconversion alone has no benefit at all. Substepping the accretion can cut the errors by an order of magnitude, while substepping autoconversion as well improves the error by only another factor of 2.

However, we find, somewhat suprisingly, a near-perfect first-order convergence below a timestep of about 100s, a much larger timestep than we would have expected to be necessary to capture this process. We suspect that this is partially due to the fact that,




**Figure 10.** Convergence plot for mean error in cluster 9 grid cells under different substepping strategies for rain evaporation and self-collection. Runs substepping all of MG2 are shown in blue. Runs substepping evaporation and self-collection together, with total MG2 timestep held fixed at 300s, shown in green. Substepping evaporation alone, with all other processes at 300s, shown in red. The dotted reference line has slope 1.

while many grid cells have shorter associated time scales, the rate of accretion is largest in grid cells with large masses of both cloud liquid and rain, which have longer time scales.

## 7 Discussion

Our analysis of MG2, based on examining the eigenvalues of a numerically-derived Jacobian, shows that there are many situations where MG2 would appear to be unstable when using E3SM’s forward Euler method at a 300s timestep. It is stable in practice due to nonlinearities in MG2, and especially due to the presence of limiters, but we would expect these limiters to reduce the accuracy with which E3SM solves the equations that MG2 is intended to implement.

This leads us to three concerns. First, is the time discretization error first-order, so that we can readily trade off increased computational cost for a proportional decrease in error? Or are there key unresolved timescales present in the system, so that the solutions we find at a 300s time step qualitatively different from the converged results, e.g.



**Figure 11.** Convergence plot for mean error in cluster 6 grid cells under different substepping strategies for accretion and autoconversion. Runs substepping all of MG2 are shown in blue. Runs substepping accretion and autoconversion together, with total MG2 timestep held fixed at 300s, shown in yellow. Substepping autoconversion alone is shown in green, and substepping accretion alone is shown in red.

governed largely by limiters? Our results suggest that the answer to these questions depends significantly on the regime and the set of active processes involved. The time scales involved in the production and growth of snow, for instance, seem to be generally much longer than the 300s time step MG2 typically runs at. For processes involving rain or cloud ice, however, MG2 relies on limiters for stability, and in the particular case of rain evaporation, we have shown that the 300s time step is an order of magnitude too large to achieve first order convergence.

Our second concern is whether we can use a process-based analysis of MG2 to significantly improve accuracy without substepping the entire microphysics. Our results again suggest that it is possible to significantly improve accuracy in this way in certain regimes, but that it takes some care to do so effectively. For instance, substepping MG2's rain evaporation at a much shorter timestep, say 10s, produces a significant increase in model cost due to a number of *pow* calls, while only moderately reducing the error. But by also substepping rain self-collection (a less computationally-intensive process) in the same loop, the error can be reduced by an order of magnitude. We believe that sedimentation may also be relevant to these rain processes, and hope that a future study will investigate in-

cluding rain evaporation and self-collection as part of the sedimentation solver itself, rather than sequentially splitting sedimentation from all other microphysical processes.

Third, we can broaden our view and ask whether it is necessary at this point to accept the increased computational cost in order to improve MG2’s accuracy at all. To answer this question, we ran a simulation using E3SM v1 with MG2 substepped at a 1s time step (compset F1850C5AV1C-04P2, grid ne30\_ne30) for four simulated years. Figure 12 shows a Taylor diagram that compares the spatial variability of several key variables in this run to a control run with the MG2 default time step of 300s. The differences here appear to be quite minor. Figure 13 shows the spatial distribution of precipitation, which likewise is quite similar and shows no obvious systematic differences.

We can see significant differences, however, if we look specifically at the vertical distribution of rain mass. Figure 14 shows the difference between the zonally averaged in-area rain mass, which shows a reduction of rain in the upper cloud and near the surface, as well as a significant increase in the lower cloud. Our current interpretation of this is that this is due to a combination of reduced production of rain, reduced evaporation, and the effects of these differences on the rain fall speed. However, the exact reason for these differences is obscure, and in need of further study. We also believe that the coupling of sedimentation to these processes should be further examined.

## 8 Conclusions

By numerically calculating the eigenspectrum of the Jacobian of the MG2 microphysics, we were able to associate a set of time scales to these microphysical processes, and found that MG2’s equations are not always accurately modeled at the default E3SM time step. We were able to associate the short time scales to several processes, and demonstrated that in some regimes, decreasing the MG2 timestep can lead to significant differences in rain-related processes. We have also demonstrated that rain evaporation and self-collection should not be treated independently in MG2’s formulation, though accretion appears to be more independent.

The importance of these results, however, seems to depend significantly on the variables of interest. Several key variables show that we achieve essentially the same climate even with all of MG2 run at a short timestep, and in particular average rainfall is not significantly affected. However, there is some evidence that this is due to an effect involving cancellation of errors, since the vertical distribution of rain differs significantly at shorter time steps. Therefore we conclude that experiments which examine the cloud microphysics from a process-oriented perspective, and experiments looking at particular case studies (i.e. single-column runs) should consider the effects of time resolution on model results.


## Acknowledgments

Data will be provided through Argonne National Laboratory’s Petrel service. That transfer process is in progress, and a URL will be available shortly.

The TAR file at this location contains the following files:

1. MG2\_data\_collection.cam.h1.0001-01-06-00000.nc – Data set used to provide inputs to MG2.
2. MG2\_data\_collection.10\_cluster\_labels.0001-01-06-00000.nc – Labels input grid points according to clusters generated via k-means.
3. Jacobian\_cutoff\_\* – Eigenvalues derived from the numerical derivative of MG2 for each grid point. These were calculated in parallel and placed in 12 separate files.



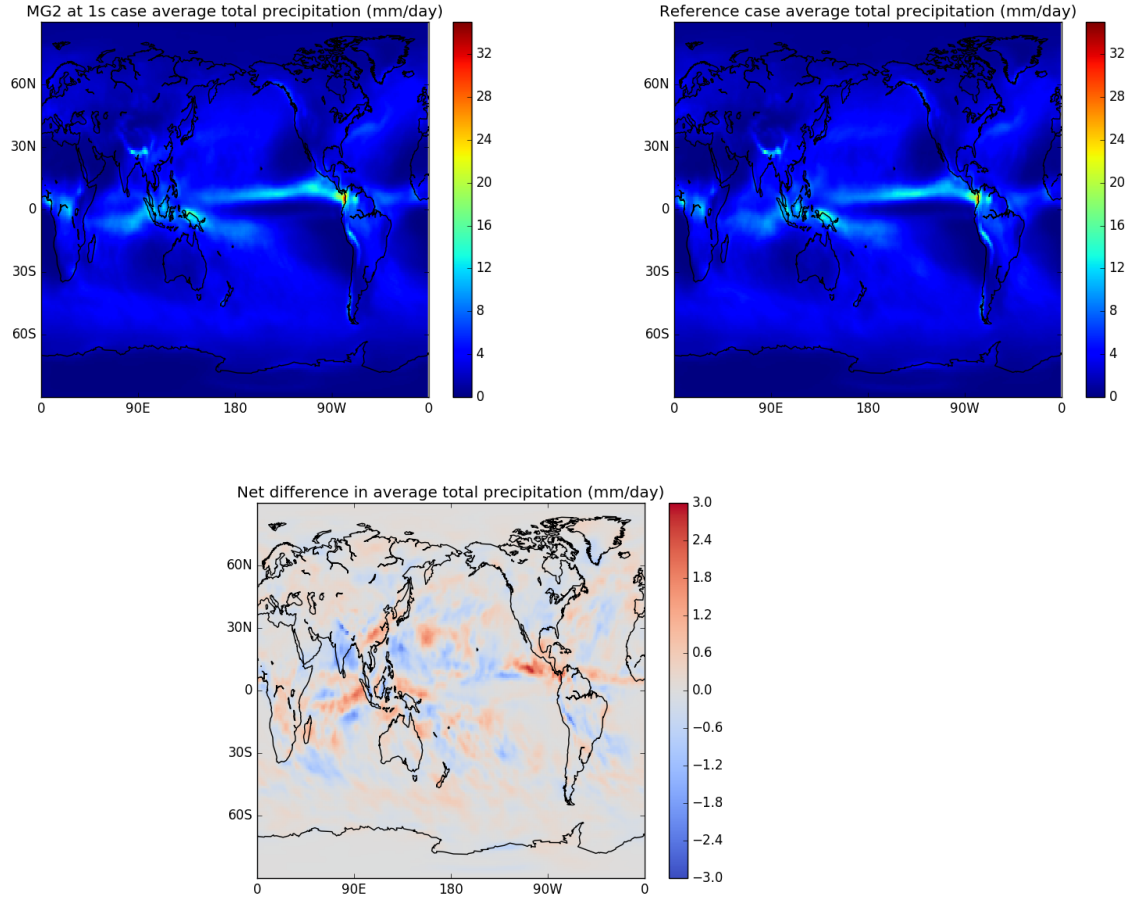


**Figure 12.** Taylor diagram comparing run with MG2 running at 1s to run at default settings. Only years 2-4 of the simulation were used.

- 624 4. nsubr\_control\_ANN\_climo.nc – Averaged data for control run used for Figures 12-14.
- 625
- 626 5. short\_timestep\_ANN\_climo.nc – Averaged data for short MG2 timestep run used for Figures 12-14.
- 627

628 This research was funded by the U.S. Department of Energy, Office of Science, Of-  
629 fice of Biological and Environmental Research.

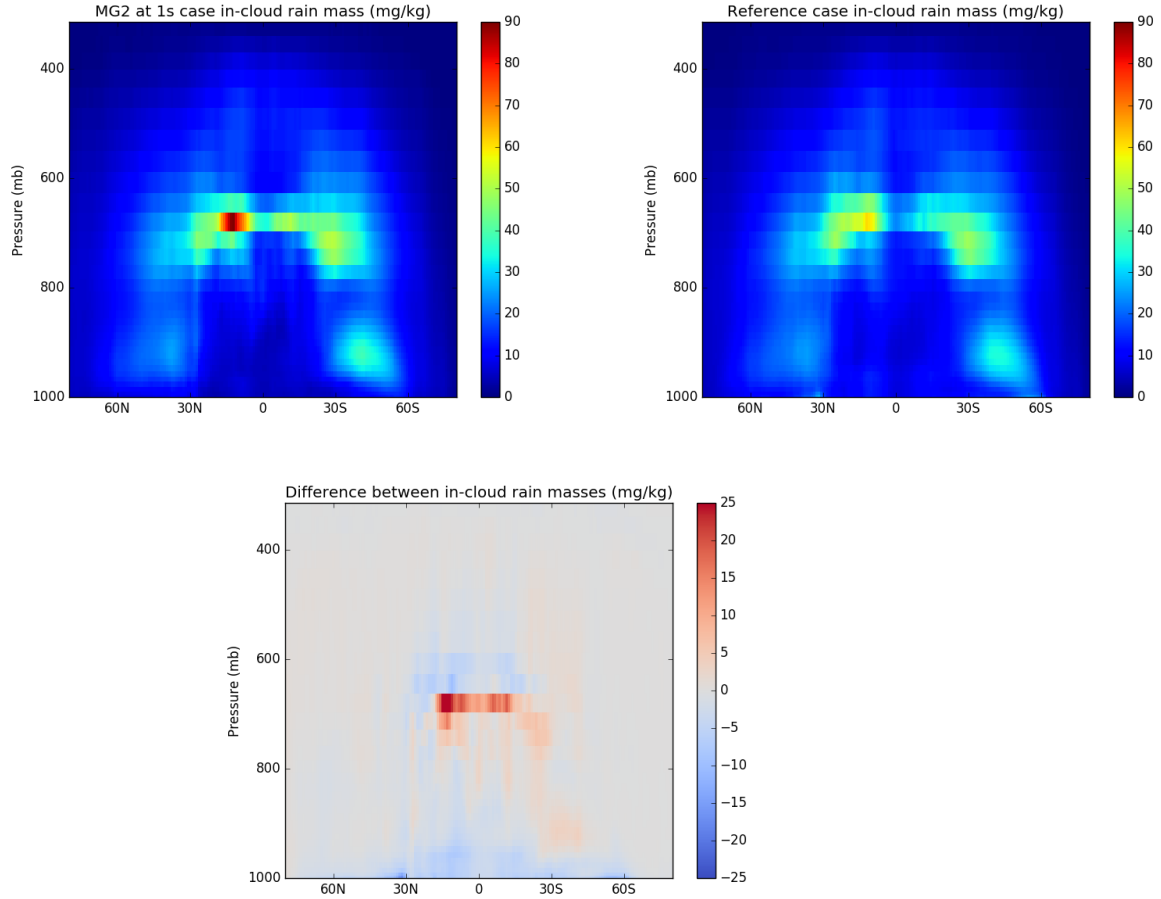
630 Sean Patrick Santos was further supported during this research by the Seattle chap-  
631 ter of the ARCS Foundation, which provided a scholarship to support his doctoral stud-  
632 ies.



**Figure 13.** Comparison of total precipitation at the surface using MG2 at a 1s (left) and 300s (right) timestep over three years, with the difference also plotted (center).

## References

- Gettelman, A., & Morrison, H. (2015). Advanced two-moment bulk microphysics for global models. part i: Off-line tests and comparison with other schemes. *Journal of Climate*, 28(3), 1268-1287. Retrieved from <https://doi.org/10.1175/JCLI-D-14-00102.1> doi: 10.1175/JCLI-D-14-00102.1
- Gettelman, A., Morrison, H., Santos, S., Bogenschutz, P., & Caldwell, P. (2015). Advanced Two-Moment Bulk Microphysics for Global Models. Part II: Global Model Solutions and AerosolCloud Interactions. *Journal of Climate*, 28(3), 1288-1307. Retrieved from <http://journals.ametsoc.org/doi/full/10.1175/JCLI-D-14-00103.1> doi: 10.1175/JCLI-D-14-00103.1
- Golaz, J.-C., Caldwell, P. M., Van Roekel, L. P., Petersen, M. R., Tang, Q., Wolfe, J. D., ... Zhu, Q. (2019). The doe e3sm coupled model version 1: Overview and evaluation at standard resolution. *Journal of Advances in Modeling Earth Systems*, 11(7), 2089-2129. Retrieved from <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018MS001603> doi: 10.1029/2018MS001603
- LeVeque, R. J. (2007). *Finite difference methods for ordinary and partial differential equations: Steady state and time dependent problems*. Society for Industrial



**Figure 14.** Comparison of zonally-averaged rain mass in E3SM runs using MG2 at a 1s (left) and 300s (right) timestep over three years, with the difference also plotted (center).

- and Applied Mathematics.
- Rasch, P. J., Xie, S., Ma, P.-L., Lin, W., Wang, H., Tang, Q., ... Yang, Y. (2019). An overview of the atmospheric component of the energy exascale earth system model. *Journal of Advances in Modeling Earth Systems*, 11(8), 2377-2411. Retrieved from <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001629> doi: 10.1029/2019MS001629
- Wan, H., Rasch, P. J., Taylor, M. A., & Jablonowski, C. (2015). Short-term time step convergence in a climate model. *Journal of Advances in Modeling Earth Systems*, 7(1), 215-225. Retrieved from <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014MS000368> doi: 10.1002/2014MS000368
- Wan, H., Rasch, P. J., Zhang, K., Qian, Y., Yan, H., & Zhao, C. (2014). Short ensembles: an efficient method for discerning climate-relevant sensitivities in atmospheric general circulation models. *Geoscientific Model Development*, 7(5), 1961-1977. Retrieved from <https://www.geosci-model-dev.net/7/1961/2014/> doi: 10.5194/gmd-7-1961-2014
- Xie, S., Lin, W., Rasch, P. J., Ma, P.-L., Neale, R., Larson, V. E., ... Zhang, Y. (2018). Understanding cloud and convective characteristics in version 1 of the e3sm atmosphere model. *Journal of Advances in Modeling Earth Systems*, 10(10), 2618-2644. Retrieved from <https://agupubs.onlinelibrary.wiley>

670 .com/doi/abs/10.1029/2018MS001350 doi: 10.1029/2018MS001350  
671 Zhang, K., Rasch, P. J., Taylor, M. A., Wan, H., Leung, R., Ma, P.-L., ... Xie, S.  
672 (2018). Impact of numerical choices on water conservation in the e3sm at-  
673 mosphere model version 1 (eamv1). *Geoscientific Model Development*, 11(5),  
674 1971–1988. Retrieved from [https://www.geosci-model-dev.net/11/1971/](https://www.geosci-model-dev.net/11/1971/2018/)  
675 2018/ doi: 10.5194/gmd-11-1971-2018