

# 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 S. 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>University of California, Irvine

<sup>2</sup>Lawrence Livermore National Lab

<sup>3</sup>University of Washington

<sup>4</sup>PHWorley Consulting

November 26, 2022

## 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 CESM and 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 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) HR grid columns. Two examples prove the concept, showing that minor artifacts can be introduced near the HR/LR 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.

1       **Load-balancing intense physics calculations to embed**  
2       **regionalized high-resolution cloud resolving models in**  
3       **the E3SM and CESM climate models**

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

6               <sup>1</sup>Department of Earth System Science, University of California, Irvine, California, USA

7               <sup>2</sup>Lawrence Livermore National Laboratory, Livermore, California, USA

8               <sup>3</sup>Department of Atmospheric Sciences, University of Washington, Seattle, Washington, USA

9               <sup>4</sup>PHWorley Consulting, Oak Ridge, Tennessee, USA

10       **Key Points:**

- 11       • We adapt E3SM/CESM so most of a processor pool can operate on just a sub-  
12       set of demanding physics columns.
- 13       • Load balancing theory finds optimal computational scale for regionalized high in-  
14       tensity physics work.
- 15       • Multi-Domain CRM tests with LES resolution embedded in 30% of the superpa-  
16       rameterized E3SM succeed with few artifacts.

---

Corresponding author: Michael Pritchard, [mspritch@uci.edu](mailto:mspritch@uci.edu)

**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 CESM and 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 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) HR grid columns. Two examples prove the concept, showing that minor artifacts can be introduced near the HR/LR 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 vs. 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 and stratocumuli require 10-meter horizontal resolution to resolve faithfully. Handling all of these scales on a uniformly high-resolution (LES) 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

67 times to nest smoothly from horizontal grid of 25- to 100-km in an exterior planetary  
68 atmospheric model to the 250-m needed to quasi-resolve boundary layer dynamics, and  
69 experience attests that considerable computational expense must be expended in the tran-  
70 sition zones between the nested meshes (Gustafson Jr et al., 2020). Meanwhile, phys-  
71 ical trouble can also be expected in the lateral nesting transition regions given that con-  
72 ventional physics parameterizations have strong “grey-zone” sensitivities (Hagos et al.,  
73 2013) to exterior resolution. In this context it would be useful if individual users of Multi-  
74 scale climate Modeling Frameworks (MMFs; Grabowski (2004); M. F. Khairoutdinov &  
75 Randall (2001); Hannah et al. (2020)) could focus embedded resolution where they think  
76 it matters for specific problems of interest.

77 Unfortunately, even MMF simulations, which exhibit gentler sensitivities to external  
78 resolution (Kooperman et al., 2016a) and sidestep the typical resolution-nesting prob-  
79 lems, cannot be reconfigured to vary regionally due to simple load-balancing issues as-  
80 sociated with geographic workload variance that have yet to be fully solved. A symp-  
81 tom of this problem is that, with one exception (Jansson et al., 2019), all MMF tests  
82 have deployed cloud superparameterization (SP) *globally* despite evidence that classi-  
83 cal SP mostly improves the simulated rainfall distribution near the equator and over sum-  
84 mer continents (Kooperman et al., 2016a,b), and despite the fact that even more costly  
85 forms of refined-resolution SP such as “ultra-parameterization (UP)” were designed to  
86 improve the simulation of convection over small regions of the planet where marine low  
87 clouds tend to occur. One might naturally wonder then why SP and HR have not been  
88 deployed *regionally*, at reduced computational expense, for instance to unburden the cost  
89 of long, ocean-coupled simulations, or to open up computational room to afford a relax-  
90 ation of the MMF’s historical idealizations (e.g. 2D, small-domain, coarse-resolution CRMs).  
91 One key issue is that the limits of the *load balancing* software infrastructure inherited  
92 from the host CESM and E3SM climate models have prevented experimentation of this  
93 variety.

94 Load balancing is not a new issue. Global climate models have many levels of par-  
95 allelism that can be exploited. Constituent component models (atmosphere, ocean, land,  
96 etc.) can be run in parallel, i.e. on distinct subsets of processing elements, with some  
97 or all of the other components, and load balancing focuses on how computing resources  
98 are allocated between the components (Worley et al., 2011; Dennis et al., 2012). For in-  
99 dividual component models, parallelism is typically introduced first via a domain decom-  
100 position of the horizontal computational grid. For the numerical methods used in the  
101 E3SM and CESM for dynamics and tracer advection, compact geographical patches are  
102 most efficient. For the parameterized physics in the CESM and E3SM atmosphere mod-  
103 els – which dominate the overall expense and scalability of MMFs (M. Khairoutdinov  
104 et al., 2005) owing to the superparameterization approach of embedding expensive lo-  
105 cal cloud-resolving calculations – computation of vertical columns (fixed horizontal co-  
106 ordinates) are independent, and subsets of columns can be grouped and assigned to pro-  
107 cessing units, both to enhance vectorization and to improve load balance (Worley & Drake,  
108 2005). This group need not be geographically contiguous (Worley, 2006): indeed an ad-  
109 jacent geographic assignment of columns into processing elements is generally not de-  
110 sirable as it introduces, for example, load imbalance due to vertical columns correspond-  
111 ing to daytime requiring more intensive (shortwave plus longwave) radiation calculation  
112 than nighttime columns (longwave alone). Columns from the same geographic region will  
113 also more likely exhibit similar simulation-dependent physical processes, such as micro-  
114 physical calculations in locations frequently covered in cloud. Therefore, by default, the  
115 atmosphere models in the US climate models CESM and E3SM do not assign neighbor-  
116 ing columns to the same processing element. Instead, a space-filling-curve or a longitude-  
117 latitude global ordering of columns are “dealt” to the processing elements in a wrap map-

118 ping, mixing up columns geographically.<sup>1</sup> Such an approach attempts to load balance  
 119 by reducing the variability in computational cost between processing elements arising  
 120 from the diurnal and seasonal cycles, columns over land versus ocean, columns over the  
 121 arctic versus the equator, etc., in order to obtain improved performance.<sup>2</sup> This wrap map  
 122 approach to load balancing assigns approximately the same number of columns to each  
 123 processing element. If the average cost per column assigned to these processing elements  
 124 (local average) is approximately the same as the average over all columns (global aver-  
 125 age) at a given instance in simulation time, then this mapping is near optimal. Whether  
 126 this condition on the local and global average costs is achieved is difficult to determine,  
 127 but, based on the known static sources of load imbalance, the CESM/E3SM approach  
 128 is reasonable. Another justification for the approach is that, for the typical physical pa-  
 129 rameterizations used in the atmosphere in CESM and E3SM, the computational cost per  
 130 processing element is most strongly a function of the number of columns assigned, and  
 131 assigning the individual columns in a load-balanced fashion subject to an equidistribu-  
 132 tion of the number of columns simply improves upon this initial optimization. In any  
 133 case, the CESM/E3SM approach has proven to be better than not load balancing at all.

134 However, the performance of this approach to load balancing the physical param-  
 135 eterizations will degrade if the difference in the computational cost between individual  
 136 columns is large. This becomes especially true as the number of processing elements ap-  
 137 proaches the number of physical grid columns. For instance, at extreme computational  
 138 scale, when each processing element has only a single grid column, all processing elements  
 139 will wait for the most expensive grid column to finish. Nonetheless, if the relative com-  
 140 putational cost per column is known, a load-balanced assignment of columns to process-  
 141 ing elements can instead be computed directly, not depending on the heuristics currently  
 142 used. This interesting use case naturally emerges in superparameterized climate simu-  
 143 lations in which different GCM grid columns are simulating different cloud regimes with  
 144 different characteristic eddy scales.

145 The current load balancing heuristics likewise may be inappropriate for next-generation  
 146 climate simulation for applications beyond global cloud feedback, where innovations in  
 147 regionalized, high-intensity sub-grid physics could also be helpful. For instance, a hu-  
 148 man impact modeler may benefit from embedding regionally intensive atmospheric physics  
 149 and embedded boundary layer calculations only along corridors that produce damaging  
 150 extreme events to vulnerable populations. Those interested in climate feedbacks linked  
 151 to long-range transport of biogeochemical feedbacks might want to embed high resolu-  
 152 tion near sources of emissions, perhaps including effects of urban building configuration,  
 153 road traffic, or vegetation nearby megacities (Buccolieri et al., 2011), in addition to ex-  
 154 plicit treatments of storm dynamics responsible for scavenging along likely transport path-  
 155 ways. Similarly, an atmospheric chemist may wish to hyper-resolve chemical reactions

---

<sup>1</sup> There is also an option to assign columns in pairs when doing the wrap mapping, where the pairs dif-  
 fer 180 degrees 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.

<sup>2</sup> Note that using different domain decompositions for the dynamics and for the physical parameteri-  
 zations incurs communication cost (Mirin & Worley, 2012), whether via explicit MPI communication or  
 memory copying or both. There are also options in CESM and E3SM to retain the geographic decompo-  
 sition 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 simula-  
 tions 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.

156 and processing of reagents by explicit turbulence just within subregions of the tropics  
 157 where nitrous oxide, odd hydrogen, and their precursors from deep convection or strato-  
 158 spheric intrusions, are vital to ozone prediction (Prather & Jacob, 1997; Pickering et al.,  
 159 1992). A tropical ecohydrologist may want to hyper-resolve the atmospheric boundary  
 160 layer only over tropical rainforests to look at surface flux exchanges and land-atmosphere  
 161 impacts and their climate teleconnections with minimal approximation. In each of the  
 162 above examples, a large computational gap exists between a small pool of *heavily-loaded*  
 163 grid columns on the physics side of a climate prediction code relative to a much larger  
 164 pool of relatively *lightly-loaded* grid columns.

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

171 The paper is structured as follows. In Section 2 we begin to build some theory and  
 172 then demonstrate via a concrete example why the existing physics column load-balancing  
 173 infrastructure in the CESM and E3SM can be inefficient when presented with highly re-  
 174 gionalized, high-intensity physics computations. In Section 2.3, we present a general the-  
 175 ory for how to load balance a binary mixture of high-intensity and low-intensity physics  
 176 work, mainly by relaxing an assumption that a fixed number of columns be used in the  
 177 processing elements that are used to parallel-decompose groups of physics columns. The  
 178 theory predicts an optimal computational scale for a generic range of problems and its  
 179 predictions are tested against actual performance measurements of the E3SM climate  
 180 model.

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

## 191 2 Methods

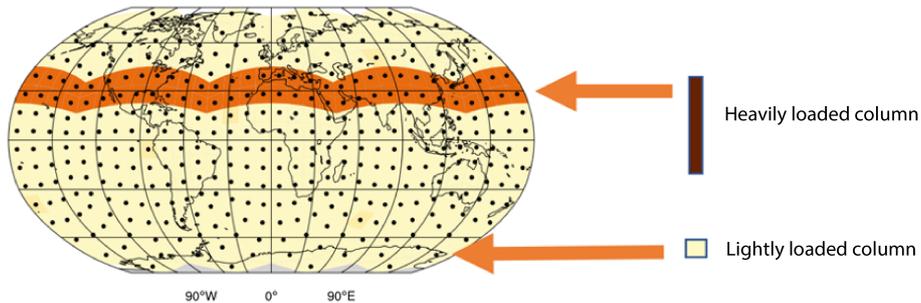
192 Let us begin by defining some parameters, considering the example of a binary mix-  
 193 ture of “lightly” ( $l$ )- vs. “heavily” ( $h$ )-loaded columns. For a given GCM horizontal res-  
 194 olution, the total number of physical columns  $C$  is fixed, and we define the number frac-  
 195 tion  $F$  of the world covered by the heavy columns as

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

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

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

198 Let us view  $F$  as a flexible parameter that different climate modelers may wish to  
 199 vary according to different applications. For simplicity we assume a geographically static  
 200 load imbalance, i.e. each physical column can be categorized as either a heavily ( $h$ ) or



**Figure 1.** Physical columns (black dots) for a coarse resolution aquaplanet with a binary mixture of heavy vs. 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).

201 lightly ( $l$ ) loaded column based on a pre-determined binary map array read by the model  
 202 during the initialization stage (e.g. see Section 2.3): we will not change the location and  
 203 the number of heavily vs. lightly loaded columns during the simulation. All physical columns  
 204 are then distributed across  $P$  processing elements.

205 The next fundamental parameter is the computational intensity ratio,

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

206 i.e. the average compute time required for one processing element to integrate a single  
 207 heavily loaded physics column ( $t_h$ ) divided by the time required for it to integrate a single  
 208 lightly loaded column ( $t_l$ ).

## 209 2.1 A coarse-resolution example to inform optimization trade-offs

210 To illustrate the issues involved in load-balancing, we now investigate an unusu-  
 211 ally coarse resolution aquaplanet configuration<sup>3</sup> of the superparameterized E3SM (Han-  
 212 nah et al., 2020) that can be run at a convenient computational scale. In the coarse aqua-  
 213 planet there are only  $C = 384$  physical columns spanning the globe (black dots in Fig-  
 214 ure 1), but it is nonetheless representative of the load balancing challenge at higher com-  
 215 putational scales.

216 To create a binary mixture of heavily vs. lightly loaded physics columns, we added  
 217 the new capability to use two separate CRM grid configurations in separate physical grid  
 218 columns. Heavily loaded columns are achieved by either increasing the number of em-  
 219 bedded CRM columns or reducing the time step relative to the standard CRM grid con-  
 220 figuration. In this way, we tested the effect on parallel decomposition of adding 5 times  
 221 extra physics work ( $\tau = 5$ ) over a narrow latitude band ( $N_h = 64$  columns) shown by  
 222 the dark orange belt (Figure 1). That is, the fraction of the world covered in heavy work  
 223  $F = 64/384 \approx 0.16$ . Note the icons at right in Figure 1, which will be referred to in  
 224 later schematics – the dark red rectangles denote these heavy-loaded horizontal grid columns.

<sup>3</sup>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).

225 Small yellow squares with five times less vertical extent denote the lightly loaded columns  
 226 that require five times fewer calculations apiece. Our final assumption is that there are  
 227  $P = 128$  processing elements, i. e.  $n \equiv C/P = 3$  physics columns per processing ele-  
 228 ment.

229 The immediate problem with the E3SM’s existing load balancing infrastructure is  
 230 exposed by the schematic in Figure 2a, which summarizes how the binary mixture of heav-  
 231 ily vs. lightly loaded columns is assigned to each of the 128 processing elements by de-  
 232 fault. Since (by design in this example) the number of heavily loaded columns is not an  
 233 integral multiple of the number of processing elements, not all tasks can be assigned an  
 234 equal mixture of heavy vs. lightly loaded columns. Instead, the first 64 processing el-  
 235 ements are assigned a triad of one-heavy plus two-lightly loaded columns. Next, having  
 236 run out of heavily loaded columns to assign, the remaining 64 processing elements are  
 237 each assigned a triad comprising three lightly loaded columns, a comparatively smaller  
 238 workload. The implication is a discrete jump in computational cost per processing el-  
 239 ement, i.e. a load imbalance, confirmed by our own task-level timing measurements (Fig-  
 240 ure 2b). Practically, this means that the processing elements assigned light work will wait  
 241 and idle until the processing elements assigned heavier work are done computing.

242 One unsatisfying solution to this problem might be to increase  $F$  such that a larger  
 243 fraction of the planet is covered by heavy work until equal groups of work can be attained  
 244 for each processing element. For instance,  $F = \frac{1}{3}$  i.e.  $N_h = 128$ , or  $F = \frac{2}{3}$  i.e.  $N_h =$   
 245  $256$  in this example would immediately avoid imbalance. Since  $N_h$  is now an integral mul-  
 246 tiple of the total number of processing elements  $P$ , this is compatible with the default  
 247 load-balancing scheme. That is, these discrete values of  $F$  work because we have assumed  
 248  $\frac{C}{P} = \frac{1}{3}$  in this example. An equally unsatisfying solution would be to decrease the num-  
 249 ber of processors  $P$ ; while this could help remove the inefficiency it is incongruous with  
 250 the idea of adding heavy regional work that is compensated with *high* computational scale,  
 251 our goal.

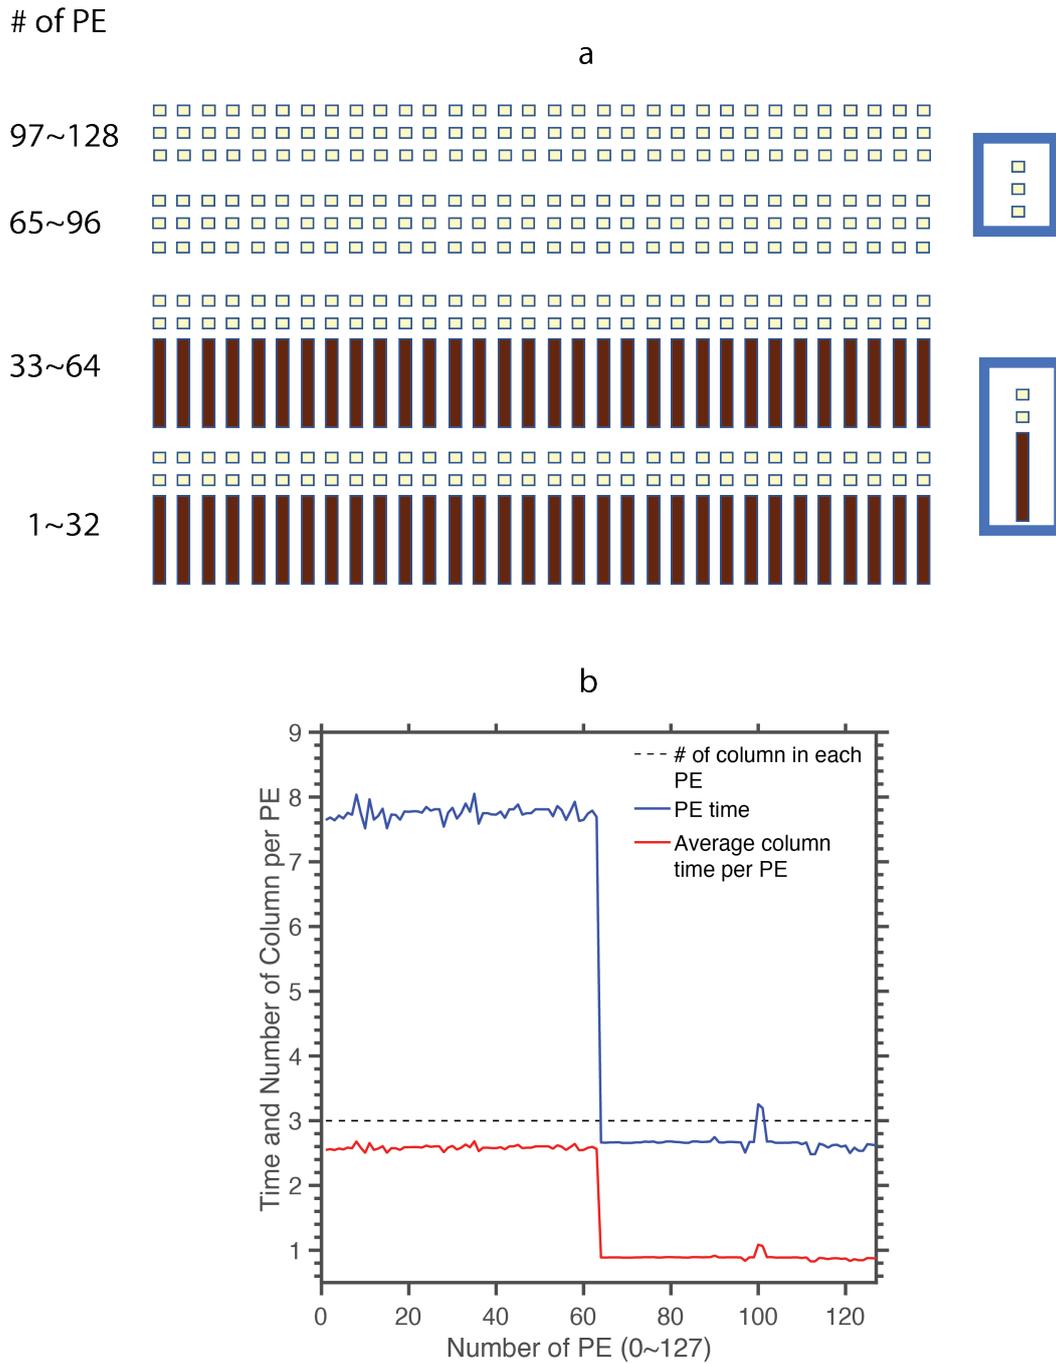
252 Our solution to this problem, shown in Figure 3, is to relax the assumption that  
 253 all processing elements have to work on equal-sized groups of physics columns. For the  
 254 current example, this means allowing the first 64 processing elements to devote them-  
 255 selves entirely to integrating a *single heavily loaded physics grid column* apiece. This dove-  
 256 tails with the idea that such physics columns, being rate-limiting, deserve maximal com-  
 257 putational resources. The remaining 64 processing elements each are then assigned equal  
 258 sized groups of five lightly loaded columns. This heavy:light ratio ( $\tau = 5$ ) of 5:1 is a  
 259 reflection of the assumed difference in computational intensity between the different re-  
 260 gions and can thus change with problem definition, but helps illustrate the issues at play  
 261 that must be considered. The new method can be adapted to other  $\tau$ . For instance, for  
 262 higher values of  $\tau$  but the same  $F$ , the model could still assign 64 heavily loaded columns  
 263 into the same first 64 processing elements but should increase the size of the groups of  
 264 lightly loaded columns to approximately  $\tau$ , requiring less overall processors.

265 For our own setup, as illustrated by the timing results in Figure 2b, the overall strat-  
 266 egy successfully achieves load balance for our chosen problem ( $F = 16.7\%$ ,  $M = C/3$ ,  $\tau =$   
 267  $5$ ), effectively allowing half of the processor pool to work on less than twenty percent of  
 268 the planet with maximum throughput and without introducing inefficiencies.

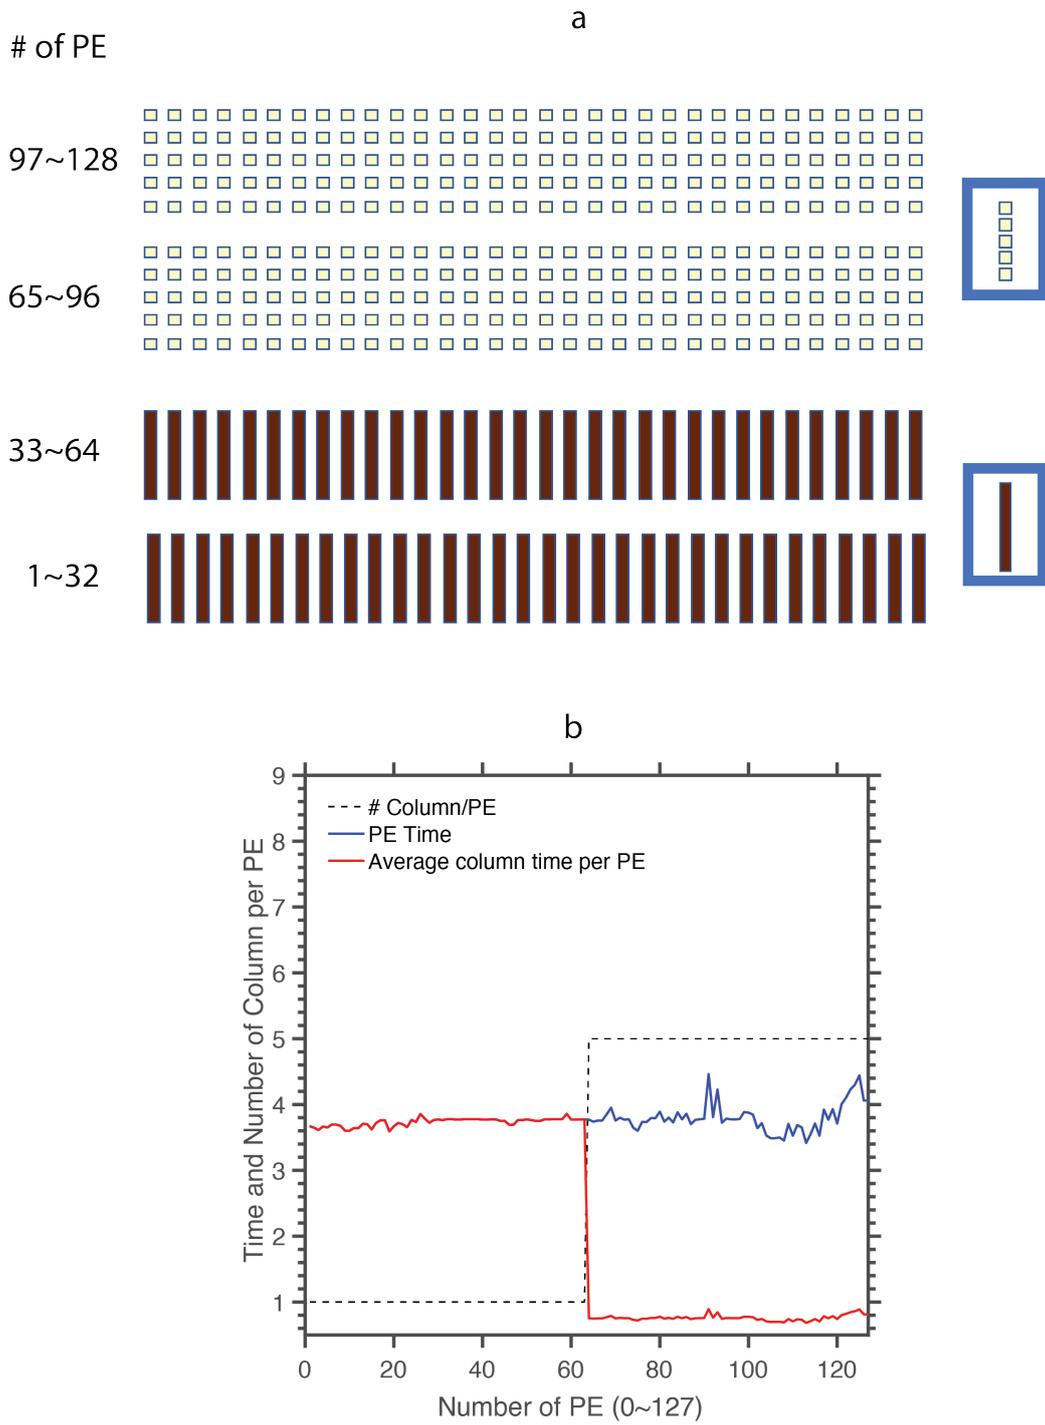
269 Informed by the qualitative lessons learned in this concrete example, we will now  
 270 derive some generalized constraints that allow a user to determine the full spectrum of  
 271 load balanced configurations ( $F, M, \tau$ ) that this approach opens HR.

## 272 2.2 Generalized theory of the problem

273 Under the current assumption in CESM and E3SM that all processing elements  
 274 must work on equal-size groups of physics columns, we have shown above that load-balanced



**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 (processing element 1 to 64) and two lightly loaded columns compared with three lightly loaded columns (processing element 65 to 128).



**Figure 3.** Similar to Figure 2 but using Multi-Domain CRM approach.

**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 E3SM.

$n=C/P$	$n_h$	$n_l$	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>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

275 conditions can be achieved for our binary mixture of heavy and light loading columns,  
 276 but only for a limited set of discrete conditions that will turn out to have unsatisfying  
 277 general properties. To see this, let

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

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

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

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

284 Table 1 summarizes the possible permutations of integers. In the simplest case,  $n =$   
 285 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  
 286 using half as many processors as there are physical grid cells, covering half the world in  
 287 heavy work is the only viable solution. For the next simplest case ( $n = 3$ ) two possi-  
 288 bilities exist -  $(n_h, n_l) = (1, 2)$  or  $(2, 1)$ . Thus if using one third as many processors as  
 289 there are physical grid cells,  $F = (\frac{1}{3}, \frac{2}{3})$  are both viable, as we found in the preceding  
 290 section. Likewise, for  $n = 4$ ,  $F = (\frac{1}{4}, \frac{1}{2}, \frac{3}{4})$  are viable and for  $n = 5$ ,  $F = (\frac{1}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5})$   
 291 work too, and so on.

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

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

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

299

### 2.3 Generalized solution

300  
301  
302  
303  
304  
305  
306

To solve this problem we separate the total processing elements into two pools of processors, i.e.  $P = P_h + P_l$ . We allow a subset of processors,  $P_h = N_h$  to be assigned individual<sup>4</sup> 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  $\tau$  (Eq.3),

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

307

Combining Eq.1, Eq.2, and Eq.7,  $P$  can be solved as

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

308  
309  
310  
311  
312  
313  
314  
315  
316  
317  
318  
319

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 processing elements ( $N_{HR} \approx P$ ) 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).

320  
321  
322  
323  
324  
325  
326  
327  
328

The more interesting and practical parameter dependence is observed for the regime  $2 \leq \tau \leq 100$  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.

329  
330  
331  
332  
333  
334

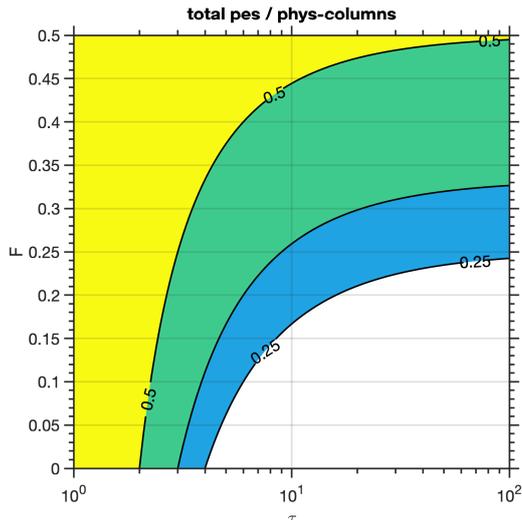
As a test of the theory, we now measure model throughput at computational scales in the vicinity of the predicted optimal  $P$  using two actual configurations of the E3SM climate model adapted to include our "Multi-Domain CRM" capability – an extremely course resolution aquaplanet (ne4pg2,  $C = 384$ , Hannah et al. (2021)) and an approximately 2.8-degree real-geography simulation (ne16pg2 horizontal resolution,  $C = 6144$ ) global grid (Figure 4).

335  
336  
337  
338

As in the previous example, workload imbalance is imposed using different, i.e. "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 Eq. 8 are shown by the vertical black dashed line in Fig-

---

<sup>4</sup>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



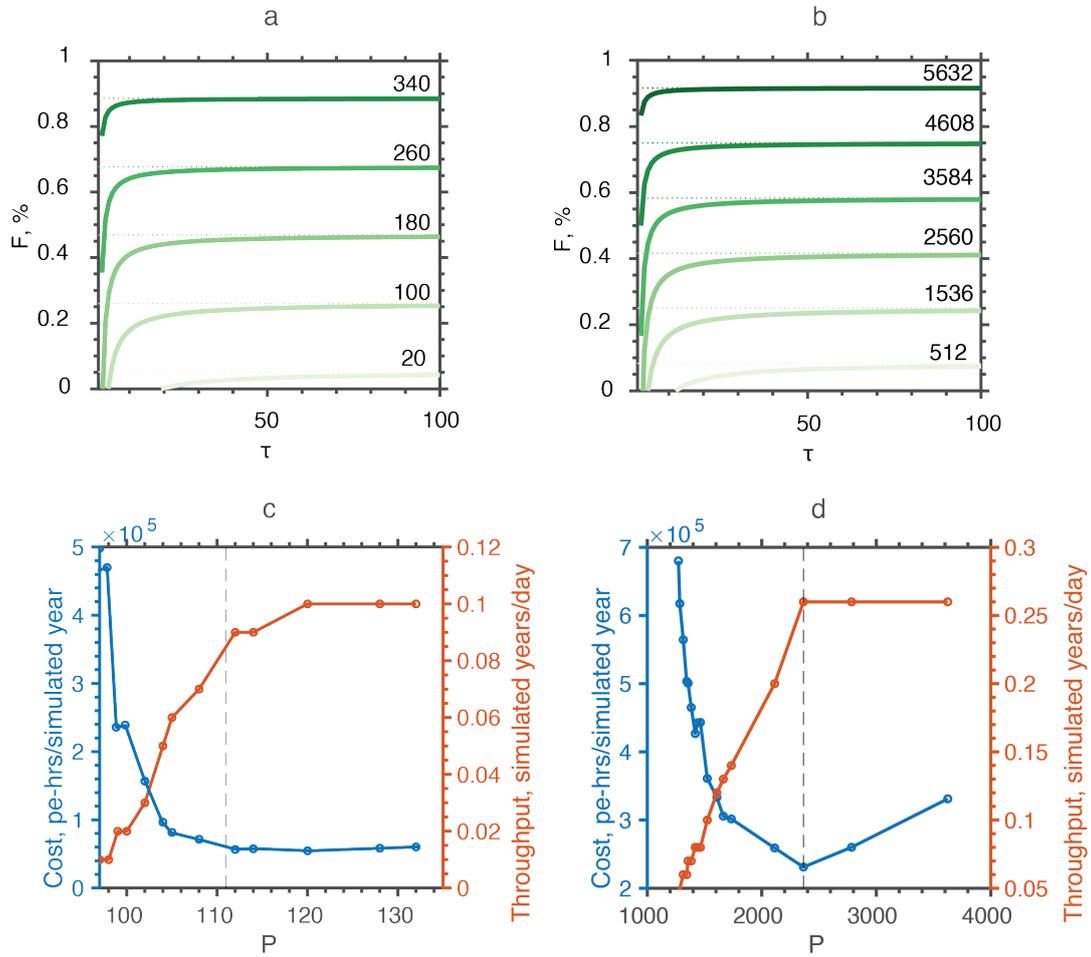
**Figure 4.** The most efficient total number of processing elements, expressed as a fraction of the total number of grid columns ( $\frac{P}{C}$ ) contoured as a function of the fraction of the world covered by heavily loaded columns ( $F$ ) and the heavy:light workload ratio ( $\tau$ ). Calculation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

339 ure 5cd, which summarizes the measured performance statistics across a range of adja-  
 340 cent choices of  $P$ . Confirming the theory, the scalability results of both experiments show  
 341 the predicted value indeed corresponds to the lowest model cost and approximately the  
 342 highest model throughput. Using a larger than optimal  $P$  spreads the pool of light work-  
 343 ing columns across more processing elements, but the overall throughput is unchanged,  
 344 since the rate is limited by the pool of  $P_h$  cores, thus needlessly increasing the total simu-  
 345 lation cost in units of processing element-hours or equivalent electrical energy burden.

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

### 355 3 Results

356 Whether it actually makes sense to use a binary mixture of CRM grids in opera-  
 357 tional climate simulation depends on how severe the artifacts induced at the grid tran-  
 358 sition boundary are. Such grid transition artifacts have been endemic to variable reso-  
 359 lution models with localized mesh refinement and limited domain CRM simulations that  
 360 exhibit a resolution dependence of precipitation and wind fields (Hagos et al., 2013; Rauscher  
 361 et al., 2013). We therefore perform a series of experiments and analyze similar trade-offs  
 362 of the Multi-Domain CRM approach, beginning with an aquaplanet to maximize detectabil-  
 363 ity of artifacts relative to a statistically steady background state.



**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 Eq. 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 Eq. 8.

364

### 3.1 Aquaplanet experiments results

365

366

367

368

369

370

371

372

373

374

375

376

377

378

379

380

381

382

383

384

385

386

387

388

We performed five integrations using a ne16pg2 (total 6144 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 LRCTRL a coarse CRM horizontal resolution of 1200 m and 32 column CRM arrays, characteristic of classical superparameterization, is used. In the second baseline test, HRCTRL, a horizontal resolution of 200 m and 64 column CRM arrays are used instead, i.e. 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 ten for the time step. 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.

389

390

391

392

393

394

395

396

397

398

399

400

401

402

403

404

405

406

407

408

409

410

411

412

413

414

415

A first look at maps of time-mean low cloud fraction and absorbed shortwave radiation 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 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. Figure 6a,b shows these expected baseline signals – as we refine the CRM horizontal resolution from 1200m (LRCTRL) to 200m (HRCTRL), less cloud coverage results in more absorbed shortwave radiation (ASR), systematically. When HR is regionalized, the low cloud fraction and cloud brightness differences are just as expected locally - as evidenced by negligible anomalies between the Multi-Domain CRM and HRCTRL simulations within the green-boxed regions (Figure 6c-f). 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 (Figure 6 a,b). 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 flag boundary in the 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 (Figure 6d,h). 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.

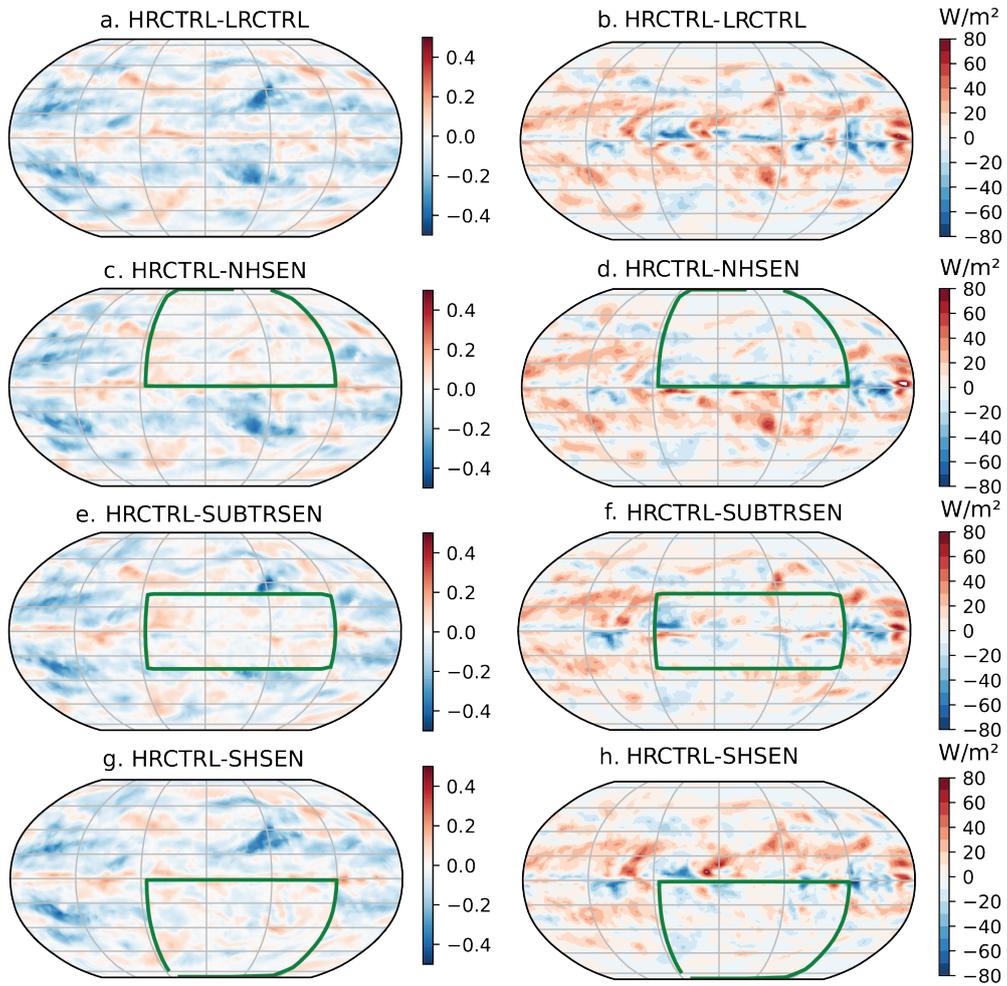
416 Zonally averaged vertical cross sections (Figure 7) of the vertically resolved cloud  
 417 fraction and vertical velocity variance ( $w'w'$ ) in the lower troposphere also show intended  
 418 changes. The cross section is chosen intentionally interior to the heavily loaded zonal sub-  
 419 region, with corresponding longitudes (latitudes) at bottom left and upper right of 160°E  
 420 (70°S) and 85°W (70°N) respectively; the green vertical lines delineate the grid-transition  
 421 boundaries of the three Multi-Domain CRM experiments. Within the regionalized HR  
 422 location, the Multi-Domain CRM simulations (Figure 7c-e) capture the expected  $w'w'$   
 423 enhancement (red contours) and low cloud fraction reduction of the HRCTRL simula-  
 424 tion relative to LRCTRL for all three sensitivity experiments. Our experiments also con-  
 425 firm that the differences between the HR and Multi-Domain CRM configurations are mostly  
 426 minor within the regionalized HR location (inside the green box), while the differences  
 427 between LR and Multi-Domain CRM are minor elsewhere (outside the green box). One  
 428 minor artifact however is that the magnitude of  $w'w'$  within the HR flag area is slightly  
 429 enhanced compared with HRCTRL, especially in the NHSEN and SHSEN experiments  
 430 on just one side of the equator, again hinting at a secondary trade-off when using those  
 431 configurations.

432 Based on the above findings, we hypothesize that when HR is regionalized in a hemi-  
 433 spherically asymmetric manner, this otherwise equatorially symmetric aquaplanet is prone  
 434 to exhibiting artificial ITCZ migrations coupled to deep convection and Hadley circula-  
 435 tion cells. To demonstrate this danger of using Multi-Domain CRM, Figure 8 shows  
 436 the total zonal mean meridional overturning circulation between 10°S and 10°N and the  
 437 precipitation rate focusing on the extreme case of NHSEN where the HR/LR boundary  
 438 is placed over the equator. In NHSEN (SHSEN) the HR grid is used exclusively in the  
 439 northern (southern) hemisphere, causing it to dim preferentially due to HR's reduced  
 440 low cloud fraction. The expectation should then be a shift of the ITCZ to the relatively  
 441 absorptive southern (northern) hemisphere, which is consistent with the zonal mean pre-  
 442 cipitation peaking south (north) of the equator in NHSEN (SHSEN) unlike LRCTRL  
 443 and HRCTRL. Thus while the gross features of the Hadley cell appear similar in the NHSEN  
 444 and HRCTRL simulations (panels a,b; low-level convergence and divergence at both mid-  
 445 and high-levels), a close comparison reveals a  $\approx 16\%$  magnitude anomaly circulation in  
 446 NHSEN-minus-HRCTRL anomaly (Figure 8c), clockwise (via its upper- and lower-most  
 447 branches) in the height-latitude plane, consistent with cross-equatorial flow anomalies  
 448 carrying energy away from the warmer southern hemisphere. This same ITCZ response  
 449 also predicts the cloud fraction and brightness changes noted earlier, where NHSEN tends  
 450 to have less ASR on the southern flank of the equator (Figure 8a) (due to meridionally  
 451 shifted reflection from repositioned deep ITCZ clouds) (Figure 6d) compared with HRC-  
 452 CTRL.

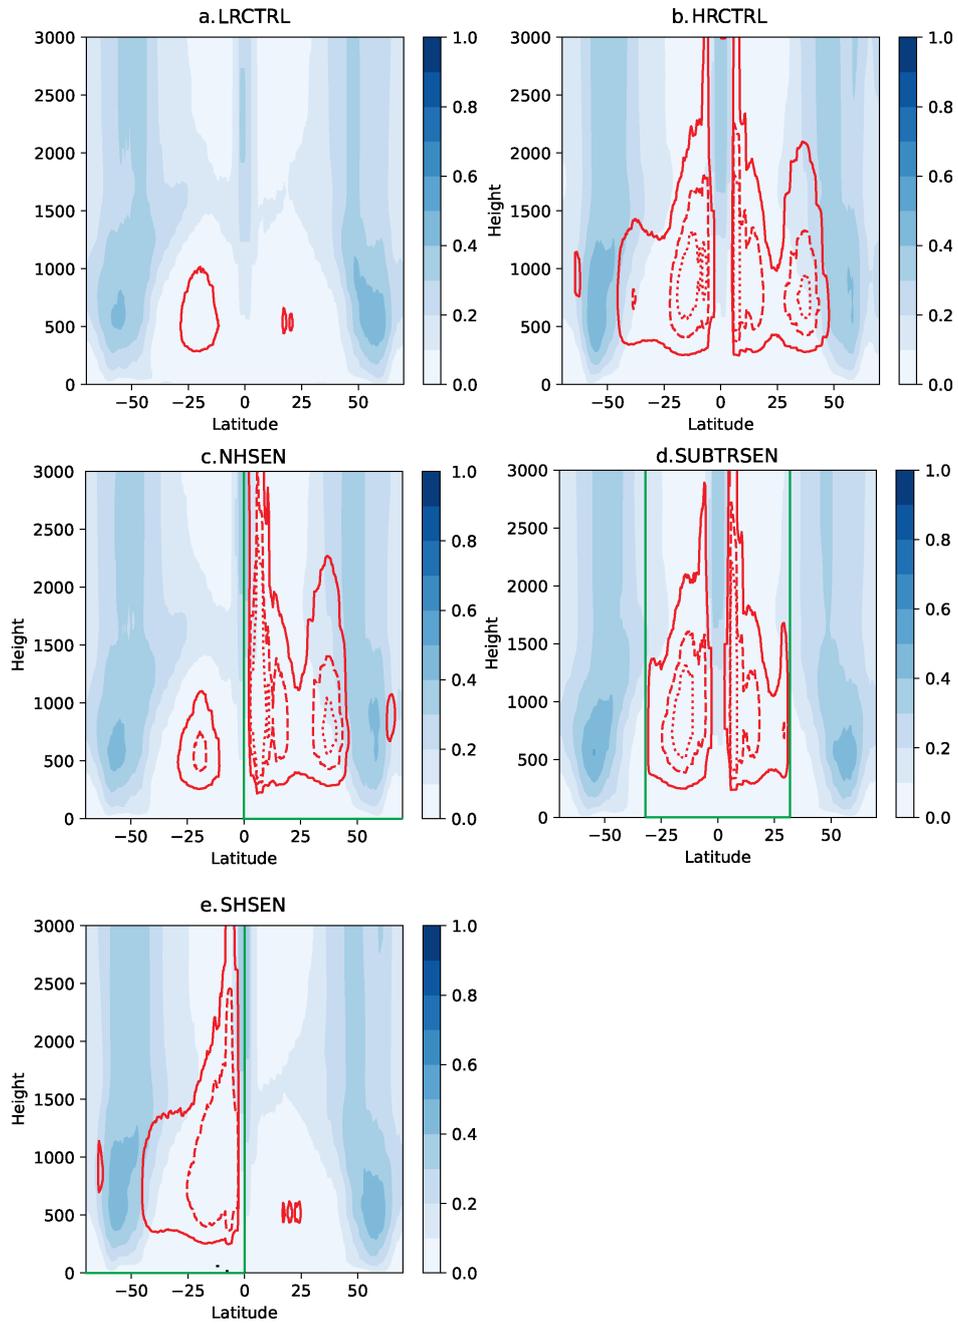
453 The overall assessment of these aquaplanet results is that regionalized HR produces  
 454 cloud brightness and fraction statistics remarkably similar to global HR. While grid tran-  
 455 sition artifacts exist, they do not produce major biases that would argue against the Multi-  
 456 Domain CRM methodology. We can induce artifacts especially in a delicate aquaplanet  
 457 setting by placing CRM grid transitions right on the equator (SHSEN,NHSEN) or in ways  
 458 that imply hemispheric asymmetry, i.e. demanding ITCZ shift responses, but we can eas-  
 459 ily avoid them with a more judicious choice of the regional HR placement (SUBTRSENS).

### 460 3.2 Real-geography hindcast experiments results

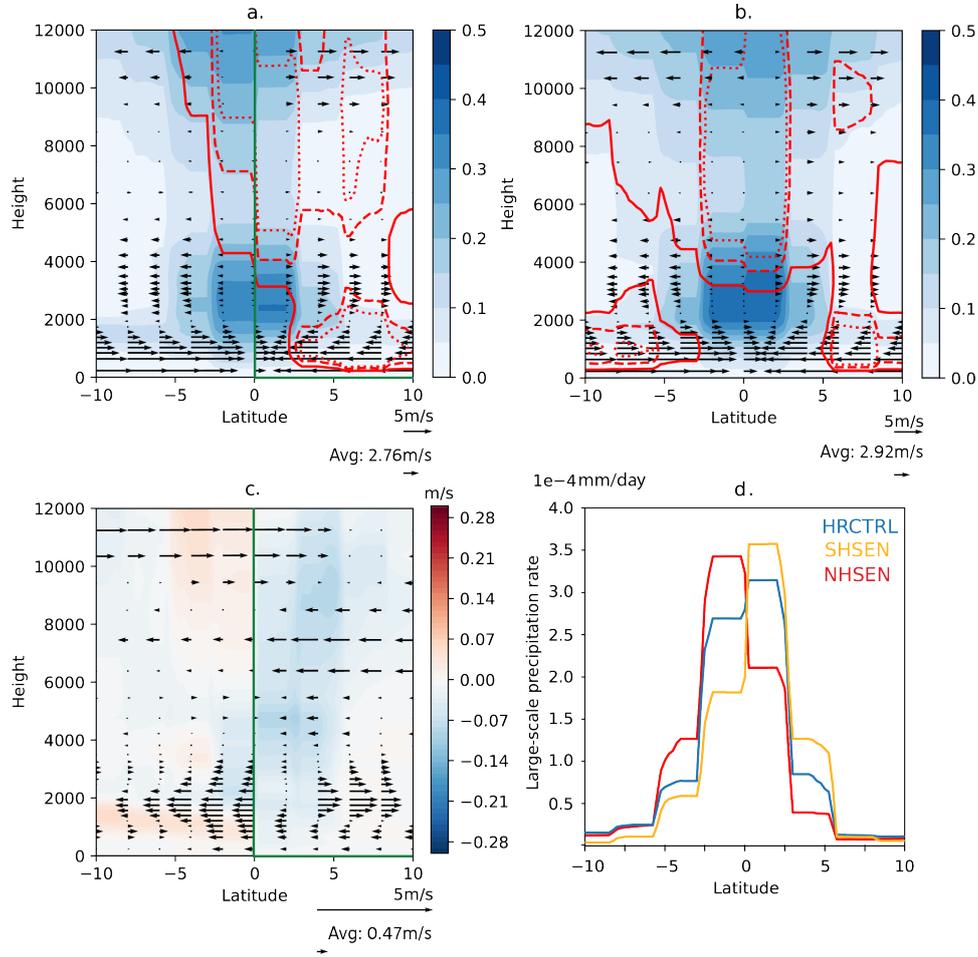
461 The above lessons inform our strategy in the next phase of analysis that transitions  
 462 to a real-geography hindcast simulation setup. Unlike the symmetric aquaplanet, this  
 463 class of experiment does not have any idealizations of meridional symmetry and should  
 464 be expected to be less delicate, but contains additional degrees of freedom for grid tran-  
 465 sition artifacts requiring independent investigation. Similar to our aquaplanet simula-  
 466 tions, there are 6,144 physical columns spanning the globe in these new tests, but now  
 467  $F = 30\%$  (2048 columns) of the planet use the HR configuration. The model is initial-



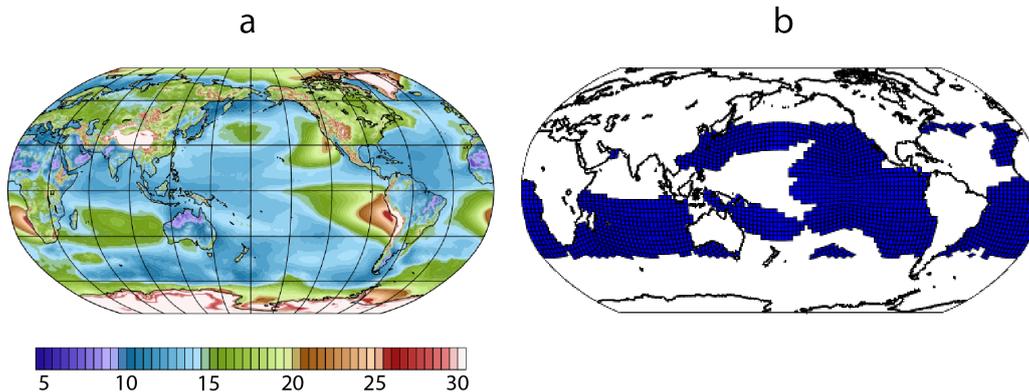
**Figure 6.** The low cloud fraction difference (left) between (a) HRCTRL and LRCTRL, (c) HRCTRL and NHSEN, (e) HRCTRL and SUBTRSEN, (g) HRCTRL and SHSEN. The absorbed shortwave radiation (right) difference between (a) HRCTRL and LRCTRL, (c) HRCTRL and NHSEN, (e) HRCTRL and SUBTRSEN, (g) HRCTRL and SHSEN. The HR mask region is inside the area encompassed by the green line.



**Figure 7.** The zonal averaged crosssection plot for (a) LRCTRL, (b) HRCTRL, (c) NHSEN, (d) SUBTRSEN, and (e) SHSEN. The red contours mark the 0.5 (solid), 0.8 (dashed), and 1.0 (dotted) intra-CRM vertical velocity variance. The green lines represents the grid-transition boundaries.



**Figure 8.** The zonal mean cloud fraction,  $v$  and  $w$  wind components for (a) NHSEN, (b) HRCTRL. The zonal mean cloud fraction,  $v$  and  $w$  wind components differences between (c) NHSEN and HRCTRL. The zonal mean precipitation rate for (d) NHSEN (red solid line), LRC-TRL (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.



**Figure 9.** October climatology (2008-2018) of (a) average lower tropospheric stability (LTS) from ERA5 reanalysis and (b) derived Multi-Domain CRM 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.

468      interpolated ERA5 reanalysis data and forced with prescribed sea-surface tem-  
 469      peratures from the NOAA Optimally Interpolated daily SST dataset (Reynolds et al.,  
 470      2007) for the given initial date, persisted in time. Five separate 7-day simulations are  
 471      performed using Oct. 1 initial conditions taken from independent years spanning 2008  
 472      to 2012.

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

487      We now compare the results of control MMF-HR, MMF-LR and Multi-Domain simu-  
 488      lations, focusing on the ensemble mean of the 7-day hindcast climatology. Figure 10  
 489      shows the ASR biases relative to regrided CERES-SYN daily mean estimates from satel-  
 490      lite (Wielicki et al., 1996). The LRCTRL and HRCTRL simulations create different char-  
 491      acteristic patterns and magnitudes of shortwave biases outside/inside the heavily-loaded  
 492      sub-region; we will return to this point shortly. As in the aquaplanet, and as expected  
 493      from Parishani et al. (2017), as we increase the CRM horizontal resolution, the HRC-  
 494      TRL simulation tends to produce less low cloud fraction (Figure A4) and a positive ASR  
 495      (dim) bias compared with LRCTRL (Figure 10a,b). The LRCTRL simulation tells us  
 496      that LR clouds are systematically too bright in the mid-latitude and trade cumulus re-  
 497      gions, with the exception of some stratocumulus dim biases (Figure 10e,f), whereas the  
 498      HRCTRL simulation has less severe bright biases in the mid-latitudes and a *dim* bias  
 499      ( $0.22 \text{ W/m}^2$ ) throughout most of the subtropics.

500 It is convenient that these two control simulations have different structures because  
 501 it allows us to quantitatively test the Multi-Domain CRM method by separately calcu-  
 502 lating the area-weighted global root-mean square error (RMSE) within those two regions  
 503 ; see subpanel titles in Figure 10. Aggregated *within* the HR masked region (interior to  
 504 the green contour), the spatial root-mean-squared error of the Multi-Domain CRM en-  
 505 semble mean hindcast bias pattern (Figure 10c,  $\text{RMSE}=24.89 \text{ W/m}^2$ ) is a closer match  
 506 to the HR-control simulation (Figure 10g,  $\text{RMSE}=24.92 \text{ W/m}^2$ ) for the same region com-  
 507 pared to the LR control simulation (Figure 10e,  $\text{RMSE}=23.46 \text{ W/m}^2$ ). Meanwhile, the  
 508 Multi-Domain CRM looks very similar to the LRCTRL simulation outside the trade Cu  
 509 region. It has the same brightness bias (Figure 10d,  $\text{BIAS}=0.22 \text{ W/m}^2$  and Figure 10f,  
 510  $\text{BIAS}=0.62 \text{ W/m}^2$ ). That is, Multi-Domain CRM (Figure 10d,  $\text{RMSE}=17.08 \text{ W/m}^2$ )  
 511 has similar RMSE compared with the LRCTRL (Figure 10f,  $\text{RMSE}=17.06 \text{ W/m}^2$ ) sim-  
 512 ulation. A complementary analysis of the lower magnitude OLR biases and RMSE statis-  
 513 tics is available in (Figure A3).

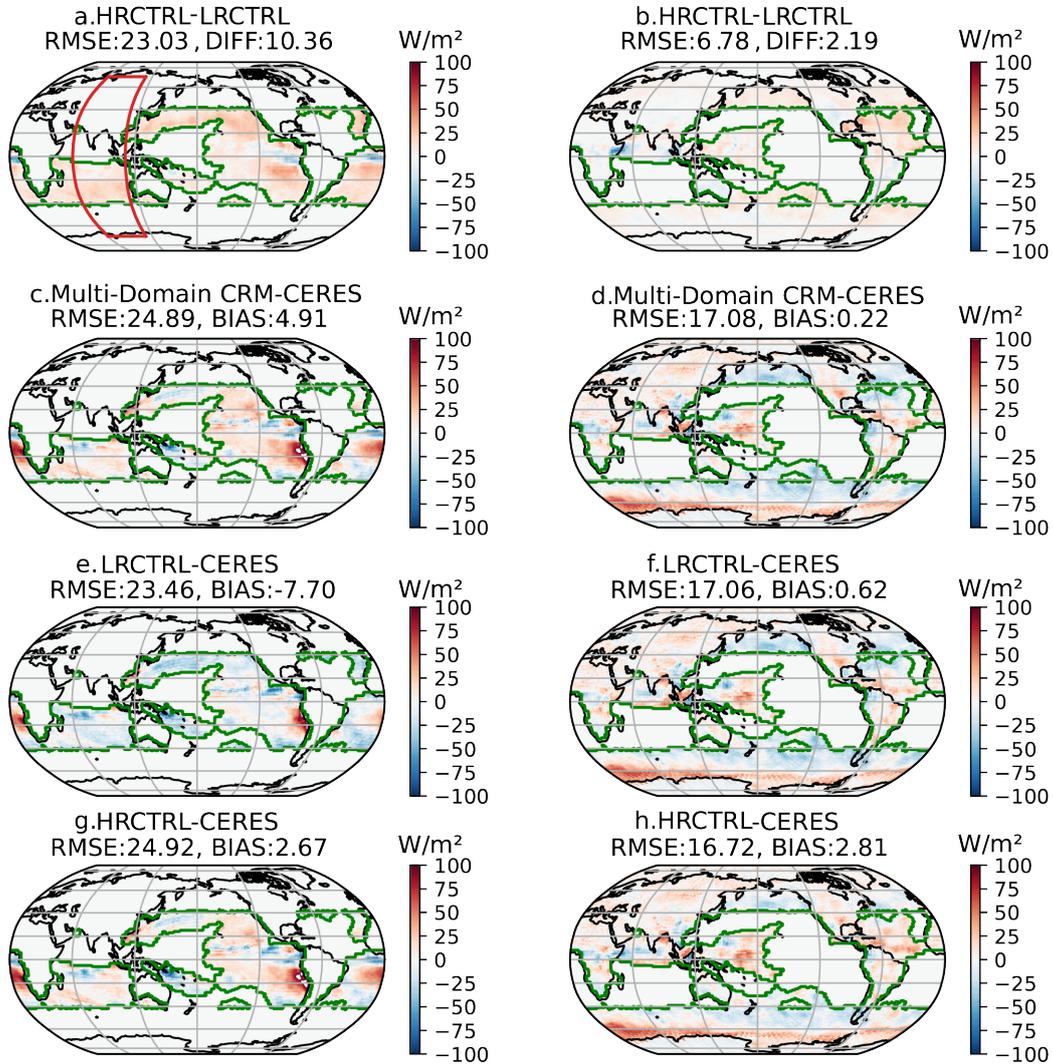
514 Recalling that in the Indian Ocean subregion, we expect the Multi-Domain CRM  
 515 to produce artificial cross-equatorial circulations, we now measure their magnitude. Based  
 516 on the 7-day ensemble mean, the zonal average of the ensemble mean cloud fraction and  
 517 vertical velocity variance are shown in Figure 11, with corresponding longitudes (lati-  
 518 tudes) at bottom left and upper right of  $60^\circ\text{E}$  ( $70^\circ\text{S}$ ) and  $110^\circ\text{E}$  ( $70^\circ\text{N}$ ) respectively (Fig-  
 519 ure 10a). Within the cross section (interior to the green vertical lines) both the cloud  
 520 fraction and vertical velocity variance from Multi-Domain CRM simulation are similar  
 521 to HRCTRL (compare contours and shading in Figure 11a,b). The circulation is mostly  
 522 similar in the Multi-Domain CRM (Figure 11a), but the anomaly vector field in Figure  
 523 11b does reveal a weak anomaly of the zonal mean circulation. Its magnitude is only  $\approx$   
 524  $4\%$ .

## 525 4 Discussion and Conclusions

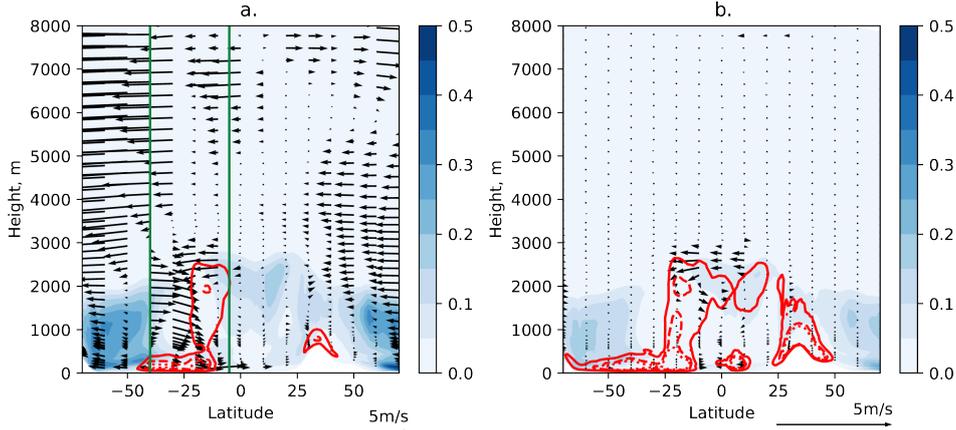
526 By exploring an unusual configuration of a superparameterized climate simulation  
 527 that uses a binary mixture of heavily- versus lightly-loaded physics columns, we have iden-  
 528 tified a limitation of the current parallel load balancing infrastructure in the E3SM and  
 529 CESM: Extreme, rate limiting calculations embedded on the physics side of the code can-  
 530 not be regionalized to small fractions of the planet in ways that make efficient use of am-  
 531 bitious computational resources.

532 We have solved the technical problem by relaxing the current assumption that all  
 533 processing elements be assigned equal sized groups of physics columns. This allows rate-  
 534 limiting regionalized calculations to be unthrottled to their maximum throughput (at  
 535 the upper limit of one pe per physics column) while balancing the remaining load through  
 536 unusually large (many physics columns per pe) column assignments elsewhere. This con-  
 537 cept was motivated by processor-level timing measurements applied to a coarse-resolution  
 538 aquaplanet example, which illustrated the constraints. This informed a general theory  
 539 to predict optimally load balanced and performant conditions under the new assump-  
 540 tion of a binary mixture of heavily- vs. lightly-loaded physics columns. The theory (Equa-  
 541 tion 8) predicts the optimal scale  $P$  for a given climate model with total grid columns  
 542  $C$  containing a fraction  $F$  of heavily-loaded grid columns, each with work overhead  $\tau$ ,  
 543 to achieve maximal throughput and minimal total expense. Predictions are successfully  
 544 validated by actual timing measurements in test configurations that vary both  $F$  and  
 545  $\tau$ .

546 Our own narrow scientific motivation in developing this capability has to do with  
 547 liberating computational resources to meet the cost and throughput needs of Multi-scale  
 548 climate Modeling Framework (MMF) simulations of explicit low cloud feedback. The em-  
 549 bedded CRM resolution requirements of these physics are challenging, since faithfully  
 550 simulating shallow clouds requires resolving small-scale (less than 100-meter) turbulent



**Figure 10.** The map of the ensemble mean absorbed shortwave radiation (ASR) differences between HRCTRL and LRCTRL for (a) inside (b) outside the HR flag area, between Multi-Domain CRM and CERES observations for (c) inside (d) outside the HR flag area, between LRCTRL and CERES observations for (e) inside (f) outside the HR flag area, and between HRCTRL and CERES for (g) inside (h) outside the HR flag area. The HR flag 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.



**Figure 11.** The 7-day ensemble mean zonal mean cloud fraction,  $v$  and  $w$  wind components for (a) Multi-Domain CRM. The ensemble mean zonal mean cloud fraction for (b) 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.

551 eddies in the boundary layer. This is only marginally possible on global scales, even with  
 552 efficiency of MMF, due to expense constraints on the testable resolution and dimension-  
 553 ality of the embedded CRM arrays. As such, recent attempts at high-resolution MMF  
 554 (MMF-HR; Parishani et al. (2017, 2018); Terai et al. (2020)) are physically unsatisfy-  
 555 ing compared to the large eddy simulations that inspire them: The embedded turbulence  
 556 arrays are too small and low-dimensional to exhibit appropriate cellular cloud formation  
 557 organization, and do not allow enough room for organized structures to allow a seam-  
 558 less transition from shallow to deep convection. The interior grid resolution (20-m ver-  
 559 tical spacing, 200-m horizontal) is still too coarse to resolve the spectrum of boundary  
 560 layer eddies that we would like. Yet if, as we desire, we increase the dimensionality, re-  
 561 fine the resolution, or extend the domain size, the computational cost becomes too high  
 562 for the long multi-month simulations needed to study aerosol-cloud feedback.

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

568 Anticipating such applications, we thus performed a set of simulations to exam-  
 569 ine the emergent trade-offs, such as artifacts induced at CRM grid transition boundaries,  
 570 when Multi-Domain MMF is exploited for a binary mixture of low- and high-resolution  
 571 CRM domains. Comparison of Multi-Domain results against globally homogenous stand-  
 572 ard MMF-LR and (expensive) MMF-HR simulations shows that it produces remark-  
 573 ably similar low cloud fractions and shortwave radiative fluxes as the standard MMF-  
 574 HR configuration over the heavily loaded sub-region, as intended, while the rest of the  
 575 globe stays similar to the MMF-LR baseline. Although artifacts can be induced such as  
 576 by positioning the meridional boundary of a CRM grid transition on the equator in an  
 577 otherwise symmetric aquaplanet, in real-geography tests that use sensible sub-regions  
 578 to contain the high resolution, unintended consequences are not obvious.

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

- 581 • (our own interest) increasingly avoid approximation of shallow cloud systems by  
582 embedding increasingly high-resolution and even 3D CRMs in strategic locations  
583 to study their broader climate system interactions at maximum throughput.
- 584 • Increase the throughput and reduce the cost of classical 2D LR by avoiding its use  
585 over areas such as the extratropical ocean where it appears to have marginal ben-  
586 efits (Kooperman et al., 2016a).
- 587 • Afford the expense of large 3D CRMs in classical superparameterization by region-  
588 alizing them just to the tropics, using standard 2D SP or even conventional pa-  
589 rameterization elsewhere; this could assist tuning of coupled LR simulations by  
590 admitting explicit shallow cumulus momentum transport that acts as a throttle  
591 on the equatorial cold tongue (Woelfle et al., 2018), instead of approximating mo-  
592 mentum effects via parameterized scalar transport (Tulich, 2015).
- 593 • Deploy especially large, high-fidelity storm-resolving 3D CRM domains only along  
594 corridors that connect such events’ genesis regions to their main human impacts  
595 sectors, such as where vulnerable societies or valuable infrastructure exist.

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

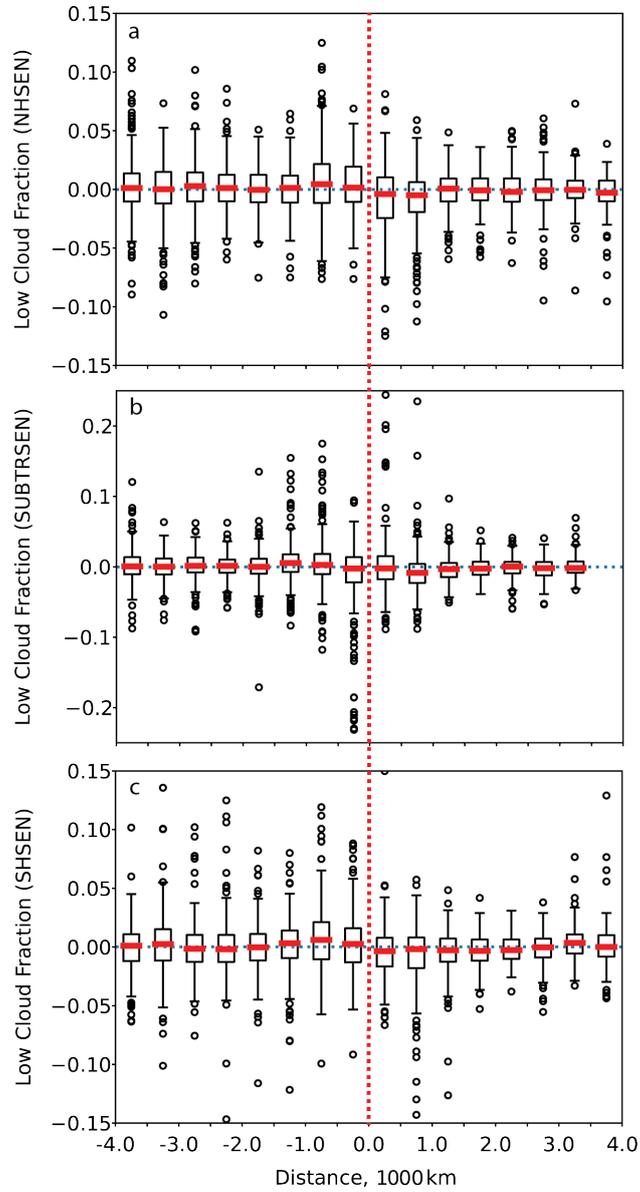
601 More generally, the same load balancing infrastructure that already enables our Multi-  
602 Domain CRM approach should immediately allow any interested E3SM or CESM de-  
603 veloper to focus any form of highly intense, rate-limiting atmospheric physics calcula-  
604 tions over small regions of the planet where this could be helpful to science. To this end,  
605 all code modifications needed to implement our approach within a legacy fork of the E3SM  
606 MMF climate model are available at [https://github.com/mspritch/E3SM/commits/pritch/multiCRM-](https://github.com/mspritch/E3SM/commits/pritch/multiCRM-openmp4.5)  
607 [openmp4.5](https://github.com/mspritch/E3SM/commits/pritch/multiCRM-openmp4.5), in which we also include a flexible infrastructure that allows the user to pro-  
608 vide an auxiliary input file specifying the geographic locations of intense work at run time.  
609 We hope this code and its documentation (10.5281/zenodo.5521784) will enable new ex-  
610 perimentation for other classes of GCM physics where pushing high performance com-  
611 puting limits amidst geographically heterogeneous calculation could be advantageous.

## 612 **Appendix A Additional metric comparing Multi-Domain CRM with** 613 **LRCTRL and HRCTRL**

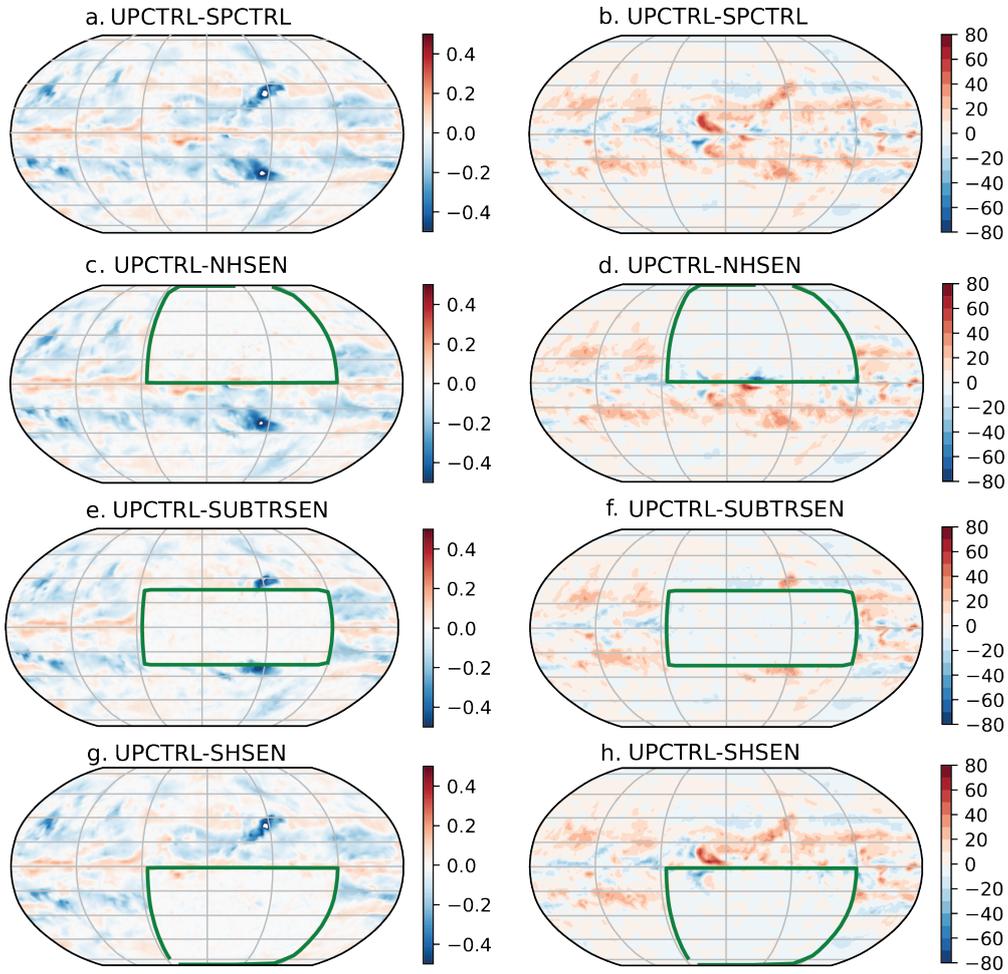
614 To compare variable variations across the HR/LR flag area for aquaplanet simu-  
615 lations, we define the shortest distance of each grid point relative to the HR mask bound-  
616 ary as a new coordinate, i.e. grid points with a smaller distance are closer to the HR/LR  
617 boundary. Inside (outside) the HR flag region, grid point has a positive (negative) dis-  
618 tance. We further distribute all grid points across 16 bins span from -4000 km to 4000  
619 km with a 500 m interval. The box plot (Figure A1) shows the median, lower quartile,  
620 upper quartile, and the outliers of grid points inside each bin. The y axis of the box plot  
621 represents the difference between HRCTRL and Multi-Domain CRM for a positive dis-  
622 tance and LRCTRL and Multi-Domain CRM for a negative distance. This comparison  
623 further confirms that HR/LR boundary has limited impact on the cloud fraction.

## 624 **Acknowledgments**

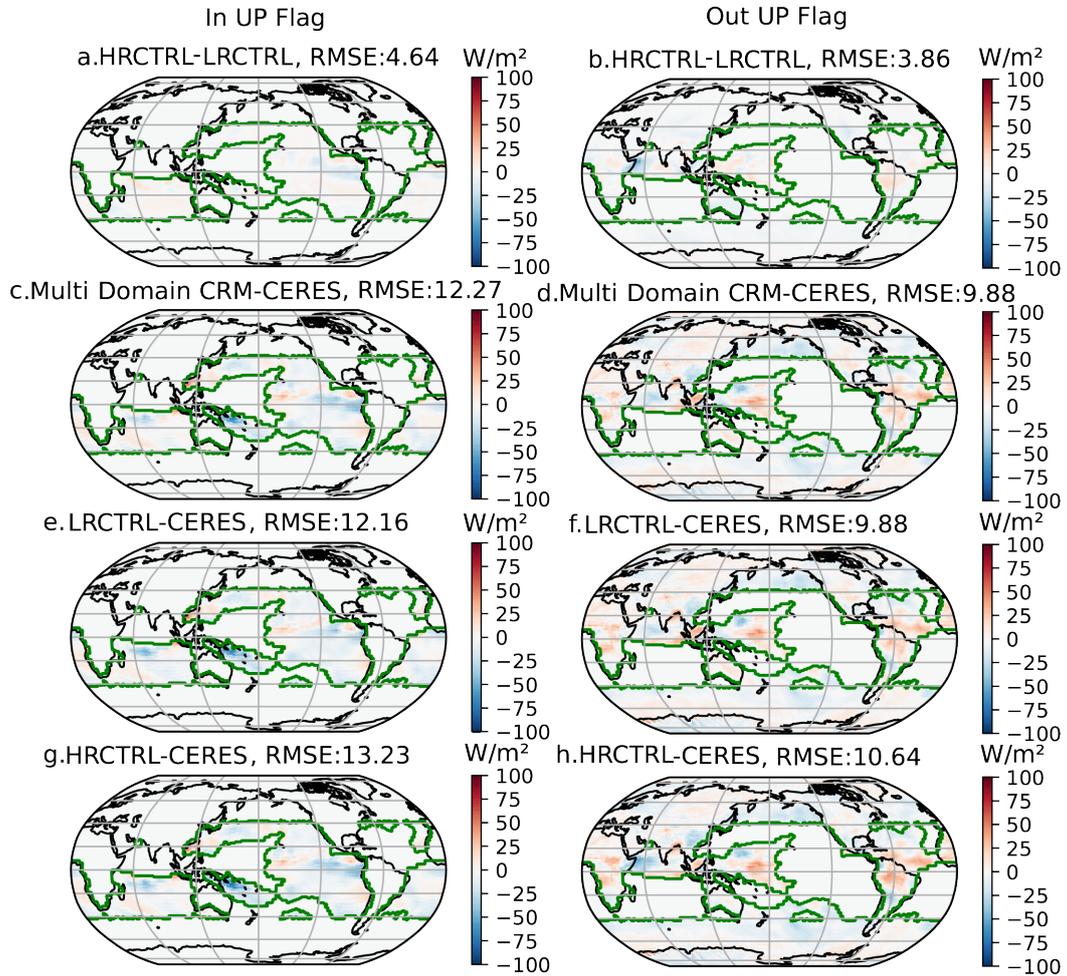
625 LP,MP,PB and CB acknowledge funding from the NSF Climate and Large-scale Dynam-  
626 ics program under grants AGS-1912134 and AGS-1912130; WH,LP and MP acknowl-  
627 edge the work was also performed under the auspices of the U.S. Department of Energy  
628 by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This  
629 research used computational resources from the Texas Advanced Computing Center (TACC)



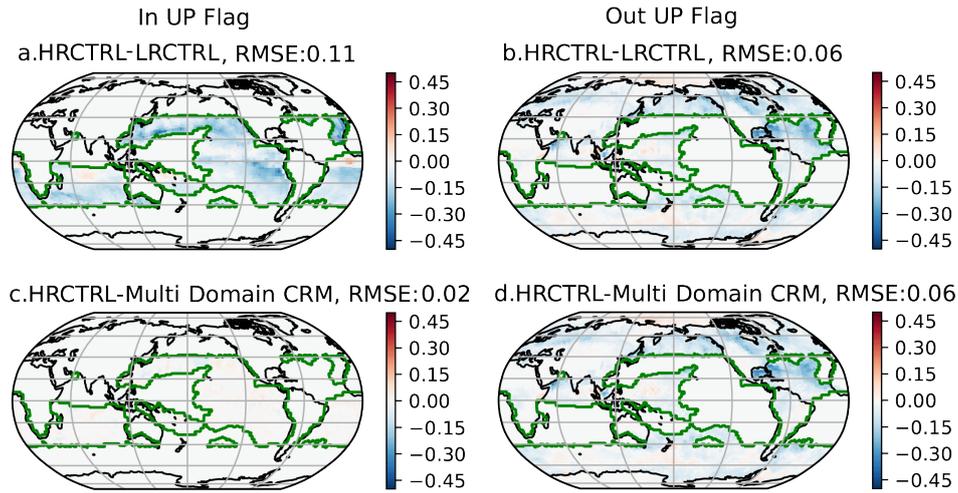
**Figure A1.** Box plot of the cloud fraction differences based on (a) NHSEN, (b) SUBTRSEN, and (c) SHSEN. Columns with positive (negative) distance represent in (out of) the HR flag region. The cloud fraction difference is between HRC<sub>CTRL</sub> and Multi-Domain CRM for positive distances and LR<sub>CTRL</sub> 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) HRCTRL and LRCTRL, (c) HRCTRL and NHSEN, (e) HRCTRL and SUBTRSEN, (g) HRCTRL and 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 HR mask region is inside the area encompassed by the green line.



**Figure A3.** The map of the ensemble mean outgoing longwave radiation (OLR) differences between HRCTRL and LRCTRL for (a) inside (b) outside the HR flag area, between Multi-Domain CRM and CERES observations for (c) inside (d) outside the HR flag area, between LRCTRL and CERES observations for (e) inside (f) outside the HR flag area, and between HRCTRL and CERES for (g) inside (h) outside the HR flag area. The HR flag area is encompassed by the green solid line.



**Figure A4.** The map of the ensemble mean low cloud fraction differences between HRCTRL and LRCTRL for (a) inside (b) outside the HR flag area, and between HRCTRL and Multi-Domain CRM for (c) inside (d) outside the HR flag area.

630 at The University of Texas at Austin under allocation ATM-20009 and from the Extreme  
 631 Science and Engineering Discovery Environment (XSEDE), which is supported by Na-  
 632 tional Science Foundation Grant ACI-1548562 (Towns et al., 2014), under allocation TG-  
 633 ATM190002. The simulation data in this study are stored and can be made available  
 634 to scientific researchers upon request. We gratefully acknowledge many useful conver-  
 635 sations with Xingqiu Yuan, Azamat M. Mametjanov, Tom Beucler, Andrea Jenney, Xi-  
 636 aojuan Liu, Jacob Michael Hendrickson, Griffin Mooers, and Jerry Lin and anonymous  
 637 reviewers for their constructive comments and suggestions.

638 **References**

- 639 Bader, D., Collins, W., Jacob, R., Jones, P., Rasch, P., Taylor, M., . . . Williams, D.  
640 (2014). *Accelerated climate modeling for energy. us department of energy.* Tech.  
641 Rep.
- 642 Buccolieri, R., Salim, S. M., Leo, L. S., Di Sabatino, S., Chan, A., Ielpo, P., . . .  
643 Gromke, C. (2011). Analysis of local scale tree-atmosphere interaction on pol-  
644 lutant concentration in idealized street canyons and application to a real urban  
645 junction. *Atmospheric Environment*, *45*(9), 1702–1713.
- 646 Dennis, J. M., Vertenstein, M., Worley, P. H., Mirin, A. A., Craig, A. P., Ja-  
647 cob, R., & Mickelson, S. A. (2012). Computational performance of ultra-  
648 high-resolution capability in the Community Earth System Model. *Interna-  
649 tional Journal of High Performance Computing Applications*, *26*(1), 5–16. doi:  
650 10.1177/1094342012436965
- 651 Grabowski, W. W. (2004). An improved framework for superparameterization. *Jour-  
652 nal of the atmospheric sciences*, *61*(15), 1940–1952.
- 653 Gustafson Jr, W. I., Vogelmann, A. M., Li, Z., Cheng, X., Dumas, K. K., Endo,  
654 S., . . . Xiao, H. (2020). The large-eddy simulation (les) atmospheric radiation  
655 measurement (arm) symbiotic simulation and observation (lasso) activity for con-  
656 tinental shallow convection. *Bulletin of the American Meteorological Society*,  
657 *101*(4), E462–E479.
- 658 Hagos, S., Leung, R., Rauscher, S. A., & Ringler, T. (2013). Error characteristics  
659 of two grid refinement approaches in aquaplanet simulations: Mpas-a and wrf.  
660 *Monthly Weather Review*, *141*(9), 3022–3036.
- 661 Hannah, W. M., Bradley, A. M., Guba, O., Tang, Q., Golaz, J.-C., & Wolfe, J.  
662 (2021). Separating physics and dynamics grids for improved computational effi-  
663 ciency in spectral element earth system models. *Journal of Advances in Modeling  
664 Earth Systems*, *13*(7), e2020MS002419.
- 665 Hannah, W. M., Jones, C. R., Hillman, B. R., Norman, M. R., Bader, D. C., Tay-  
666 lor, M. A., . . . others (2020). Initial results from the super-parameterized e3sm.  
667 *Journal of Advances in Modeling Earth Systems*, *12*(1).
- 668 Hurrell, J. W., Holland, M. M., Gent, P. R., Ghan, S., Kay, J. E., Kushner, P. J.,  
669 . . . others (2013). The community earth system model: a framework for collabora-  
670 tive research. *Bulletin of the American Meteorological Society*, *94*(9), 1339–1360.
- 671 Jansson, F., van den Oord, G., Pelupessy, I., Grönqvist, J. H., Siebesma, A. P., &  
672 Crommelin, D. (2019). Regional superparameterization in a global circulation  
673 model using large eddy simulations. *Journal of Advances in Modeling Earth Sys-  
674 tems*, *11*(9), 2958–2979.
- 675 Khairoutdinov, M., Randall, D., & DeMott, C. (2005). Simulations of the at-  
676 mospheric general circulation using a cloud-resolving model as a superparam-  
677 eterization of physical processes. *Journal of the Atmospheric Sciences*, *62*(7),  
678 2136–2154.
- 679 Khairoutdinov, M. F., & Randall, D. A. (2001). A cloud resolving model as a  
680 cloud parameterization in the ncar community climate system model: Preliminary  
681 results. *Geophysical Research Letters*, *28*(18), 3617–3620.
- 682 Kooperman, G. J., Pritchard, M. S., Burt, M. A., Branson, M. D., & Randall, D. A.  
683 (2016a). Impacts of cloud superparameterization on projected daily rainfall intensi-  
684 ty climate changes in multiple versions of the community earth system model.  
685 *Journal of Advances in Modeling Earth Systems*, *8*(4), 1727–1750.
- 686 Kooperman, G. J., Pritchard, M. S., Burt, M. A., Branson, M. D., & Randall, D. A.  
687 (2016b). Robust effects of cloud superparameterization on simulated daily rainfall  
688 intensity statistics across multiple versions of the community earth system m  
689 odel. *Journal of Advances in Modeling Earth Systems*, *8*(1), 140–165.
- 690 Mirin, A. A., & Worley, P. H. (2012). Improving the performance scalability of the  
691 Community Atmosphere Model. *International Journal of High Performance Com-*

- 692     *puting Applications*, 26(1), 17–30. doi: 10.1177/1094342011412630
- 693 Neale, R. B., & Hoskins, B. J. (2000). A standard test for agcms including their  
694 physical parametrizations: I: The proposal. *Atmospheric Science Letters*, 1(2),  
695 101–107.
- 696 Parishani, H., Pritchard, M. S., Bretherton, C. S., Terai, C. R., Wyant, M. C.,  
697 Khairoutdinov, M., & Singh, B. (2018). Insensitivity of the cloud response to  
698 surface warming under radical changes to boundary layer turbulence and cloud  
699 microphysics: Results from the ultraparameterized cam. *Journal of Advances in*  
700 *Modeling Earth Systems*, 10(12), 3139–3158.
- 701 Parishani, H., Pritchard, M. S., Bretherton, C. S., Wyant, M. C., & Khairoutdinov,  
702 M. (2017). Toward low-cloud-permitting cloud superparameterization with explicit  
703 boundary layer turbulence. *Journal of Advances in Modeling Earth Systems*, 9(3),  
704 1542–1571.
- 705 Pickering, K. E., Thompson, A. M., Scala, J. R., Tao, W.-K., Dickerson, R. R., &  
706 Simpson, J. (1992). Free tropospheric ozone production following entrainment  
707 of urban plumes into deep convection. *Journal of Geophysical Research: Atmo-*  
708 *spheres*, 97(D16), 17985–18000.
- 709 Prather, M. J., & Jacob, D. J. (1997). A persistent imbalance in hox and nox pho-  
710 tochemistry of the upper troposphere driven by deep tropical convection. *Geophys-*  
711 *ical Research Letters*, 24(24), 3189–3192.
- 712 Rauscher, S. A., Ringler, T. D., Skamarock, W. C., & Mirin, A. A. (2013). Ex-  
713 ploring a global multiresolution modeling approach using aquaplanet simulations.  
714 *Journal of Climate*, 26(8), 2432–2452.
- 715 Reynolds, R. W., Smith, T. M., Liu, C., Chelton, D. B., Casey, K. S., & Schlax,  
716 M. G. (2007). Daily high-resolution-blended analyses for sea surface temperature.  
717 *Journal of climate*, 20(22), 5473–5496.
- 718 Schneider, T., Lan, S., Stuart, A., & Teixeira, J. (2017). Earth system modeling 2.0:  
719 A blueprint for models that learn from observations and targeted high-resolution  
720 simulations. *Geophysical Research Letters*, 44(24), 12–396.
- 721 Terai, C. R., Pritchard, M. S., Blossey, P., & Bretherton, C. S. (2020). The impact  
722 of resolving subkilometer processes on Aerosol-Cloud interactions of Low-Level  
723 clouds in global model simulations. *Journal of Advances in Modeling Earth Sys-*  
724 *tems*, 12(11).
- 725 Tulich, S. (2015). A strategy for representing the effects of convective momentum  
726 transport in multiscale models: Evaluation using a new superparameterized ver-  
727 sion of the weather research and forecast model (sp-wrf). *Journal of Advances in*  
728 *Modeling Earth Systems*, 7(2), 938–962.
- 729 Wielicki, B. A., Barkstrom, B. R., Harrison, E. F., Lee III, R. B., Smith, G. L., &  
730 Cooper, J. E. (1996). Clouds and the earth’s radiant energy system (ceres):  
731 An earth observing system experiment. *Bulletin of the American Meteorological*  
732 *Society*, 77(5), 853–868.
- 733 Woelfle, M., Yu, S., Bretherton, C., & Pritchard, M. (2018). Sensitivity of coupled  
734 tropical pacific model biases to convective parameterization in cesm1. *Journal of*  
735 *Advances in Modeling Earth Systems*, 10(1), 126–144.
- 736 Worley, P. H. (2006). Benchmarking using the community atmospheric model. In  
737 *Proceedings of the 2006 spec benchmark workshop*.
- 738 Worley, P. H., & Drake, J. B. (2005). Performance portability in the physical pa-  
739 rameterizations of the community atmospheric model. *The International Journal*  
740 *of High Performance Computing Applications*, 19(3), 187–201.
- 741 Worley, P. H., Vertenstein, M., & Craig, A. P. (2011). Community Climate Sys-  
742 tem Model. In D. Padua (Ed.), *Encyclopedia of parallel computing* (pp. 342–351).  
743 Boston, MA: Springer US. doi: 10.1007/978-0-387-09766-4\_376
- 744 Wyant, M. C., Bretherton, C. S., & Blossey, P. N. (2009). Subtropical low cloud  
745 response to a warmer climate in a superparameterized climate model. part i:

746  
747

Regime sorting and physical mechanisms. *Journal of Advances in Modeling Earth Systems*, 1(3).