


**RESEARCH ARTICLE**

10.1029/2021MS002841

**Key Points:**

- We adapt Energy Exascale Earth Model (E3SM)/Community Earth System Model so most of a processor pool can operate on just a subset of demanding physics columns
- Load balancing theory finds optimal computational scale for regionalized high intensity physics work
- Multi-Domain cloud resolving model tests with Large-Eddy Simulation resolution embedded in 30% of the superparameterized E3SM succeed with few artifacts

**Correspondence to:**

 M. Pritchard,  
 mspritch@uci.edu

**Citation:**

Peng, L., Pritchard, M., Hannah, W. M., Blossey, P. N., Worley, P. H., & Bretherton, C. S. (2022). Load-balancing intense physics calculations to embed regionalized high-resolution cloud resolving models in the E3SM and CESM climate models. *Journal of Advances in Modeling Earth Systems*, 14, e2021MS002841. <https://doi.org/10.1029/2021MS002841>

Received 23 SEP 2021

Accepted 26 MAR 2022

**Author Contributions:**

**Conceptualization:** Michael Pritchard, Christopher S. Bretherton  
**Funding acquisition:** Michael Pritchard  
**Methodology:** Michael Pritchard, Walter M. Hannah, Peter N. Blossey, Patrick H. Worley, Christopher S. Bretherton  
**Project Administration:** Michael Pritchard  
**Software:** Michael Pritchard, Walter M. Hannah, Patrick H. Worley  
**Supervision:** Michael Pritchard

© 2022 The Authors. *Journal of Advances in Modeling Earth Systems* published by Wiley Periodicals LLC on behalf of American Geophysical Union. This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](#), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

# Load-Balancing Intense Physics Calculations to Embed Regionalized High-Resolution Cloud Resolving Models in the E3SM and CESM Climate Models

 Liran Peng<sup>1</sup> , Michael Pritchard<sup>1</sup> , Walter M. Hannah<sup>2</sup> , Peter N. Blossey<sup>3</sup> ,  
 Patrick H. Worley<sup>4</sup> , and Christopher S. Bretherton<sup>3</sup> 

<sup>1</sup>Department of Earth System Science, University of California, Irvine, CA, USA, <sup>2</sup>Lawrence Livermore National Laboratory, Livermore, CA, USA, <sup>3</sup>Department of Atmospheric Sciences, University of Washington, Seattle, WA, USA, <sup>4</sup>PHWorley Consulting, Oak Ridge, TN, USA

**Abstract** We design a new strategy to load-balance high-intensity sub-grid atmospheric physics calculations restricted to a small fraction of a global climate simulation's domain. We show why the current parallel load balancing infrastructure of Community Earth System Model (CESM) and Energy Exascale Earth Model (E3SM) cannot efficiently handle this scenario at large core counts. As an example, we study an unusual configuration of the E3SM Multiscale Modeling Framework (MMF) that embeds a binary mixture of two separate cloud-resolving model grid structures that is attractive for low cloud feedback studies. Less than a third of the planet uses high-resolution (MMF-HR; sub-km horizontal grid spacing) relative to standard low-resolution (MMF-LR) cloud superparameterization elsewhere. To enable MMF runs with Multi-Domain cloud resolving models (CRMs), our load balancing theory predicts the most efficient computational scale as a function of the high-intensity work's relative overhead and its fractional coverage. The scheme successfully maximizes model throughput and minimizes model cost relative to precursor infrastructure, effectively by devoting the vast majority of the processor pool to operate on the few high-intensity (and rate-limiting) high-resolution (HR) grid columns. Two examples prove the concept, showing that minor artifacts can be introduced near the HR/low-resolution CRM grid transition boundary on idealized aquaplanets, but are minimal in operationally relevant real-geography settings. As intended, within the high (low) resolution area, our Multi-Domain CRM simulations exhibit cloud fraction and shortwave reflection convergent to standard baseline tests that use globally homogenous MMF-LR and MMF-HR. We suggest this approach can open up a range of creative multi-resolution climate experiments without requiring unduly large allocations of computational resources.

**Plain Language Summary** The atmospheric physics parameterizations of traditional climate models do not spend radically different amounts of computational power and time across different parts of the globe. However, there are some physical processes, such as low cloud formation, that require much higher resolution than we currently use to model climate change. Using the traditional framework, this becomes too expensive because of computational limitations. In this work, we develop a way to efficiently enhance the interior resolution of a multi-scale climate model over only parts of the world. Specifically, we develop a way for a supercomputer to best split up the work of simulating the cloud physics over the parts of the world with coarser versus higher resolution so that it runs much faster. We also show that this task of using coarse resolution in some places and high resolution in others doesn't produce unwanted side effects.

## 1. Introduction

Users of the Community Earth System Model (CESM; Hurrell et al. (2013)) and the Energy Exascale Earth Model (E3SM; Bader et al. (2014)) currently enjoy limited options to flexibly vary the computational intensity of sub-grid physics calculations, such as focusing extreme, rate-limiting calculations within small geographic subregions of interest (Jansson et al., 2019). From a scientific standpoint this is unsatisfying. Consider for instance the practicalities of multi-scale climate modeling in which one aims to embed subsamples of plausible explicit moist convection. The problem is that the computational requirements of explicitly resolving cloud-forming eddies varies by orders of magnitude across geographic regimes. The largest of deep cumulonimbus clouds require 1–4 km horizontal resolution to resolve, but trade cumulus clouds require hundred-meter horizontal resolution

**Validation:** Michael Pritchard, Walter M. Hannah, Peter N. Blossey

**Visualization:** Peter N. Blossey

**Writing – review & editing:** Michael Pritchard, Walter M. Hannah, Peter N. Blossey, Patrick H. Worley, Christopher S. Bretherton

and stratocumuli require 10-m horizontal resolution to resolve faithfully. Handling all of these scales on a uniformly high-resolution Large-Eddy Simulation mesh is computationally impractical and expected to remain so for decades (Schneider et al., 2017). Recent developments in regional mesh refinement (Hagos et al., 2013), while appealing, also cannot fully meet the need. The resolution would need to be doubled at least seven times to nest smoothly from horizontal grid of 25- to 100-km in an exterior planetary atmospheric model to the 250-m needed to quasi-resolve boundary layer dynamics, and experience attests that considerable computational expense must be expended in the transition zones between the nested meshes. Meanwhile, physical trouble can also be expected in the lateral nesting transition regions given that conventional physics parameterizations have strong “gray-zone” sensitivities (Hagos et al., 2013) to exterior resolution. In this context it would be useful if individual users of Multi-scale climate Modeling Frameworks (MMFs; Grabowski (2004); M. F. Khairoutdinov and Randall (2001); Hannah et al. (2020)) could focus embedded resolution where they think it matters for specific problems of interest.

Unfortunately, even MMF simulations, which exhibit gentler sensitivities to external resolution (Kooperman et al., 2016a) and sidestep the typical resolution-nesting problems, cannot be reconfigured to vary regionally due to simple load-balancing issues associated with geographic workload variance that have yet to be fully solved. A symptom of this problem is that, with one exception (Jansson et al., 2019), all MMF tests have deployed cloud superparameterization (SP) *globally* despite evidence that classical SP mostly improves the simulated rainfall distribution near the equator and over summer continents (Kooperman et al., 2016a, 2016b), and despite the fact that even more costly forms of refined-resolution SP such as “ultra-parameterization” were designed to improve the simulation of convection over small regions of the planet where marine low clouds tend to occur. One might naturally wonder then why SP and high-resolution (HR) have not been deployed *regionally*, at reduced computational expense, for instance to unburden the cost of long, ocean-coupled simulations, or to open up computational room to afford a relaxation of the MMF’s historical idealizations (e.g., 2D, small-domain, coarse-resolution cloud resolving models [CRMs]). One key issue is that the limits of the *load balancing* software infrastructure inherited from the host CESM and E3SM climate models have prevented experimentation of this variety.

Load balancing is not a new issue. Global climate models have many levels of parallelism that can be exploited. Constituent component models (atmosphere, ocean, land, etc.) can be run in parallel, that is, on distinct subsets of processing elements (PE), with some or all of the other components, and load balancing focuses on how computing resources are allocated between the components (Dennis et al., 2012; Worley et al., 2011). For individual component models, parallelism is typically introduced first via a domain decomposition of the horizontal computational grid. For the numerical methods used in the E3SM and CESM for dynamics and tracer advection, compact geographical patches are most efficient. For the parameterized physics in the CESM and E3SM atmosphere models—which dominate the overall expense and scalability of MMFs (M. Khairoutdinov et al., 2005) owing to the SP approach of embedding expensive local cloud-resolving calculations—computation of vertical columns (fixed horizontal coordinates) are independent, and subsets of columns can be grouped and assigned to processing units, both to enhance vectorization and to improve load balance (Worley & Drake, 2005). This group need not be geographically contiguous (Worley, 2006): indeed an adjacent geographic assignment of columns into PE is generally not desirable as it introduces, for example, load imbalance due to vertical columns corresponding to daytime requiring more intensive (shortwave plus longwave) radiation calculation than nighttime columns (longwave alone). Columns from the same geographic region will also more likely exhibit similar simulation-dependent physical processes, such as microphysical calculations in locations frequently covered in cloud. Therefore, by default, the atmosphere models in the US climate models CESM and E3SM do not assign neighboring columns to the same processing element. Instead, a space-filling-curve or a longitude-latitude global ordering of columns are “dealt” to the PE in a wrap mapping, mixing up columns geographically.<sup>1</sup> Such an approach attempts to load balance by reducing the variability in computational cost between PE arising from the diurnal and seasonal cycles, columns over land versus ocean, columns over the arctic versus the equator, etc., in order to obtain improved performance.<sup>2</sup> This wrap map approach to load balancing assigns approximately the same number of columns to each processing element. If the average cost per column assigned to these PE (local average) is approximately the same as the average over all columns (global average) at a given instance in simulation time, then this mapping is near optimal. Whether this condition on the local and global average costs is achieved is difficult to determine, but, based on the known static sources of load imbalance, the CESM/E3SM approach is reasonable. Another justification for the approach is that, for the typical physical parameterizations used in the atmosphere in CESM and E3SM, the computational cost per processing element is most strongly a function of

the number of columns assigned, and assigning the individual columns in a load-balanced fashion subject to an equidistribution of the number of columns simply improves upon this initial optimization. In any case, the CESM/E3SM approach has proven to be better than not load balancing at all.

However, the performance of this approach to load balancing the physical parameterizations will degrade if the difference in the computational cost between individual columns is large. This becomes especially true as the number of PE approaches the number of physical grid columns. For instance, at extreme computational scale, when each processing element has only a single grid column, all PE will wait for the most expensive grid column to finish. Nonetheless, if the relative computational cost per column is known, a load-balanced assignment of columns to PE can instead be computed directly, not depending on the heuristics currently used. This interesting use case naturally emerges in superparameterized climate simulations in which different GCM grid columns are simulating different cloud regimes with different characteristic eddy scales.

The current load balancing heuristics likewise may be inappropriate for next-generation climate simulation for applications beyond global cloud feedback, where innovations in regionalized, high-intensity sub-grid physics could also be helpful. For instance, a human impact modeler may benefit from embedding regionally intensive atmospheric physics and embedded boundary layer calculations only along corridors that produce damaging extreme events to vulnerable populations. Those interested in climate feedbacks linked to long-range transport of biogeochemical feedbacks might want to embed high resolution near sources of emissions, perhaps including effects of urban building configuration, road traffic, or vegetation nearby megacities (Buccolieri et al., 2011), in addition to explicit treatments of storm dynamics responsible for scavenging along likely transport pathways. Similarly, an atmospheric chemist may wish to hyper-resolve chemical reactions (There is also an option to assign columns in pairs when doing the wrap mapping, where the pairs differ  $180^\circ$  in longitude and the same absolute distance in latitude from the equator but with opposite signs. For a longitude-latitude mesh, this is very effective at eliminating load imbalances due to diurnal and seasonal cycles. It has proven less useful for cubed sphere meshes or when using regional refinement. Note that using different domain decompositions for the dynamics and for the physical parameterizations incurs communication cost (Mirin & Worley, 2012), whether via explicit message passing interface communication or memory copying or both. There are also options in CESM and E3SM to retain the geographic decomposition assignment of columns used in the atmosphere dynamics for the physical parameterizations, or to perform the wrap map over only subsets of the columns and processing elements (local load balancing), in order to eliminate or decrease the communication cost of the per physics timestep redistribution of columns between the physics and dynamics phases of the computation. For current production simulations on current production systems, the performance improvement from global load balancing exceeds the overhead of the remapping, but this has not always been true, and the default load balancing option has changed over time) and processing of reagents by explicit turbulence just within subregions of the tropics where nitrous oxide, odd hydrogen, and their precursors from deep convection or stratospheric intrusions, are vital to ozone prediction (Pickering et al., 1992; Prather & Jacob, 1997). A tropical ecohydrologist may want to hyper-resolve the atmospheric boundary layer only over tropical rainforests to look at surface flux exchanges and land-atmosphere impacts and their climate teleconnections with minimal approximation. In each of the above examples, a large computational gap exists between a small pool of *heavily-loaded* grid columns on the physics side of a climate prediction code relative to a much larger pool of relatively *lightly-loaded* grid columns.

The purpose of this paper is to point out why the current load balancing strategies in the CESM and E3SM atmosphere models are ineffective for such scenarios and then demonstrate a technical strategy, in the context of superparameterized climate simulation, that can meet the need in a way that should also open a range of new, flexible options for the interested climate modeler, atmospheric dynamicist, or human impacts researcher.

The paper is structured as follows. In Section 2 we begin to build some theory and then demonstrate via a concrete example why the existing physics column load-balancing infrastructure in the CESM and E3SM can be inefficient when presented with highly regionalized, high-intensity physics computations. In Section 2.4, we present a general theory for how to load balance a binary mixture of high-intensity and low-intensity physics work, mainly by relaxing an assumption that a fixed number of columns be used in the PE that are used to parallel-decompose groups of physics columns. The theory predicts an optimal computational scale for a generic range of problems and its predictions are tested against actual performance measurements of the E3SM climate model.

In Section 3, we exploit this technique to investigate a new hybrid form of MMF that embeds high-resolution (MMF-HR) cloud-resolving model calculations over a fraction of the globe and a standard low resolution CRM (MMF-LR) elsewhere. Trade-offs of this new “Multi-Domain CRM” approach are then measured focusing on conditions such as artificial circulations that can set up near the low-resolution (LR)/HR boundary in the context of a simplified aquaplanet in which such issues tend to be maximally detectable. Informed by this experience, we design a scientifically relevant Multi-Domain CRM configuration and test its performance in real-world hindcasts assessed against satellite observations, commenting on the outlook for this approach to enable new kinds of cloud feedback simulation. The conclusions and broader outlook are summarized in Section 4.

## 2. Methods

Let us begin by defining some parameters, considering the example of a binary mixture of “lightly” ( $l$ )- versus “heavily” ( $h$ )-loaded columns. For a given GCM horizontal resolution, the total number of physical columns  $C$  is fixed, and we define the number fraction  $F$  of the world covered by the heavy columns as

$$F \equiv \frac{N_h}{C}, \quad (1)$$

where  $N_h$  and  $N_l$  represent the total number of heavy and lightly loaded columns respectively. That is, the total number of physical columns  $C$  can then be written as

$$C = N_h + N_l. \quad (2)$$

Let us view  $F$  as a flexible parameter that different climate modelers may wish to vary according to different applications. For simplicity we assume a geographically static load imbalance, that is, each physical column can be categorized as either a heavily ( $h$ ) or lightly ( $l$ ) loaded column based on a pre-determined binary map array read by the model during the initialization stage (e.g., see Section 2.4): we will not change the location and the number of heavily versus lightly loaded columns during the simulation. All physical columns are then distributed across  $P$  PE.

The next fundamental parameter is the computational intensity ratio,

$$\tau \equiv t_h/t_l, \quad (3)$$

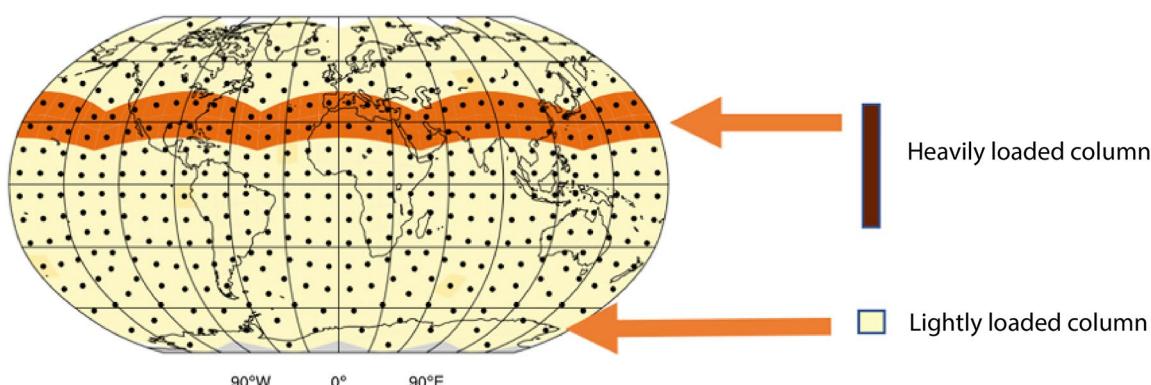
that is, the average compute time required for one processing element to integrate a single heavily loaded physics column ( $t_h$ ) divided by the time required for it to integrate a single lightly loaded column ( $t_l$ ).

### 2.1. Model Description

All simulations reported in this study use a superparameterized version of the E3SM (Hannah et al., 2020) that can be run with flexible exterior and interior resolution. In each GCM grid column, a laterally periodic 2D CRM aligned in the north-south direction replaces the conventional convective and stratiform cloud parameterizations allowing explicit computation of the global cloud fraction used for radiation computations (M. F. Khairoutdinov & Randall, 2001) down to the km-scale (or in the case of MMF-HR, 250-m) scale. The embedded CRM is based on the System for Atmospheric Modeling and uses a simplified one-moment microphysics (M. Khairoutdinov et al., 2005). In real-geography simulations, prescribed aerosol is based on present day climatological conditions. As to what is assumed beneath the CRM’s resolved scale, we do not run SP in conjunction with sophisticated sub-grid scale (SGS) parameterizations involving higher-order closures for the sub-km scale. Rather we use a classical approach that relies on a simple Smagorinsky 1.5-order TKE scheme to handle the SGS (Smolarkiewicz & Grabowski, 1990).

### 2.2. A Coarse-Resolution Example to Inform Optimization Trade-Offs

To illustrate the issues involved in load-balancing, we now investigate an unusually coarse resolution aquaplanet configuration3 of the superparameterized E3SM (Hannah et al., 2020) that can be run at a convenient



**Figure 1.** Physical columns (black dots) for a coarse resolution aquaplanet with a binary mixture of heavy versus light workload in the physics package. The dark orange belt represents locations in which physics columns with five times as much computational intensity ( $\tau = 5$ ) are imposed over a small fraction ( $F = 0.167$ ) of the globe (see text).

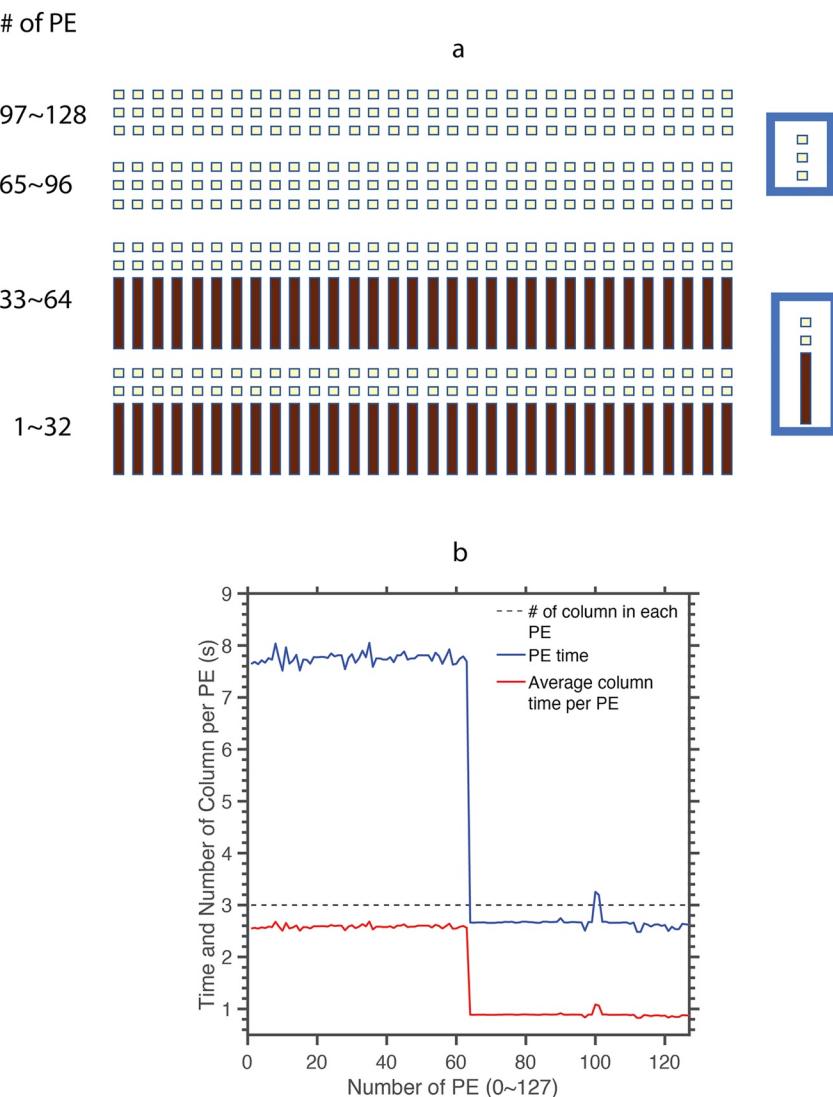
computational scale. In the coarse aquaplanet there are only  $C = 384$  physical columns spanning the globe (black dots in Figure 1), but it is nonetheless representative of the load balancing challenge at higher computational scales.

To create a binary mixture of heavily versus lightly loaded physics columns, we added the new capability to use two separate CRM grid configurations in separate physical grid columns. Heavily loaded columns are achieved by either increasing the number of embedded CRM columns or reducing the time step relative to the standard CRM grid configuration. In this way, we tested the effect on parallel decomposition of adding 5 times extra physics work ( $\tau = 5$ ) over a narrow latitude band ( $N_h = 64$  columns) shown by the dark orange belt (Figure 1). That is, the fraction of the world covered in heavy work  $F = 64/384 \approx 0.16$ . Note the icons at right in Figure 1, which will be referred to in later schematics—the dark red rectangles denote these heavy-loaded horizontal grid columns. Small yellow squares with five times less vertical extent denote the lightly loaded columns that require five times fewer calculations apiece. Our final assumption is that there are  $P = 128$  PE, i.e.,  $n \equiv C/P = 3$  physics columns per processing element.

The immediate problem with the E3SM's existing load balancing infrastructure is exposed by the schematic in Figure 2a, which summarizes how the binary mixture of heavily versus lightly loaded columns is assigned to each of the 128 PE by default. Since (by design in this example) the number of heavily loaded columns is not an integral multiple of the number of PE, not all tasks can be assigned an equal mixture of heavy versus lightly loaded columns. Instead, the first 64 PE are assigned a triad of one-heavy plus two-lightly loaded columns. Next, having run out of heavily loaded columns to assign, the remaining 64 PE are each assigned a triad comprising three lightly loaded columns, a comparatively smaller workload. The implication is a discrete jump in computational cost per processing element, that is, a load imbalance, confirmed by our own task-level timing measurements (Figure 2b). Practically, this means that the PE assigned light work will wait and idle until the PE assigned heavier work are done computing.

One unsatisfying solution to this problem might be to increase  $F$  such that a larger fraction of the planet is covered by heavy work until equal groups of work can be attained for each processing element. For instance,  $F = \frac{1}{3}$  that is,  $N_h = 128$ , or  $F = \frac{2}{3}$  that is,  $N_h = 256$  in this example would immediately avoid imbalance. Since  $N_h$  is now an integral multiple of the total number of PE  $P$ , this is compatible with the default load-balancing scheme. That is, these discrete values of  $F$  work because we have assumed  $\frac{C}{P} = \frac{1}{3}$  in this example. An equally unsatisfying solution would be to decrease the number of processors  $P$ ; while this could help remove the inefficiency it is incongruous with the idea of adding heavy regional work that is compensated with *high* computational scale, our goal.

Our solution to this problem, shown in Figure 3, is to relax the assumption that all PE have to work on equal-sized groups of physics columns. For the current example, this means allowing the first 64 PE to devote themselves entirely to integrating a *single heavily loaded physics grid column* apiece. This dovetails (The E3SMv1 used the same spectral element grid for dynamics and physics calculations, but the version discussed here uses a new method for putting physics calculations on a coarser finite volume grid (Hannah et al., 2021)) with the idea that

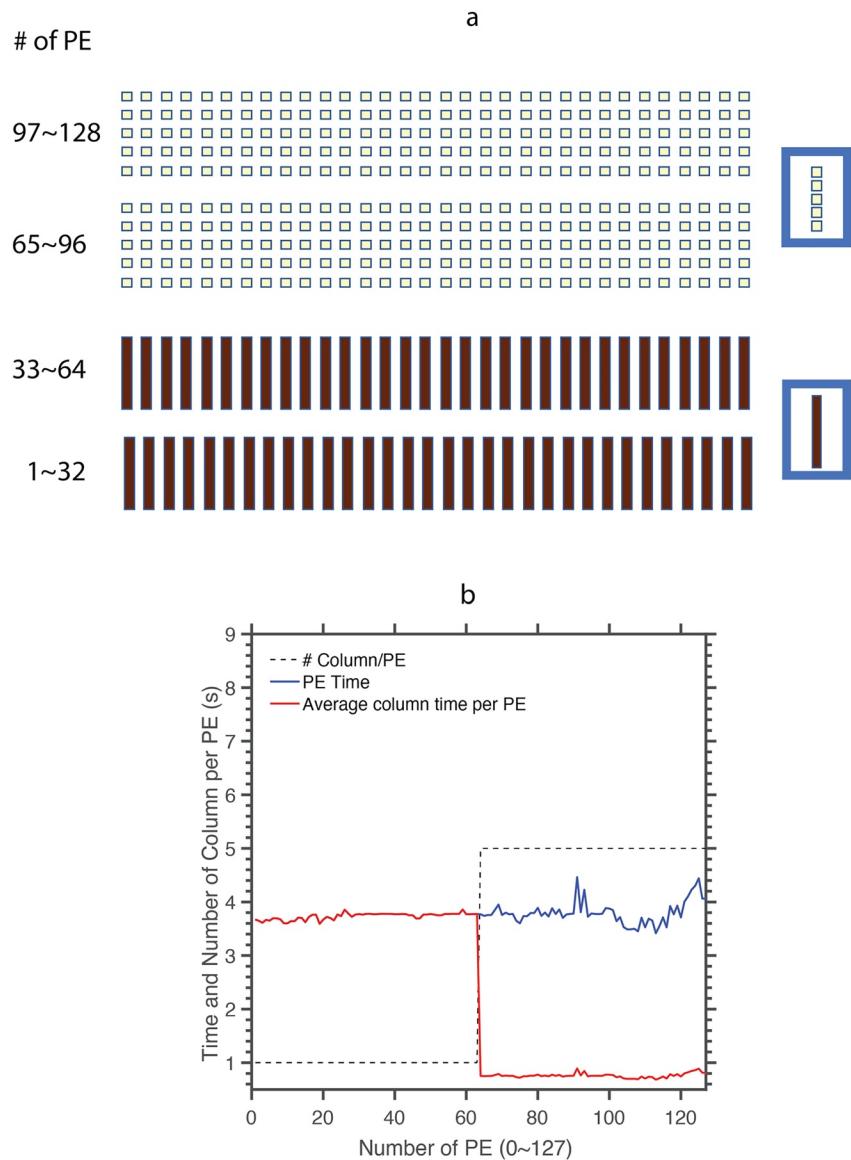


**Figure 2.** Schematic illustrating the default plan to (a) load-balance geographically heterogeneous physical column work by grouping them into processing elements (PE) and (b) the corresponding computational cost for one heavily loaded columns (PE 1~64) and two lightly loaded columns compared with three lightly loaded columns (PE 65~128).

such physics columns, being rate-limiting, deserve maximal computational resources. The remaining 64 PE each are then assigned equal sized groups of five lightly loaded columns. This heavy: light ratio ( $\tau = 5$ ) of 5:1 is a reflection of the assumed difference in computational intensity between the different regions and can thus change with problem definition, but helps illustrate the issues at play that must be considered. The new method can be adapted to other  $\tau$ . For instance, for higher values of  $\tau$  but the same  $F$ , the model could still assign 64 heavily loaded columns into the same first 64 PE but should increase the size of the groups of lightly loaded columns to approximately  $\tau$ , requiring less overall processors.

For our own setup, as illustrated by the timing results in Figure 2b, the overall strategy successfully achieves load balance for our chosen problem ( $F = 16.7\%$ ,  $P = C/3$ ,  $\tau = 5$ ), effectively allowing half of the processor pool to work on less than 20 percent of the planet with maximum throughput and without introducing inefficiencies, achieving in this case a 50%4.

Informed by the qualitative lessons learned in this concrete example, we will now derive some generalized constraints that allow a user to determine the full spectrum of load balanced configurations ( $F$ ,  $M$ ,  $\tau$ ) that this approach opens HR.



**Figure 3.** Similar to Figure 2 but using Multi-Domain cloud resolving model approach.

### 2.3. Generalized Theory of the Problem

Under the current assumption in CESM and E3SM that all PE must work on equal-size groups of physics columns, we have shown above that load-balanced conditions can be achieved for our binary mixture of heavily and lightly loaded columns, but only for a limited set of discrete conditions that will turn out to have unsatisfying general properties. To see this, let

$$n \equiv \frac{C}{P}, n = 2, 3, 4, 5 \dots \quad (4)$$

be the number of physics columns per processing element, assumed a fixed positive integer. The only possible path to load-balanced conditions is for each processing element to handle the same discrete mixture of heavy and light columns; in this case the cost ratio  $\tau$  is irrelevant but we can write

$$n = n_h + n_l \quad (5)$$

**Table 1**

*Constraints Connecting the Physics Columns per Task (Left,  $n = C/P$ ) With the Fraction of Earth That can be Covered in Heavy Work Under the Default Load-Balancing Setup of Energy Exascale Earth Model*

$n = C/P$	$n_h$	$n_l$	F
<b>2</b>	<b>1</b>	<b>1</b>	<b>1/2</b>
<b>3</b>	<b>1</b>	<b>2</b>	<b>1/3</b>
3	2	1	2/3
<b>4</b>	<b>1</b>	<b>3</b>	<b>1/4</b>
4	2	2	1/2
4	3	1	3/4
<b>5</b>	<b>1</b>	<b>4</b>	<b>1/5</b>
5	2	3	2/5
5	3	2	3/5
5	4	1	4/5
<b>6</b>	<b>1</b>	<b>5</b>	<b>1/6</b>
6	2	4	1/3
6	3	3	1/2
6	4	2	2/3
6	5	1	5/6

*Note.* Bold numbers highlight the minimum  $n_h$  value for each  $n$  category.

in which  $n_h \equiv N_h/P$  and  $n_l \equiv N_l/P$  are positive integers representing the number of heavy and light loading columns per processing element.

Table 1 summarizes the possible permutations of integers. In the simplest case,  $n = 2$  there is only one possibility:  $n_h = n_l = 1$  which implies  $F = \frac{n_h}{n_h+n_l} = \frac{1}{2}$ . That is, if using half as many processors as there are physical grid cells, covering half the world in heavy work is the only viable solution. For the next simplest case ( $n = 3$ ) two possibilities exist— $(n_h, n_l) = (1, 2)$  or  $(2, 1)$ . Thus if using one third as many processors as there are physical grid cells,  $F = \left(\frac{1}{3}, \frac{2}{3}\right)$  are both viable, as we found in the preceding section. Likewise, for  $n = 4$ ,  $F = \left(\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\right)$  are viable and for  $n = 5$ ,  $F = \left(\frac{1}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5}\right)$  work too, and so on. (Suppose we wish to perform a 1-year simulation using these lightweight toy example setups ( $F = 0.167$ ). In the default load-balancing the rate of the whole model is limited by those tasks with slowest throughput (approximately 8 wall-seconds per 30-min GCM timestep; Figure 2). That is, the cost of the default simulation in units of core-hours would be  $(P \text{ cores} \times 8 \text{ wall-seconds per timestep} / 3,600 \text{ wall-seconds per wall-hour} \times 48 \text{ timesteps per sim-day} \times 365 \text{ sim-days per sim-year}) \approx 200 \text{ core-hours}$ . In contrast the load-balanced setup (Figure 3) integrates twice as fast for the exact same problem on the same sized processor pool so it costs approximately 100 core-hours, that is, a savings of 50%. However this savings is problem-specific and would obviously change for different  $(F, \tau, C)$  but this calculation should help guide readers on how to predict their own cost savings ratio for specific problem setups).

The bold rows in Table 1 reveal the key problem: The *minimum viable fraction* of the planet that can be covered with heavy work is constrained by

$$F_{\min} = \frac{1}{n} = \frac{P}{C}. \quad (6)$$

so the problem is that for a fixed external grid resolution  $C$ , decreasing the fraction of the world covered in heavy work means sacrificing how many processors  $P$  are deployed. In short, the current load balancing infrastructure of the CESM/E3SM cannot deploy most of the processors on a small fraction of the planet experiencing especially high computational intensity.

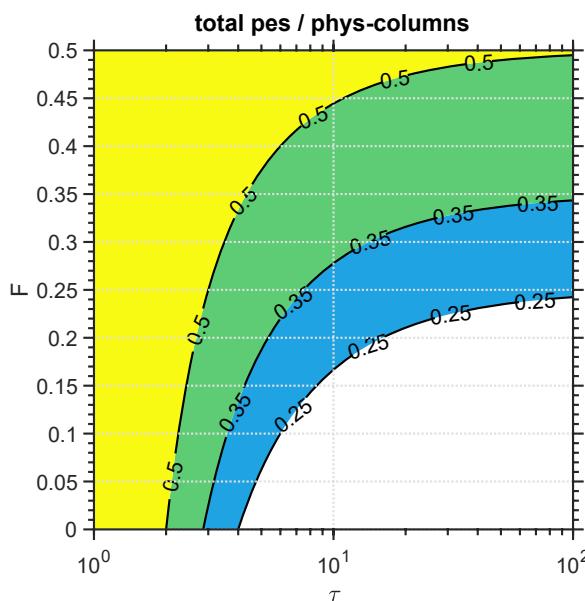
## 2.4. Generalized Solution

To solve this problem we separate the total PE into two pools of processors, that is,  $P = P_h + P_l$ . We allow a subset of processors,  $P_h = N_h$  to be assigned individual5 heavily-loaded columns, while the remaining processors handle multiple lightly loaded columns. Our new condition of load-balancing is then that the computational cost for *one* processor to integrate a *single* heavily loaded column should balance the total computational cost for an equivalent processor to integrate a discrete set of lightly loaded columns. Thus, by the definition of  $\tau$  (Equation 3),

$$P_l = P - P_h = \frac{N_l}{\tau}. \quad (7)$$

(Though it is not our focus, this assumption could easily be relaxed to create an even more general framework that additionally allows small groups comprising, for instance, two or three heavily-loaded columns, instead of one; this could be attractive for simulations with a very large number of total physical columns). Combining Equation 1, Equation 2, and Equation 7,  $P$  can be solved as

$$\frac{P}{C} = F + \frac{1-F}{\tau}. \quad (8)$$



**Figure 4.** The most efficient total number of processing elements, expressed as a fraction of the total number of grid columns ( $\frac{P}{C}$ ) contoured as a function of the fraction of the world covered by heavily loaded columns ( $F$ ) and the heavy: light workload ratio ( $\tau$ ). Calculation is based on Equation 8. Contoured values of  $\frac{P}{C}$  are labeled.

Equation 8 predicts the optimal number of processors for load-balancing given the fraction  $F$  of the world to be covered with heavy work, the intensity overhead  $\tau$  of that regionalized workload, and the total number of columns  $C$ . Figure 4 reveals the parametric dependence of  $\frac{P}{C}$  as a function of  $F$  and  $\tau$ , focusing on the regime of nontrivial computational ambition ( $P \geq C/4$ ) and scientific interest ( $F < 0.5$ ). While asymptotic behavior of  $\frac{P}{C} \approx F$  occurs for  $\frac{F}{\tau} \gg 1$  this is not a practically interesting regime, since at this limit nearly all PE ( $N_h \approx P$ ) are now working intensively on the HR columns, while the rest of the light columns are combined into one processing element ( $N_l \approx C - P$ ). At this limit, as few as one or two single processors would have to be assigned the majority of physical grid columns, which in our experience can result in such large array sizes that memory and its bandwidth become new limiting factors (not shown).

The more interesting and practical parameter dependence is observed for the regime  $2 < \tau < 100$  where Figure 4 provides some helpful guidance in simulation design. For instance, for regionalized heavy work that is ten-fold in its relative intensity ( $\tau = 10$ ), it is evidently strategic to choose  $F$  to be 17%, if one wishes to deploy one quarter as many processors as there are physical grid cells, whereas with twice as many processors,  $F \approx 26\%$  is a better choice. In our science context, when  $F$  is fixed by the problem at hand, analogous discrete predictions for the optimal simulation scale  $P$  can be made, as will be explored below. This illustrates the sort of practical considerations that the load-balancing theory can enable.

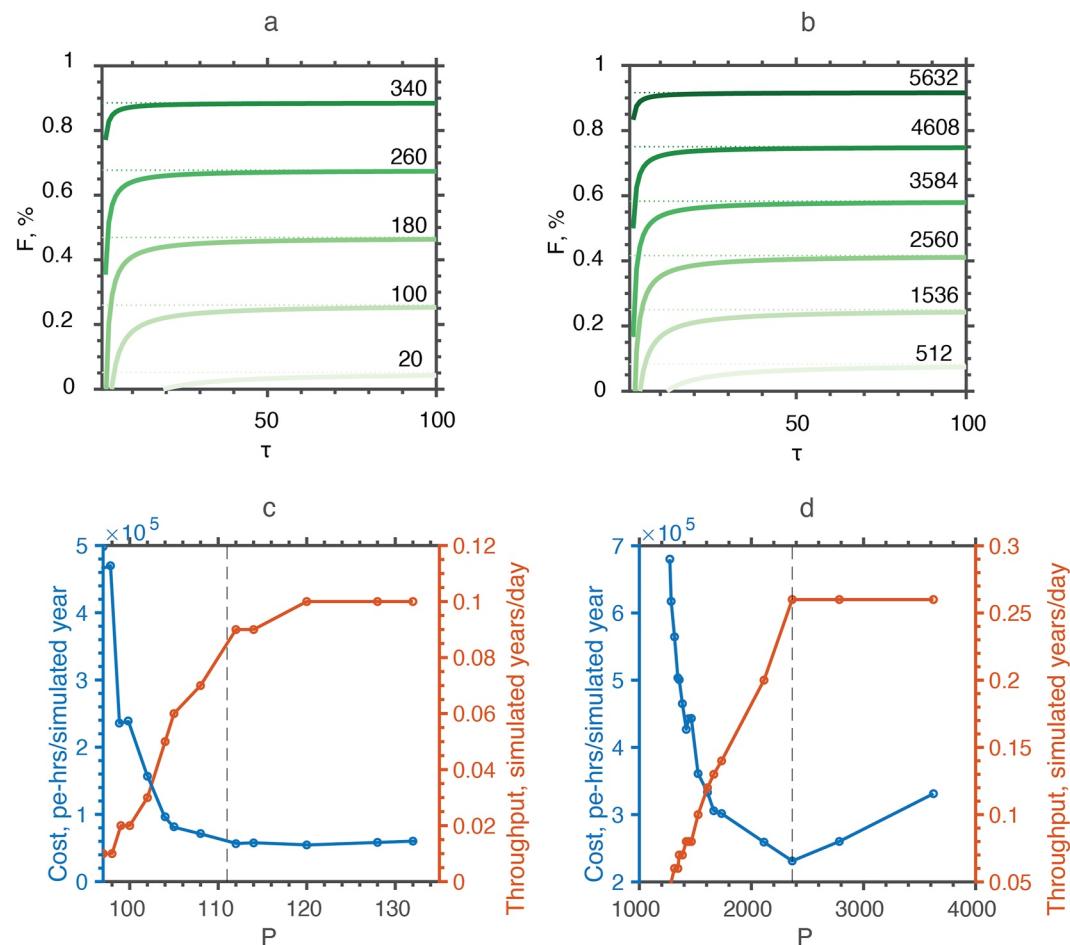
As a test of the theory, we now measure model throughput at computational scales in the vicinity of the predicted optimal  $P$  using two actual configurations of the E3SM climate model adapted to include our “Multi-Domain CRM” capability—an extremely coarse resolution aquaplanet (ne4pg2,  $C = 384$ , Hannah et al. (2021)) and an approximately 2.8-degree real-geography simulation (ne16pg2 horizontal resolution,  $C = 6144$ ) global grid. After substituting these values of  $C$  into Equation 8 the optimal total number of processors can then be viewed as a function of the additional science experiment parameters  $F$  and  $\tau$  (Figures 5a and 5b).

As in the previous example, workload imbalance is imposed using different, that is, “Multi-Domain” CRM grid configurations in different regions of the planet; the values of  $(\tau, F)$  are  $(20, 0.25)$  in the first experiment, and  $(4, 0.18)$  in the second experiment. The predicted optimal values of  $P$  using Equation 8 are shown by the vertical black dashed line in Figures 5c and 5d, which summarizes the measured performance statistics across a range of adjacent choices of  $P$ . Confirming the theory, the scalability results of both experiments show the predicted value indeed corresponds to the lowest model cost and approximately the highest model throughput. Using a larger than optimal  $P$  spreads the pool of light working columns across more PE, but the overall throughput is unchanged, since the rate is limited by the pool of  $P_h$  cores, thus needlessly increasing the total simulation cost in units of processing element-hours or equivalent electrical energy burden.

A caveat of this analysis is that it assumes the overall throughput is compute-rather than memory bandwidth- or communication-bound. While this is true to first order for superparameterized simulations, slight deviations from predicted theory are expected in Figure 5c to the extent that the memory bandwidth serves as a separate bottleneck, such as when too many lightly loaded columns are crammed on too few PE. For other classes of climate simulation in which there is not intensive, rate-limiting computational work within the physics package, the limits to performance and constraints on load balancing could be rather different and would require a different optimization approach.

### 3. Results

Whether it actually makes sense to use a binary mixture of CRM grids in operational climate simulation depends on how severe the artifacts induced at the grid transition boundary are. Such grid transition artifacts have been endemic to variable resolution models with localized mesh refinement and limited domain CRM simulations that

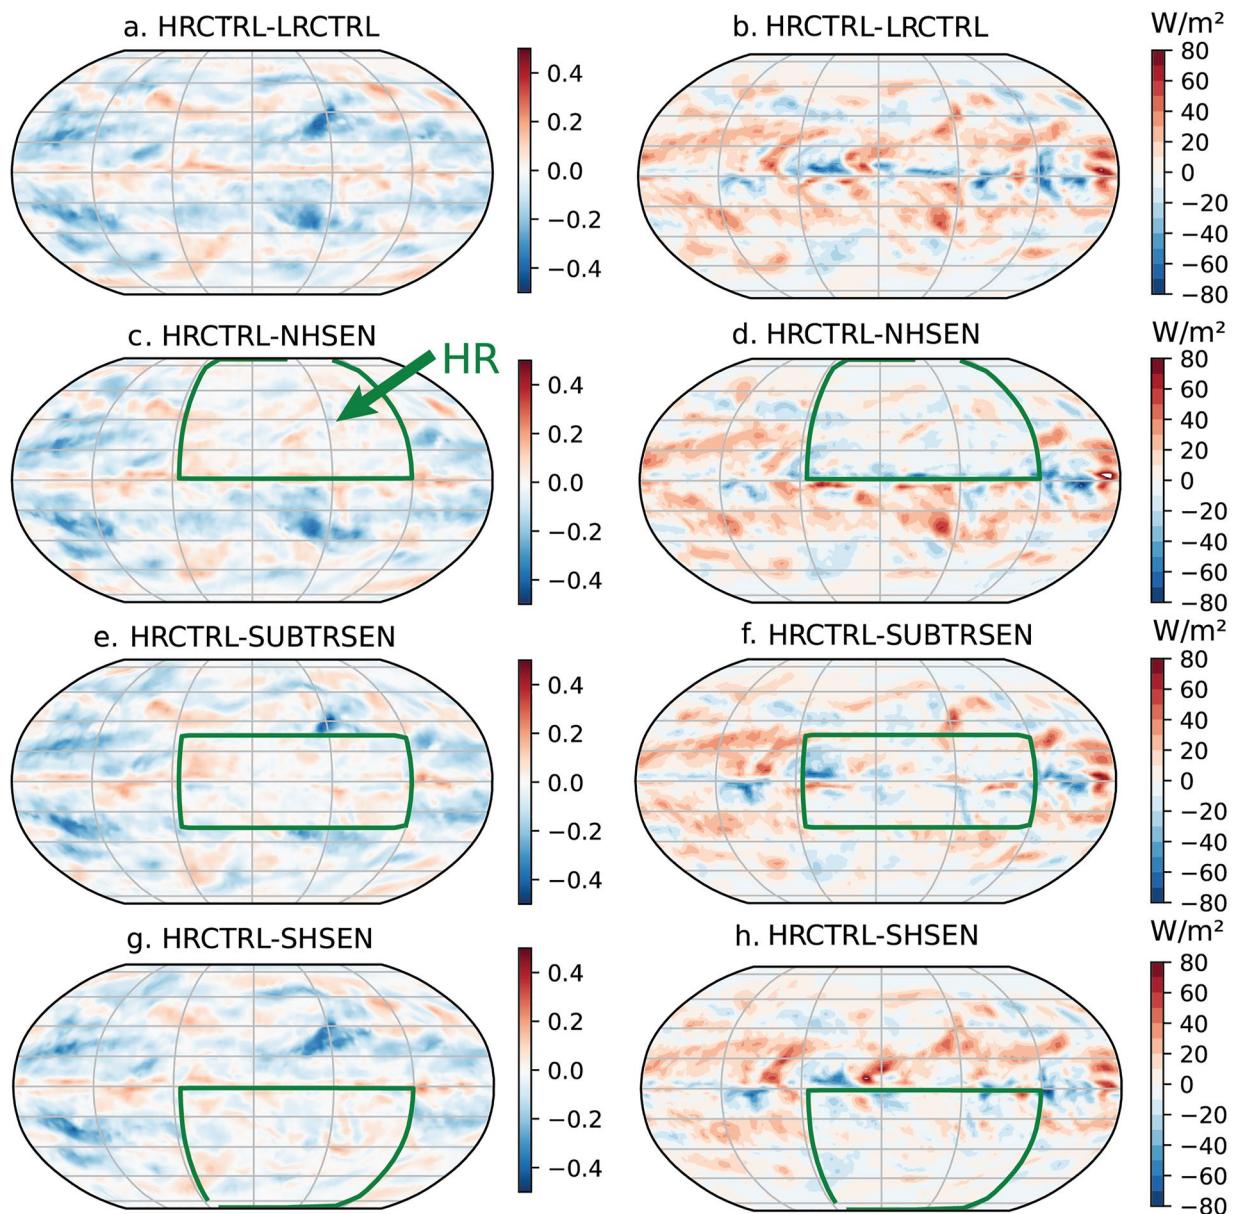


**Figure 5.** The most efficient total number of processing elements as a function of the fraction of the world covered by heavily loaded columns ( $F$ ) and the heavy: light workload ratio ( $\tau$ ) for (a) ne4pg2 and (b) ne16pg2 grids. Calculation is based on Equation 8. Contoured values of  $P$  are labeled. The computational cost (left y axis) and the model throughput (right y axis) for (c) ne4pg2 and (d) ne16pg2 grids. The vertical dash line represents the predicted value of  $P$  based on Equation 8.

exhibit a resolution dependence of precipitation and wind fields (Hagos et al., 2013; Rauscher et al., 2013). We therefore perform a series of experiments and analyze similar trade-offs of the Multi-Domain CRM approach, beginning with an aquaplanet to maximize detectability of artifacts relative to a statistically steady background state.

### 3.1. Aquaplanet Experiments Results

We performed five integrations using a ne16pg2 (total 6,144 physical columns) aquaplanet forced by prescribed zonally homogenous and meridionally symmetric SSTs (Neale & Hoskins, 2000) with a high resolution vertical grid resembling that used in Parishani et al. (2017) with 125 vertical levels approaching 20-m peak vertical resolution near the marine inversion. The first two baseline experiments apply global uniform 2-D CRM grid configurations—in low-resolution control (LRCTRL) a coarse CRM horizontal resolution of 1,200 m and 32 column CRM arrays, characteristic of classical SP, is used. In the second baseline test, high-resolution control (HRCTRL), a horizontal resolution of 200 m and 64 column CRM arrays are used instead, that is, the grid structure is identical to the “ultraparameterization” used in Parishani et al. (2017). The physical domain size of the LRCTRL simulation is three times larger than the HRCTRL but its computational domain size is half as large. Taking further into account a 10 times smaller CRM time step used for the HRCTRL (0.5 s) compared with LRCTRL (5 s), the computational intensity ratio  $\tau = 20$ —a factor of two for the added number of CRM columns and a further factor of 10 for the time step.



**Figure 6.** The low cloud fraction difference (left) between (a) high-resolution control (HRCTRL) and low-resolution control (LRCTRL), (c) HRCTRL and Northern hemisphere (NHSEN), (e) HRCTRL and Subtropics (SUBTRSEN), (g) HRCTRL and Southern Hemisphere (SHSEN). The absorbed shortwave radiation (right) difference between (a) HRCTRL and LRCTRL, (c) HRCTRL and NHSEN, (e) HRCTRL and SUBTRSEN, (g) HRCTRL and SHSEN. The high-resolution (HR) mask region (green arrow) is inside the area encompassed by the green line.

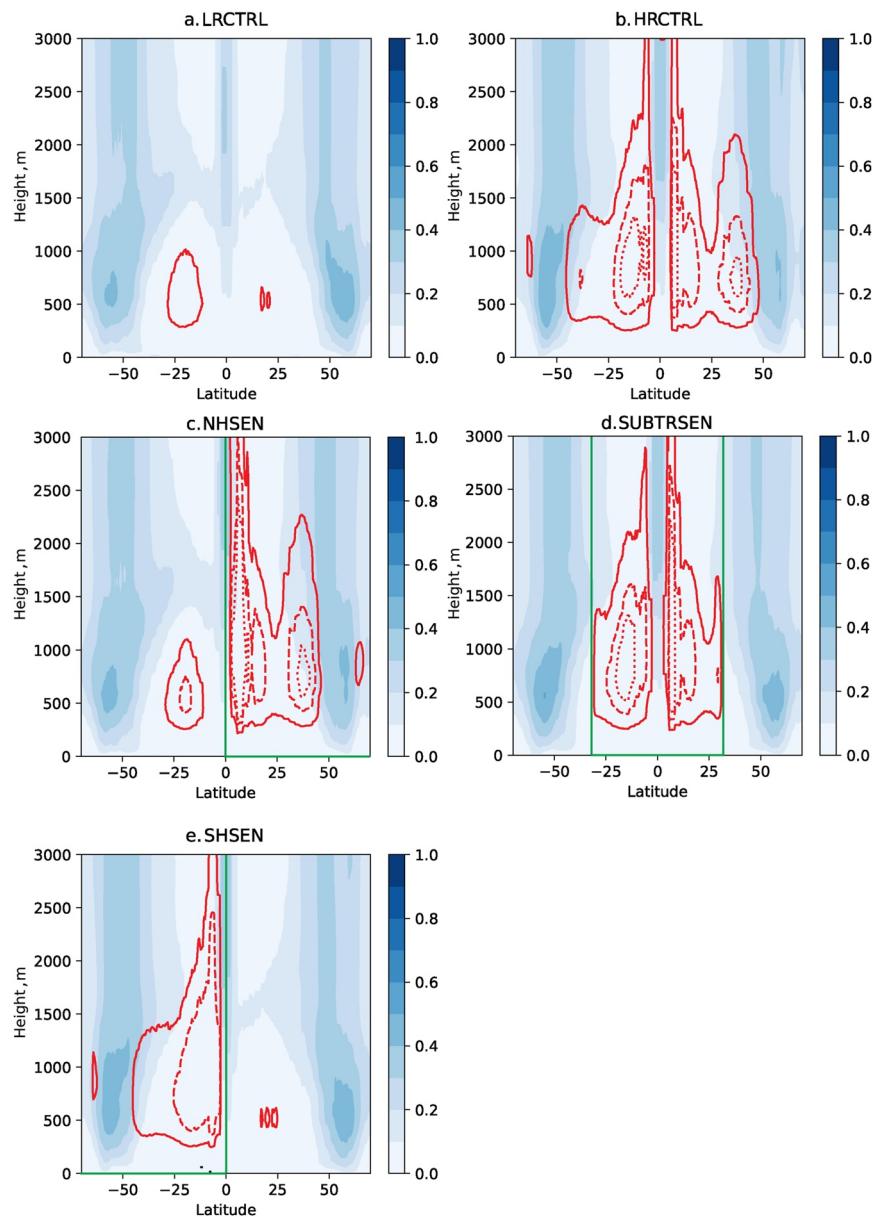
Three sensitivity tests then apply our Multi-Domain CRM approach using regionalized HR within just 1,504 of the available 6,144 columns ( $F \approx 0.25\%$ ) and using the LRCTRL grid configuration over the remaining three-quarters of the globe. The ability to use different vertical grids across the CRM instances is most likely possible, but the methods and infrastructure for this functionality will be difficult to implement, so for now we use the same high resolution vertical grid in all experiments for simplicity. The horizontal location of this region of heavy HR work (green box in Figure 6) is modified across three experiments that use varying meridional boundaries to produce three tests—Northern hemisphere (NHSEN), Subtropics (SUBTRSEN), and Southern Hemisphere (SHSEN) respectively. All simulations are 20 days in duration and produce the typical mean flow of an aquaplanet global climate model run with a realistic meridional surface temperature gradient—midlatitude westerlies and low-level tropical easterlies (Figure A3). Both LR and HR apply a classical SP approach.

A first look at maps of time-mean low cloud fraction and top-of-atmosphere (TOA) absorbed shortwave radiation (ASR) across the simulations shows that for the most part the Multi-Domain CRM method produces its intended effects locally with little distortion stemming from the grid transition boundary. We focus on low cloud fraction and TOA ASR based on previous studies that have found strong sensitivities to these properties across the same two horizontal CRM grid resolutions (Parishani et al., 2017), confirmed by our aquaplanet experiments. Figures 6a and 6b shows these expected baseline signals—as we refine the CRM horizontal resolution from 1200 (LRCTRL) to 200 m (HRCTRL), less cloud coverage results in more ASR, systematically. When HR is regionalized, the low cloud fraction and cloud brightness differences are just as expected locally - as evidenced by small anomalies between the Multi-Domain CRM and HRCTRL simulations within the green-boxed regions (Figures 6c–6f). Meanwhile, the difference between HRCTRL and Multi-Domain CRM outside the green-boxed region is nearly identical to the HRCTRL-minus-LRCTRL baseline anomaly pattern (Figures 6a and 6b). This comparison suggests that there is no systematic unintended cloud reactions due to the Multi-Domain CRM. The same finding is confirmed via a different quantitative evaluation metric combining physical columns inside and nearby the HR/LR mask boundary in Figure A1. The one potential exception occurs for the NHSEN and SHSEN experiments, which show some minor potential cloud brightness artifacts not predicted from the baseline simulations, which unless the result of internal variability may indicate issues when a grid transition is cavalierly placed directly on the equator (Figures 6d and 6h). If robust, this signal presumably associates with a deep convective response since it is visible in ASR but not in the low cloud fraction at left, which will be confirmed shortly. Note that some of the differences in the refined resolution subregion inevitably reflect internal variability; as such, a complementary view from just the first 7 days of the simulations in which memory of the initial conditions exists is available in Figure A2. As expected differences between HRCTRL and Multi-Domain CRM are even smaller on this initial timescale.

Note that some of the differences in the refined resolution subregion inevitably reflect internal variability, as opposed to long term feedback that, if present, must be second-order given the lack of meridional symmetry in the HR region between Figures 6c and 6g that would otherwise be expected. As such, a complementary view from just the first 7 days of the simulations in which memory of the initial conditions exists is available in Figure A2. As expected, differences between HRCTRL and Multi-Domain CRM are even smaller on this initial timescale, when less obscured by internal variability noise, and indeed could be viewed as negligible.

Zonally averaged vertical cross sections (Figure 7) of the vertically resolved cloud fraction and vertical velocity variance ( $w'w'$ ) in the lower troposphere also show intended changes. The cross section is chosen intentionally interior to the heavily loaded zonal subregion, with corresponding longitudes (latitudes) at bottom left and upper right of 160°E (70°S) and 85°W (70°N) respectively; the green vertical lines delineate the grid-transition boundaries of the three Multi-Domain CRM experiments. Within the regionalized HR location, the Multi-Domain CRM simulations (Figures 7c–7e) capture the expected  $w'w'$  enhancement (red contours) and subtle low cloud fraction reduction of the HRCTRL simulation relative to LRCTRL (i.e., a thin cloud layer near 1,500 m altitude) for all three sensitivity experiments. Our experiments also confirm that the differences between the HR and Multi-Domain CRM configurations are mostly minor within the regionalized HR location (inside the green box), while the differences between LR and Multi-Domain CRM are minor elsewhere (outside the green box). One minor artifact however is that the magnitude of  $w'w'$  within the HR mask area is slightly enhanced compared with HRCTRL, especially in the NHSEN and SHSEN experiments on just one side of the equator, again hinting at a secondary trade-off when using those configurations.

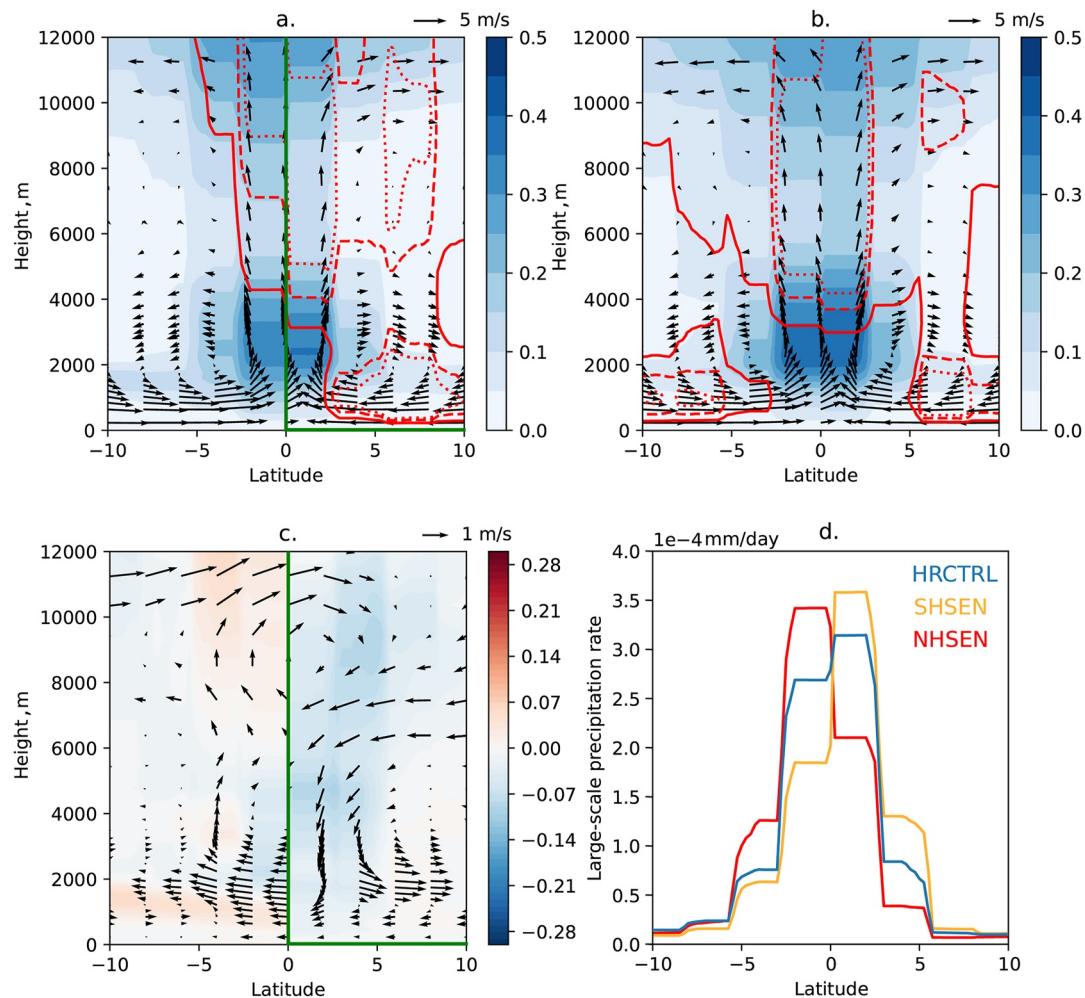
Based on the above findings, we hypothesize that when HR is regionalized in a hemispherically asymmetric manner, this otherwise equatorially symmetric aquaplanet is prone to exhibiting artificial ITCZ migrations coupled to deep convection and Hadley circulation cells. To demonstrate this danger of using Multi-Domain CRM, Figure 8 shows the total zonal mean meridional overturning circulation between 10°S and 10°N and the precipitation rate focusing on the extreme case of NHSEN where the HR/LR boundary is placed over the equator. In NHSEN (SHSEN) the HR grid is used exclusively in the northern (southern) hemisphere, causing it to dim preferentially due to HR's reduced low cloud fraction. The expectation (Hwang & Frierson, 2010) should then be a shift of the ITCZ to the relatively absorptive southern (northern) hemisphere, which is consistent with the zonal mean precipitation peaking south (north) of the equator in NHSEN (SHSEN) unlike LRCTRL and HRCTRL. Thus while the gross features of the Hadley cell appear similar in the NHSEN and HRCTRL simulations (panels a,b; low-level convergence and divergence at both mid- and high-levels), a close comparison reveals a ≈16%



**Figure 7.** The zonal averaged crosssection of the cloud fraction (shaded) for (a) low-resolution control (LRCTRL), (b) high-resolution control (HRCTRL), (c) Northern hemisphere, (d) Subtropics, and (e) Southern Hemisphere. The red contours mark the 0.5 (solid), 0.8 (dashed), and 1.0 (dotted) intra-cloud resolving model vertical velocity variance. The green lines represents the grid-transition boundaries.

magnitude anomaly circulation in NHSEN-minus-HRCTRL anomaly (Figure 8c), clockwise (via its upper- and lower-most branches) in the height-latitude plane, consistent with cross-equatorial flow anomalies carrying energy away from the warmer SHSEN. This same ITCZ response also predicts the cloud fraction and brightness changes noted earlier, where NHSEN tends to have less ASR on the southern flank of the equator (Figure 8a) (due to meridionally shifted reflection from repositioned deep ITCZ clouds) (Figure 6d) compared with HRCTRL.

The overall assessment of these aquaplanet results is that regionalized HR produces cloud brightness and fraction statistics remarkably similar to global HR. While grid transition artifacts exist, they do not produce major biases that would argue against the Multi-Domain CRM methodology. We can induce artifacts especially in a delicate aquaplanet setting by placing CRM grid transitions right on the equator (SHSEN, NHSEN) or in ways that imply

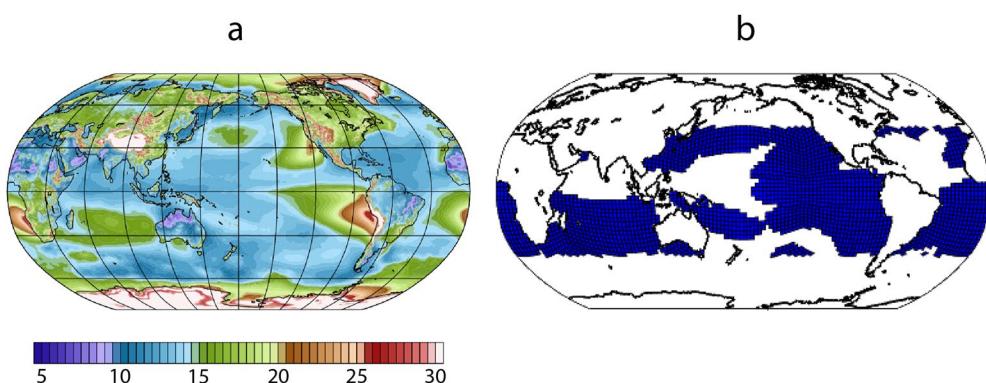


**Figure 8.** The zonal mean cloud fraction,  $v$  and  $w$  wind components for (a) Northern hemisphere (NHSEN), (b) high-resolution control (HRCTRL). The zonal mean cloud fraction,  $v$  and  $w$  wind components differences between (c) HRCTRL and NHSEN. The zonal mean precipitation rate for (d) NHSEN (red solid line), Southern Hemisphere (yellow solid line) and HRCTRL (blue solid line). The red contours mark the 0.5 (solid), 0.8 (dashed), and 1.0 (dotted) vertical velocity variance. The green lines represents the grid-transition boundaries for HNSEN. Vertical velocities have been multiplied by a factor of one hundred for visual clarity.

hemispheric asymmetry, that is, demanding ITCZ shift responses, but we can easily avoid them with a more judicious choice of the regional HR placement (SUBTRSENS).

### 3.2. Real-Geography Hindcast Experiments Results

The above lessons inform our strategy in the next phase of analysis that transitions to a real-geography hindcast simulation setup. Unlike the symmetric aquaplanet, this class of experiment does not have any idealizations of meridional symmetry and should be expected to be less delicate, but contains additional degrees of freedom for grid transition artifacts requiring independent investigation. Similar to our aquaplanet simulations, there are 6,144 physical columns spanning the globe in these new tests, but now  $F = 30\%$  (2,048 columns) of the planet use the HR configuration; thus, to perform the simulations efficiently, we used Equation 8 to obtain the optimal computational scale for this problem as  $P = 2,048$ . The model is initialized with interpolated ECMWF Reanalysis 5th Generation (ERA5) reanalysis data and forced with prescribed sea-surface temperatures from the NOAA Optimally Interpolated daily SST data set (Reynolds et al., 2007) for the given initial date, persisted in time. Five separate 7-day simulations are performed using 1 October initial conditions taken from independent years spanning 2008–2012.



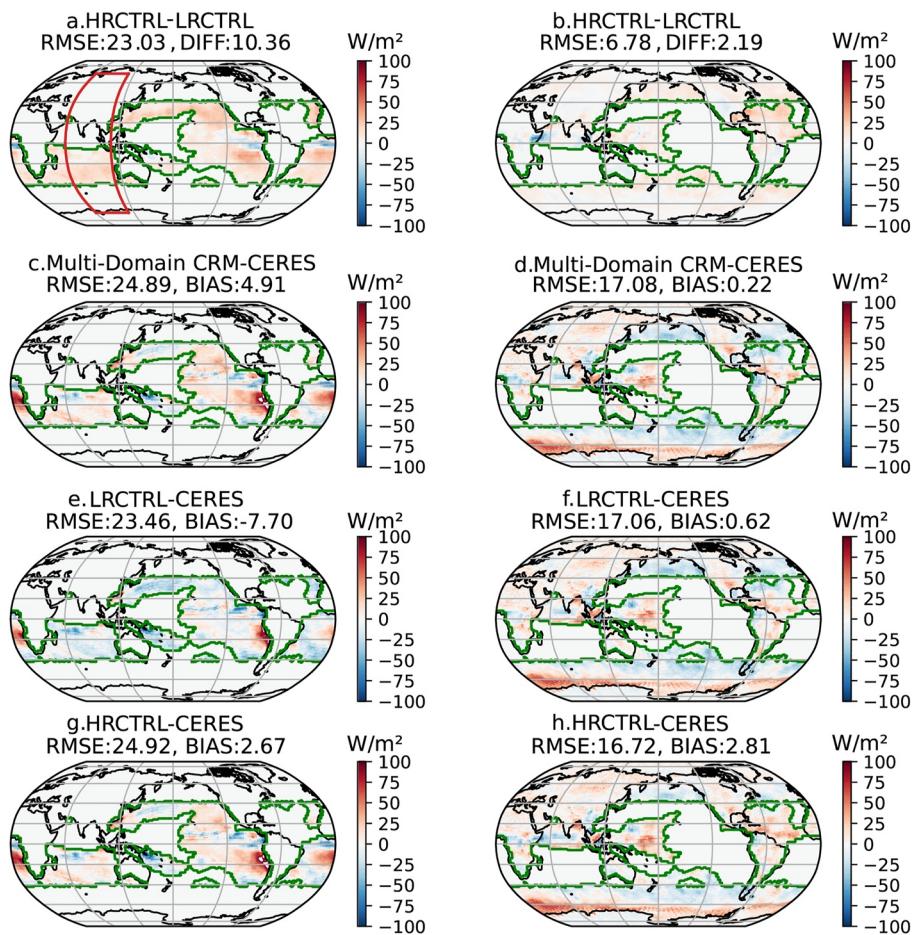
**Figure 9.** October climatology (2008–2018) of (a) average lower tropospheric stability from ECMWF Reanalysis 5th Generation reanalysis and (b) derived Multi-Domain cloud resolving model grid transition boundary for our real-geography hindcast tests. The blue area uses 200-m horizontal grid spacing, while the white area uses 1200-m horizontal grid spacing.

To define a physically strategic horizontal boundary between heavy and light cloud-resolving calculations, we used the lower tropospheric stability (LTS) as a metric to isolate regions of shallow convection, which deserve high computational intensity due to control by fine-scale eddies (Wyant et al., 2009). The LTS is defined as the difference between the potential temperature at 700 hPa and the 2-m surface air temperature; its October climatology from 2008 to 2018 based on ERA5 is shown in Figure 9a. Based on the LTS, we define a horizontal mask (Figure 9b) to contain the heavily loaded (HR) work, confined within 40 degrees of the equator. The mask covers the marine subtropical trade Cu regions in the South Atlantic, Indian Ocean, and East Pacific regions in higher-resolution CRM grid configurations. The total area covered by the HR mask is 30% of the globe by design. While the LTS conditioning results in a CRM grid transition boundary that is mostly hemispherically symmetric, it does include one subregion in the Indian Ocean that contains a near-equatorial meridional boundary; we will keep this subregion in mind based on lessons learned from the more homogeneous aquaplanet tests.

We now compare the results of control MMF-HR, MMF-LR and Multi-Domain simulations, focusing on the ensemble mean of the 7-day hindcast climatology. Figure 10 shows the ASR biases relative to regridded Clouds and the Earth's Radiant Energy System-SYN daily mean estimates from satellite (Wielicki et al., 1996). The LRCTRL and HRCTRL simulations create different characteristic patterns and magnitudes of shortwave biases outside/inside the heavily-loaded sub-region; we will return to this point shortly. As in the aquaplanet, and as expected from Parishani et al. (2017), as we increase the CRM horizontal resolution, the HRCTRL simulation tends to produce less low cloud fraction (Figure A5) and a positive ASR (dim) bias compared with LRCTRL (Figures 10a and 10b). The LRCTRL simulation tells us that LR clouds are systematically too bright in the mid-latitude and trade cumulus regions, with the exception of some stratocumulus dim biases (Figures 10e and 10f), whereas the HRCTRL simulation has less severe bright biases in the mid-latitudes and a *dim* bias ( $0.22 \text{ W/m}^2$ ) throughout most of the subtropics.

It is convenient that these two control simulations have different structures because it allows us to quantitatively test the Multi-Domain CRM method by separately calculating the area-weighted global root-mean square error (RMSE) within those two regions; see subpanel titles in Figure 10. Aggregated *within* the HR masked region (interior to the green contour), the spatial root-mean-squared error of the Multi-Domain CRM ensemble mean hindcast bias pattern (Figure 10c, RMSE =  $24.89 \text{ W/m}^2$ ) is a closer match to the HR-control simulation (Figure 10g, RMSE =  $24.92 \text{ W/m}^2$ ) for the same region compared to the LR control simulation (Figure 10e, RMSE =  $23.46 \text{ W/m}^2$ ). Meanwhile, the Multi-Domain CRM looks very similar to the LRCTRL simulation outside the trade Cu region. It has the same brightness bias (Figure 10d, BIAS =  $0.22 \text{ W/m}^2$  and Figure 10f, BIAS =  $0.62 \text{ W/m}^2$ ). That is, Multi-Domain CRM (Figure 10d, RMSE =  $17.08 \text{ W/m}^2$ ) has similar RMSE compared with the LRCTRL (Figure 10f, RMSE =  $17.06 \text{ W/m}^2$ ) simulation. A complementary analysis of the lower magnitude outgoing longwave radiation biases and RMSE statistics is available in (Figure A4).

Recalling that in the Indian Ocean subregion, we expect the Multi-Domain CRM to produce artificial cross-equatorial circulations, we now measure their magnitude. Based on the 7-day ensemble mean, the zonal average of the ensemble mean cloud fraction and vertical velocity variance are shown in Figure 11, with corresponding



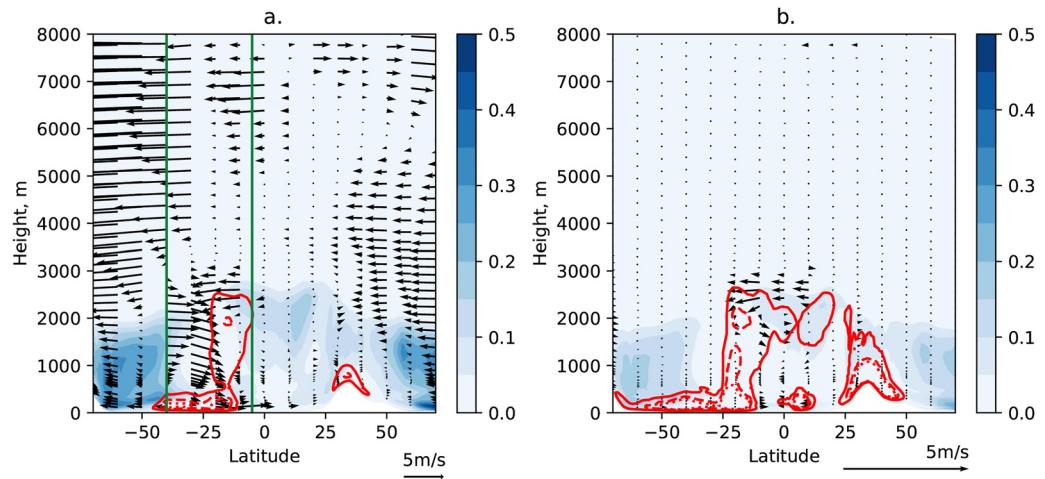
**Figure 10.** The map of the ensemble mean absorbed shortwave radiation differences between high-resolution control (HRCTRL) and low-resolution control (LRCTRL) for (a) inside (b) outside the high-resolution (HR) mask area, between Multi-Domain cloud resolving model (CRM) and Clouds and the Earth's Radiant Energy System (CERES) observations for (c) inside (d) outside the HR mask area, between LRCTRL and CERES observations for (e) inside (f) outside the HR mask area, and between HRCTRL and CERES for (g) inside (h) outside the HR mask area. The HR mask area is encompassed by the green solid line. The zonal averaged region shown in Figure 11 is encompassed by red solid line in panel (a).

longitudes (latitudes) at bottom left and upper right of 60°E (70°S) and 110°E (70°N) respectively (Figure 10a). Within the cross section (interior to the green vertical lines) both the cloud fraction and vertical velocity variance from Multi-Domain CRM simulation are similar to HRCTRL (compare contours and shading in Figures 11a and 11b). The circulation is mostly similar in the Multi-Domain CRM (Figure 11a), but the anomaly vector field in Figure 11b does reveal a weak anomaly of the zonal mean circulation. Its magnitude is only  $\approx 4\%$ .

#### 4. Discussion and Conclusions

By exploring an unusual configuration of a superparameterized climate simulation that uses a binary mixture of heavily- versus lightly-loaded physics columns, we have identified a limitation of the current parallel load balancing infrastructure in the E3SM and CESM: Extreme, rate limiting calculations embedded on the physics side of the code cannot be regionalized to small fractions of the planet in ways that make efficient use of ambitious computational resources.

We have solved the technical problem by relaxing the current assumption that all PE be assigned equal sized groups of physics columns. This allows rate-limiting regionalized calculations to be unthrottled to their maximum throughput (at the upper limit of one pe per physics column) while balancing the remaining load through unusually large (many physics columns per pe) column assignments elsewhere. This concept was motivated by



**Figure 11.** The 7-day ensemble mean zonal mean cloud fraction,  $v$  and  $w$  wind components for (a) Multi-Domain cloud resolving model (CRM). The ensemble mean zonal mean cloud fraction for (b) high-resolution control (HRCTRL) and,  $v$  and  $w$  wind components difference between Multi-Domain CRM and HRCTRL. The red contours mark the 0.5 (solid), 0.8 (dashed), and 1.0 (dotted) vertical velocity variance. The green vertical lines represents the grid-transition boundaries.

processor-level timing measurements applied to a coarse-resolution aquaplanet example, which illustrated the constraints. This informed a general theory to predict optimally load balanced and performant conditions under the new assumption of a binary mixture of heavily- versus lightly-loaded physics columns. The theory (Equation 8) predicts the optimal scale  $P$  for a given climate model with total grid columns  $C$  containing a fraction  $F$  of heavily-loaded grid columns, each with work overhead  $\tau$ , to achieve maximal throughput and minimal total expense. Predictions are successfully validated by actual timing measurements in test configurations that vary both  $F$  and  $\tau$ .

Our own narrow scientific motivation in developing this capability has to do with liberating computational resources to meet the cost and throughput needs of Multi-scale climate Modeling Framework (MMF) simulations of explicit low cloud feedback. The embedded CRM resolution requirements of these physics are challenging, since faithfully simulating shallow clouds requires resolving small-scale (less than 100-m) turbulent eddies in the boundary layer. This is only marginally possible on global scales, even with efficiency of MMF, due to expense constraints on the testable resolution and dimensionality of the embedded CRM arrays. As such, recent attempts at high-resolution MMF (MMF-HR; Parishani et al. (2017, 2018); Terai et al. (2020)) are physically unsatisfying compared to the large eddy simulations that inspire them: The embedded turbulence arrays are too small and low-dimensional to exhibit appropriate cellular cloud formation organization, and do not allow enough room for organized structures to allow a seamless transition from shallow to deep convection. The interior grid resolution (20-m vertical spacing, 200-m horizontal) is still too coarse to resolve the spectrum of boundary layer eddies that we would like. Yet if, as we desire, we increase the dimensionality, refine the resolution, or extend the domain size, the computational cost becomes too high for the long multi-month simulations needed to study aerosol-cloud feedback.

With the Multi-Domain MMF we have introduced here, and the new load-balancing theory that enables it, these problems can be somewhat offset by the cost mitigation of regionalizing HR to small fractions of the planet. Provided this does not induce unintended consequences, the technique should then allow historical idealizations of HR to be relaxed while minimizing the computational burden.

Anticipating such applications, we thus performed a set of simulations to examine the emergent trade-offs, such as artifacts induced at CRM grid transition boundaries, when Multi-Domain MMF is exploited for a binary mixture of low- and high-resolution CRM domains. Comparison of Multi-Domain results against globally homogenous standard MMF-LR and (expensive) MMF-HR simulations shows that it produces remarkably similar low cloud fractions and shortwave radiative fluxes as the standard MMF-HR configuration over the heavily loaded sub-region, as intended, while the rest of the globe stays similar to the MMF-LR baseline. Although artifacts can be induced such as by positioning the meridional boundary of a CRM grid transition on the equator in an otherwise

symmetric aquaplanet, in real-geography tests that use sensible sub-regions to contain the high resolution, unintended consequences are not obvious.

Thus, in the special case of MMF simulations, a range of immediate applications of this infrastructure can be envisioned. For instance we can:

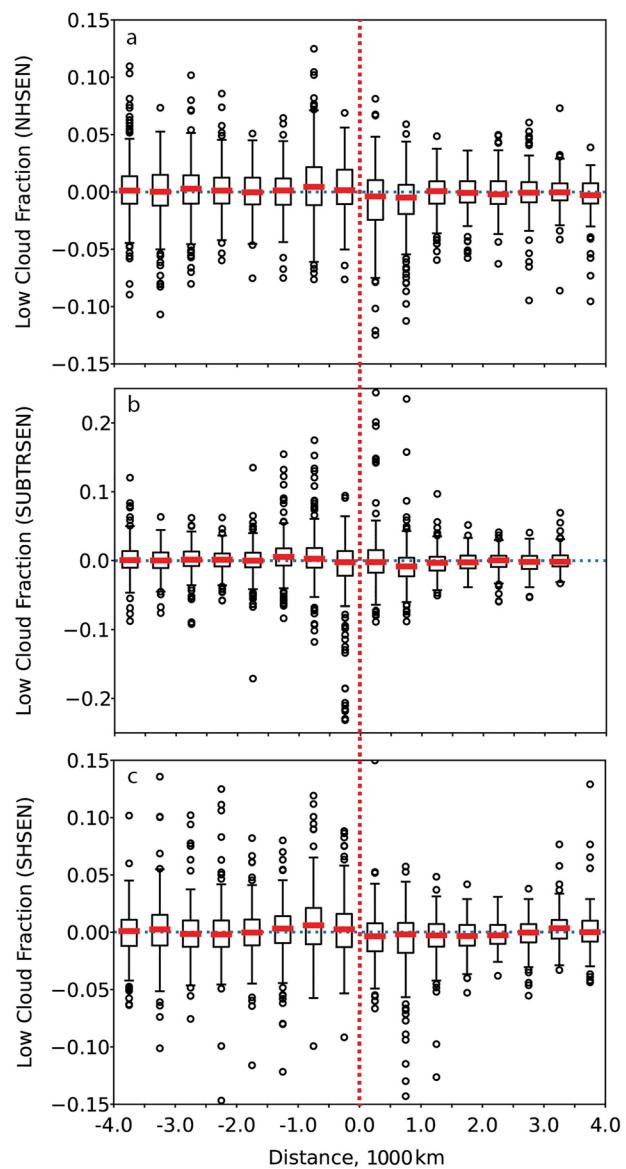
1. (our own interest) increasingly avoid approximation of shallow cloud systems by embedding increasingly high-resolution and even 3D CRMs in strategic locations to study their broader climate system interactions at maximum throughput.
2. Increase the throughput and reduce the cost of classical 2D LR by avoiding its use over areas such as the extratropical ocean where it appears to have marginal benefits (Kooperman et al., 2016a).
3. Afford the expense of large 3D CRMs in classical SP by regionalizing them just to the tropics, using standard 2D SP or even conventional parameterization elsewhere; this could assist tuning of coupled LR simulations by admitting explicit shallow cumulus momentum transport that acts as a throttle on the equatorial cold tongue (Woelfle et al., 2018), instead of approximating momentum effects via parameterized scalar transport (Tulich, 2015).
4. Deploy especially large, high-fidelity storm-resolving 3D CRM domains only along corridors that connect such events' genesis regions to their main human impacts sectors, such as where vulnerable societies or valuable infrastructure exist.

Incremental expansion of this same technique could lead to a Multi-CRM approach that includes broader diversity than a binary mixture of CRM grids, or that allows a more gradual transition between differing CRM resolutions or domain setups. With more work, instead of a static horizontal grid transition boundary, the approach could also be extended to a dynamically adaptable one.

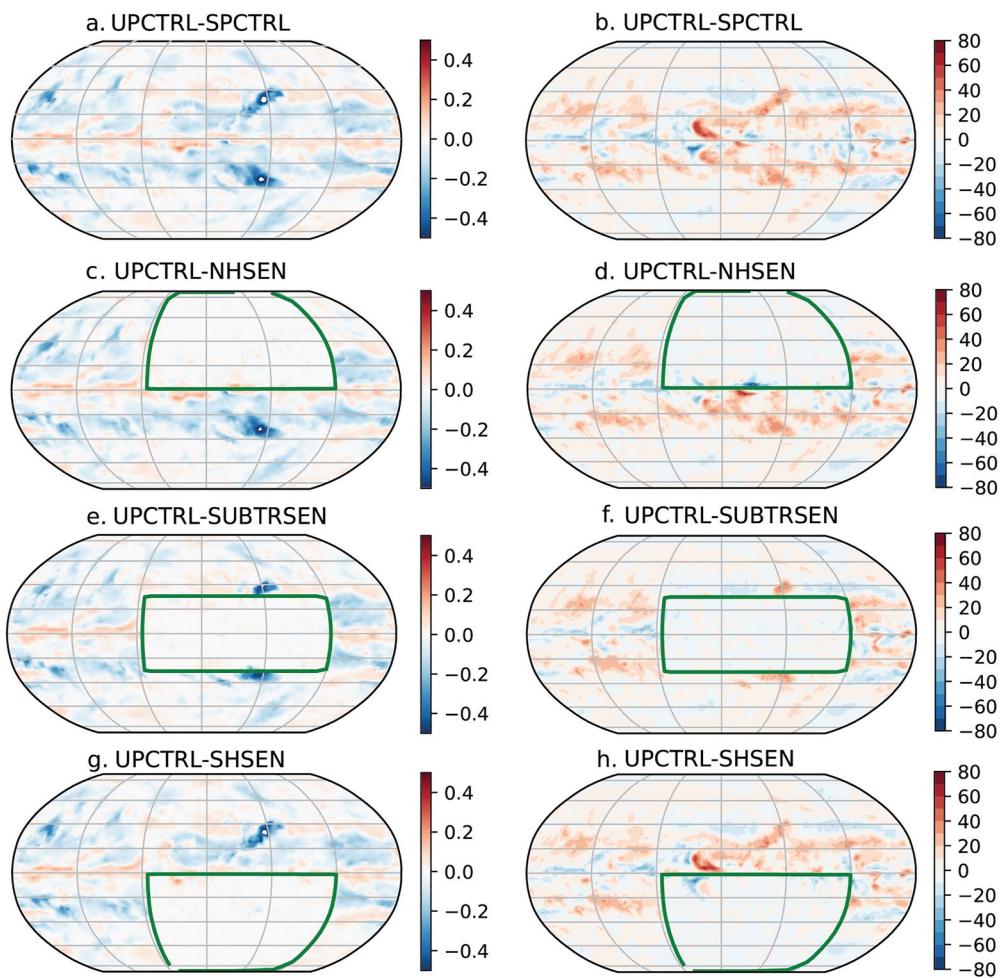
More generally, the same load balancing infrastructure that already enables our Multi-Domain CRM approach should immediately allow any interested E3SM or CESM developer to focus any form of highly intense, rate-limiting atmospheric physics calculations over small regions of the planet where this could be helpful to science. We hope our study will enable new experimentation for other classes of GCM physics where pushing high performance computing limits amidst geographically heterogeneous calculation could be advantageous.

## Appendix A: Additional Metric Comparing Multi-Domain CRM With LRCTRL and HRCTRL

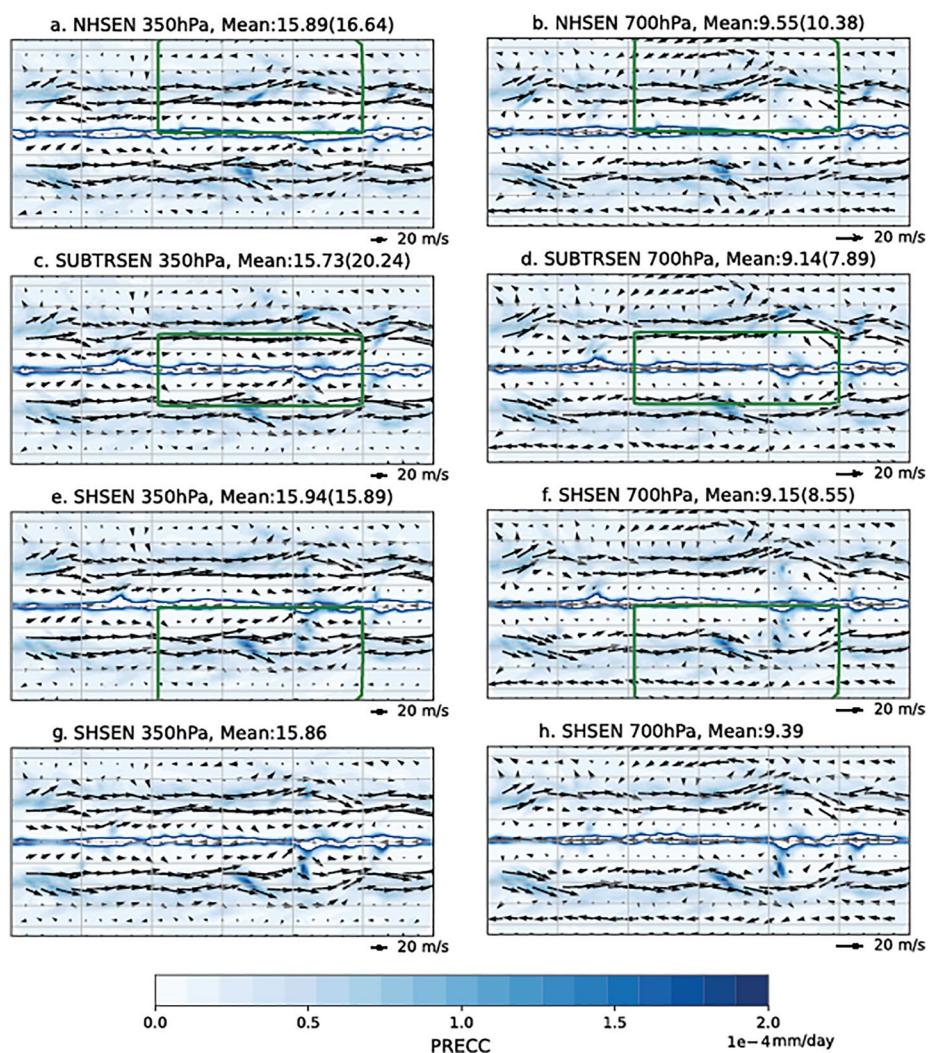
To compare variable variations across the HR/LR mask area for aquaplanet simulations, we define the shortest distance of each grid point relative to the HR mask boundary as a new coordinate, that is, grid points with a smaller distance are closer to the HR/LR boundary. Inside (outside) the HR mask region, grid point has a positive (negative) distance. We further distribute all grid points across 16 bins span from  $-4,000$  to  $4,000$  km with a 500 m interval. The box plot (Figure A1) shows the median, lower quartile, upper quartile, and the outliers of grid points inside each bin. The y axis of the box plot represents the difference between HRCTRL and Multi-Domain CRM for a positive distance and LRCTRL and Multi-Domain CRM for a negative distance. This comparison further confirms that HR/LR boundary has limited impact on the cloud fraction.



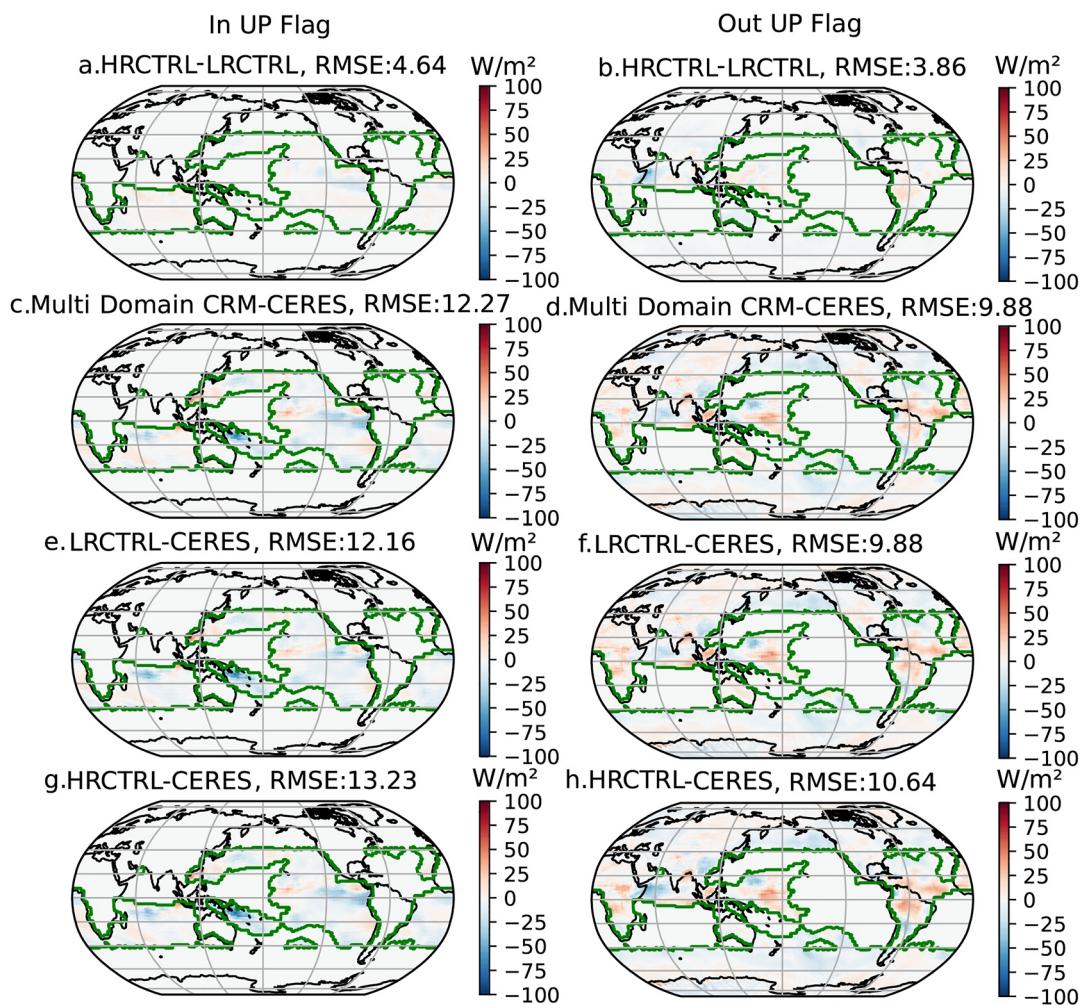
**Figure A1.** Box plot of the cloud fraction differences based on (a) Northern hemisphere, (b) Subtropics, and (c) Southern Hemisphere. Columns with positive (negative) distance represent in (out of) the high-resolution mask region. The cloud fraction difference is between high-resolution control and Multi-Domain cloud resolving model (CRM) for positive distances and low-resolution control and Multi-Domain CRM for negative distances. The red dash line represents the equator.



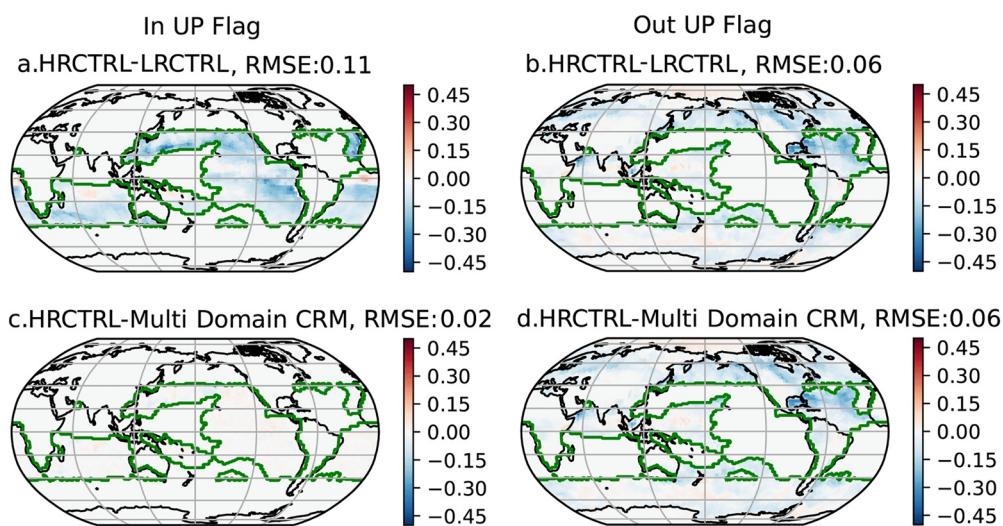
**Figure A2.** The first 7-day ensemble mean low cloud fraction difference (left) between (a) high-resolution control (HRCTRL) and low-resolution control (LRCTRL), (c) HRCTRL and Northern hemisphere (NHSEN), (e) HRCTRL and Subtropics (SUBTRSEN), (g) HRCTRL and Southern Hemisphere (SHSEN). The absorbed shortwave radiation (right) difference between (a) HRCTRL and LRCTRL, (c) HRCTRL and NHSEN, (e) HRCTRL and SUBTRSEN, (g) HRCTRL and SHSEN. The high-resolution mask region is inside the area encompassed by the green line.



**Figure A3.** The 20-day averaged horizontal wind vectors and convective precipitation rate (shaded PRECC,  $\text{m/s}$ ) at 350 hPa (left) and 700 hPa (right) for (a and b) Northern hemisphere, (c and d) Subtropics, (e and f) Southern Hemisphere, and (g and h) high-resolution control. The high-resolution (HR) mask region is inside the area encompassed by the green line, and average horizontal wind speeds are shown in the title and the number within brackets represents the average within the HR mask region.



**Figure A4.** The map of the ensemble mean outgoing longwave radiation differences between high-resolution control (HRCTRL) and low-resolution control (LRCTRL) for (a) inside (b) outside the high-resolution (HR) mask area, between Multi-Domain cloud resolving model (CRM) and Clouds and the Earth's Radiant Energy System (CERES) observations for (c) inside (d) outside the HR mask area, between LRCTRL and CERES observations for (e) inside (f) outside the HR mask area, and between HRCTRL and CERES for (g) inside (h) outside the HR mask area. The HR mask area is encompassed by the green solid line.



**Figure A5.** The map of the ensemble mean low cloud fraction differences between high-resolution control (HRCTRL) and low-resolution control (LRCTRL) for (a) inside (b) outside the high-resolution (HR) mask area, and between HRCTRL and Multi-Domain cloud resolving model (CRM) for (c) inside (d) outside the HR mask area.

## Data Availability Statement

All E3SM source code may be accessed on the GitHub repository (<https://github.com/E3SM-Project/E3SM>), and instructions are available at the website (<https://e3sm.org/model/running-e3sm/e3sm-quick-start>). All code modifications needed to implement our approach within a legacy fork of the E3SM Multiscale Modeling Framework climate model are available at <https://github.com/mspritch/E3SM/commits/pritch/multiCRM-openmp4.5>, in which we also include a flexible infrastructure that allows the user to provide an auxiliary input file specifying the geographic locations of intense work at run time. The raw output data is archived at <https://doi.org/10.5281/zenodo.5348380> (doi: 10.5281/zenodo.5348380).

## Acknowledgments

L. Peng, M. Pritchard, P. N. Blossey, and C. S. Bretherton acknowledge funding from the National Science Foundation (NSF) Climate and Large-scale Dynamics program under grants AGS-1912134 and AGS-1912130; W. M. Hannah, L. Peng, and M. Pritchard acknowledge the work was also performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This research used computational resources from the Texas Advanced Computing Center at The University of Texas at Austin under allocation ATM-20009 and from the Extreme Science and Engineering Discovery Environment, which is supported by NSF Grant ACI-1548562 (Towns et al., 2014), under allocation TG-ATM190002. We gratefully acknowledge many useful conversations with Xingqiu Yuan, Azamat M. Mametjanov, Tom Beucler, Andrea Jenney, Xiaojuan Liu, Jacob Michael Hendrickson, Griffin Mooers, and Jerry Lin and anonymous reviewers for their constructive comments and suggestions.

## References

Bader, D., Collins, W., Jacob, R., Jones, P., Rasch, P., Taylor, M., & Williams, D. (2014). Accelerated climate modeling for energy Tech. Rep. US department of energy.

Buccolieri, R., Salim, S. M., Leo, L. S., Di Sabatino, S., Chan, A., Ielpo, P., et al. (2011). Analysis of local scale tree-atmosphere interaction on pollutant concentration in idealized street canyons and application to a real urban junction. *Atmospheric Environment*, 45(9), 1702–1713. <https://doi.org/10.1016/j.atmosenv.2010.12.058>

Dennis, J. M., Vertenstein, M., Worley, P. H., Mirin, A. A., Craig, A. P., Jacob, R., & Mickelson, S. A. (2012). Computational performance of ultra-high-resolution capability in the Community Earth System Model. *International Journal of High Performance Computing Applications*, 26(1), 5–16. <https://doi.org/10.1177/1094342012436965>

Grabowski, W. W. (2004). An improved framework for superparameterization. *Journal of the Atmospheric Sciences*, 61(15), 1940–1952. [https://doi.org/10.1175/1520-0469\(2004\)061<1940:aiiffs>2.0.co;2](https://doi.org/10.1175/1520-0469(2004)061<1940:aiiffs>2.0.co;2)

Hagos, S., Leung, R., Rauscher, S. A., & Ringer, T. (2013). Error characteristics of two grid refinement approaches in aquaplanet simulations: MPAS-A and WRF. *Monthly Weather Review*, 141(9), 3022–3036. <https://doi.org/10.1175/mwr-d-12-00338.1>

Hannah, W. M., Bradley, A. M., Guba, O., Tang, Q., Golaz, J.-C., & Wolfe, J. (2021). Separating physics and dynamics grids for improved computational efficiency in spectral element Earth System Models. *Journal of Advances in Modeling Earth Systems*, 13(7), e2020MS002419. <https://doi.org/10.1029/2020MS002419>

Hannah, W. M., Jones, C. R., Hillman, B. R., Norman, M. R., Bader, D. C., Taylor, M. A., et al. (2020). Initial results from the super-parameterized E3SM. *Journal of Advances in Modeling Earth Systems*, 12(1), e2019MS001863. <https://doi.org/10.1029/2019ms001863>

Hurrell, J. W., Holland, M. M., Gent, P. R., Ghan, S., Kay, J. E., Kushner, P. J., et al. (2013). The Community Earth System Model: A framework for collaborative research. *Bulletin of the American Meteorological Society*, 94(9), 1339–1360. <https://doi.org/10.1175/bams-d-12-00121.1>

Hwang, Y.-T., & Frierson, D. M. (2010). Increasing atmospheric poleward energy transport with global warming. *Geophysical Research Letters*, 37(24), L24807. <https://doi.org/10.1029/2010gl045440>

Jansson, F., van den Oord, G., Pelupessy, I., Grönqvist, J. H., Siebesma, A. P., & Crommelin, D. (2019). Regional superparameterization in a global circulation model using large eddy simulations. *Journal of Advances in Modeling Earth Systems*, 11(9), 2958–2979. <https://doi.org/10.1029/2018ms001600>

Khairoutdinov, M., Randall, D., & DeMott, C. (2005). Simulations of the atmospheric general circulation using a cloud-resolving model as a superparameterization of physical processes. *Journal of the Atmospheric Sciences*, 62(7), 2136–2154. <https://doi.org/10.1175/jas3453.1>

Khairoutdinov, M. F., & Randall, D. A. (2001). A cloud resolving model as a cloud parameterization in the NCAR Community Climate System Model: Preliminary results. *Geophysical Research Letters*, 28(18), 3617–3620. <https://doi.org/10.1029/2001gl013552>

Kooperman, G. J., Pritchard, M. S., Burt, M. A., Branson, M. D., & Randall, D. A. (2016a). Impacts of cloud superparameterization on projected daily rainfall intensity climate changes in multiple versions of the Community Earth system Model. *Journal of Advances in Modeling Earth Systems*, 8(4), 1727–1750. <https://doi.org/10.1002/2016ms000715>

Kooperman, G. J., Pritchard, M. S., Burt, M. A., Branson, M. D., & Randall, D. A. (2016b). Robust effects of cloud superparameterization on simulated daily rainfall intensity statistics across multiple versions of the Community Earth System Model. *Journal of Advances in Modeling Earth Systems*, 8(1), 140–165. <https://doi.org/10.1002/2015ms000574>

Mirin, A. A., & Worley, P. H. (2012). Improving the performance scalability of the community atmosphere model. *International Journal of High Performance Computing Applications*, 26(1), 17–30. <https://doi.org/10.1177/1094342011412630>

Neale, R. B., & Hoskins, B. J. (2000). A standard test for AGCMs including their physical parametrizations: I: The proposal. *Atmospheric Science Letters*, 1(2), 101–107. <https://doi.org/10.1006/asle.2000.0019>

Parishani, H., Pritchard, M. S., Bretherton, C. S., Terai, C. R., Wyant, M. C., Khairoutdinov, M., & Singh, B. (2018). Insensitivity of the cloud response to surface warming under radical changes to boundary layer turbulence and cloud microphysics: Results from the ultraparameterized CAM. *Journal of Advances in Modeling Earth Systems*, 10(12), 3139–3158. <https://doi.org/10.1029/2018ms001409>

Parishani, H., Pritchard, M. S., Bretherton, C. S., Wyant, M. C., & Khairoutdinov, M. (2017). Toward low-cloud-permitting cloud superparameterization with explicit boundary layer turbulence. *Journal of Advances in Modeling Earth Systems*, 9(3), 1542–1571. <https://doi.org/10.1002/2017ms000968>

Pickering, K. E., Thompson, A. M., Scala, J. R., Tao, W.-K., Dickerson, R. R., & Simpson, J. (1992). Free tropospheric ozone production following entrainment of urban plumes into deep convection. *Journal of Geophysical Research*, 97(D16), 17985–18000. <https://doi.org/10.1029/92jd01716>

Prather, M. J., & Jacob, D. J. (1997). A persistent imbalance in HO<sub>x</sub> and NO<sub>x</sub> photochemistry of the upper troposphere driven by deep tropical convection. *Geophysical Research Letters*, 24(24), 3189–3192. <https://doi.org/10.1029/97gl03027>

Rauscher, S. A., Ringler, T. D., Skamarock, W. C., & Mirin, A. A. (2013). Exploring a global multiresolution modeling approach using aquaplanet simulations. *Journal of Climate*, 26(8), 2432–2452. <https://doi.org/10.1175/jcli-d-12-00154.1>

Reynolds, R. W., Smith, T. M., Liu, C., Chelton, D. B., Casey, K. S., & Schlax, M. G. (2007). Daily high-resolution-blended analyses for sea surface temperature. *Journal of Climate*, 20(22), 5473–5496. <https://doi.org/10.1175/2007jcli1824.1>

Schneider, T., Lan, S., Stuart, A., & Teixeira, J. (2017). Earth system modeling 2.0: A blueprint for models that learn from observations and targeted high-resolution simulations. *Geophysical Research Letters*, 44(24), 12–396. <https://doi.org/10.1002/2017gl076101>

Smolarkiewicz, P. K., & Grabowski, W. W. (1990). The multidimensional positive definite advection transport algorithm: Nonoscillatory option. *Journal of Computational Physics*, 86(2), 355–375. [https://doi.org/10.1016/0021-9991\(90\)90105-a](https://doi.org/10.1016/0021-9991(90)90105-a)

Terai, C. R., Pritchard, M. S., Blossey, P., & Bretherton, C. S. (2020). The impact of resolving subkilometer processes on Aerosol-Cloud interactions of Low-Level clouds in global model simulations. *Journal of Advances in Modeling Earth Systems*, 12(11), e2020MS002274. <https://doi.org/10.1029/2020ms002274>

Towns, J., Cockerill, T., Dahan, M., Foster, I., Gaither, K., Grimshaw, A., et al. (2014). XSEDE: Accelerating scientific discovery. *Computing in science & engineering*, 16(5), 62–74. <https://doi.org/10.1109/MCSE.2014.80>

Tulich, S. (2015). A strategy for representing the effects of convective momentum transport in multiscale models: Evaluation using a new superparameterized version of the Weather Research and Forecast model (SP-WRF). *Journal of Advances in Modeling Earth Systems*, 7(2), 938–962. <https://doi.org/10.1002/2014ms000417>

Wielicki, B. A., Barkstrom, B. R., Harrison, E. F., Lee, R. B., III, Smith, G. L., & Cooper, J. E. (1996). Clouds and the Earth's Radiant Energy System (CERES): An Earth observing system experiment. *Bulletin of the American Meteorological Society*, 77(5), 853–868. [https://doi.org/10.1175/1520-0477\(1996\)077<0853:catere>2.0.co;2](https://doi.org/10.1175/1520-0477(1996)077<0853:catere>2.0.co;2)

Woelfle, M., Yu, S., Bretherton, C., & Pritchard, M. (2018). Sensitivity of coupled tropical Pacific model biases to convective parameterization in CESM1. *Journal of Advances in Modeling Earth Systems*, 10(1), 126–144. <https://doi.org/10.1029/2017ms001176>

Worley, P. H. (2006). Benchmarking using the community atmospheric model. In *Proceedings of the 2006 spec benchmark workshop*.

Worley, P. H., & Drake, J. B. (2005). Performance portability in the physical parameterizations of the community atmospheric model. *International Journal of High Performance Computing Applications*, 19(3), 187–201. <https://doi.org/10.1177/1094342005056095>

Worley, P. H., Vertenstein, M., & Craig, A. P. (2011). Community climate system model. In D. Padua (ed.), *Encyclopedia of parallel computing* (pp. 342–351). Springer US. [https://doi.org/10.1007/978-0-387-09766-4\\_376](https://doi.org/10.1007/978-0-387-09766-4_376)

Wyant, M. C., Bretherton, C. S., & Blossey, P. N. (2009). Subtropical low cloud response to a warmer climate in a superparameterized climate model. Part I: Regime sorting and physical mechanisms. *Journal of Advances in Modeling Earth Systems*, 1(3), 7. <https://doi.org/10.3894/james.2009.1.7>