Concurrent Calculation of Radiative Transfer in the Atmospheric Simulation in ECHAM-6.3.05p2

Mohammad Reza Heidari¹, Zhaoyang Song², Enrico Degregori³, Jörg Behrens³, and Hendryk Bockelmann³

¹Department of Informatics, Universität Hamburg, Hamburg, Germany ²School of Atmospheric Sciences, Sun Yat-sen University, and Key Laboratory of Tropical Atmosphere–Ocean System, Ministry of Education, Zhuhai, China ³German Climate Computing Center (DKRZ), Hamburg, Germany

Correspondence:

1. Mohammad Reza Heidari (mohammad.reza.heidari@studium.uni-hamburg.de)

2. Zhaoyang Song (songzhaoy@mail.sysu.edu.cn)

Abstract. The scalability of the atmospheric model ECHAM6 at low resolution, as used in palaeoclimate simulations, suffers from the limited number of grid points. As a consequence, the potential of current high performance computing architectures cannot be used at full scale for such experiments, particularly within the available domain-decomposition approach. Radiation calculations are a relatively expensive part of the atmospheric simulations taking approximately up to over 50% of the total

- 5 runtime. This current level of cost is achieved by calculating the radiative transfer only once in every two simulation hours. In response, we propose to extend the available concurrency within the model further by running the radiation component in parallel with other atmospheric processes to improve scalability and performance. This paper introduces the *concurrent radiation scheme* in ECHAM6 and presents a thorough analysis of its impact on the performance of the model. It also evaluates the scientific results from such simulations. Our experiments show that ECHAM6 can achieve a speedup over 1.9x using the
- 10 concurrent radiation scheme. By performing a suite of stand-alone atmospheric experiments, we evaluate the influence of the concurrent radiation scheme on the scientific results. The simulated mean climate and internal climate variability by the concurrent radiation generally agree well with the classical radiation scheme, with minor improvements in the mean atmospheric circulation in the Southern Hemisphere and the atmospheric teleconnection to the Southern Annular Mode. This empirical study serves as a successful example that can stimulate research on other concurrent components in atmospheric modeling whenever
- 15 scalability becomes challenging.

1 Introduction

20

Earth system modeling has traditionally been a computationally demanding domain with a continual increase in complexity and resolution. It is a major application of high performance computing and represents a large variety of climate processes with diverse computational profiles and disparate performance optimization requirements. A primary subsystem of a typical Earth System Model (ESM) is the atmospheric simulation, which resolves several physical processes including radiative transfer. Radiative transfer is one of the most expensive parts for coarse and low-resolution atmospheric simulations. This process is resolved to respond to the changing state of the chemical species which interact with the radiation (Balaji et al., 2016; Salby, 1996; Wallace and Hobbs, 2006). Solar energy is the driving force for the atmosphere through the radiative transfer, which is the only physical process that is capable of exchanging energy between a planet like the Earth and the rest of the universe (Wallace

- 25 and Hobbs, 2006). Energy transfer in the atmosphere involves electromagnetic radiation that can be separated into short and long wave parts: *shortwave*, emitted by the sun, and *longwave*, emitted by the earth's surface and the atmosphere (Wallace and Hobbs, 2006; Salby, 1996). There are several atmospheric processes including greenhouse gases, aerosols and clouds that interact with electromagnetic radiation through the mechanisms of absorption, scattering and emission. The level of interaction strongly depends on the state of the atmospherics particles (evolving by advection, cloud processes and chemistry) and the
- 30 optical properties (the wavelengths and intensity) of the incident radiation (Wallace and Hobbs, 2006).

Forecasting System (IFS) model.

In principle, the absorption of solar radiation by the atmosphere and the earth's surface must be balanced by the longwave emission to space from the terrestrial radiation (Salby, 1996). It is crucial for atmospheric models to accurately represent the radiative transfer process (Rasp, 2019). Solving the problem is in essence straightforward (Wallace and Hobbs, 2006). However, this can be quite computationally demanding in practice, despite the simplifying approximations adopted in the

35 radiation component (Balaji et al., 2016; Wallace and Hobbs, 2006). As a result, in most climate models around the world, this component is not called in every time step (Balaji et al., 2016). On the contrary, it is calculated at a coarser time step than the rest of atmospheric physics, entirely pursuing a performance improvement rather than fulfilling any other technical objective (Balaji et al., 2016).

Over the years, various techniques have been used to represent radiative transfer in different models and maintain its calculation cost within acceptable limits. Morcrette (2000) discusses the effect of the temporal or spatial sampling techniques of the radiation inputs on the forecasts and analyses at the European Centre for Medium-Range Weather Forecasts (ECMWF). In a different approach, the radiative calculations are computed on a coarser grid than the one on which all other physical processes are implemented. Morcrette et al. (2008) report the implementation of a reduced grid for radiation in the ECMWF Integrated

- 45 Resolving radiation transfer on coarser time and spatial resolutions can however lead in errors in weather and climate simulations. A report by (Hogan and Bozzo, 2015) describes a computationally efficient solution to this problem. It suggests updating the surface longwave and shortwave fluxes in every time step and grid point according to the local skin temperature and albedo. Authors in (Hogan and Hirahara, 2016) propose a careful treatment of the cosine of the solar zenith angel in models that calaculate radiation every 3h. This solution aims at reducing the negative impacts of the biases that occur due to discrete
- 50 sampling of solar zenith angle. A follow-up study by (Hogan and Bozzo, 2018) introduces a flexible new radiation scheme (ecRAD) in the ECMWF model which is around 41% faster than the previous package. The report shows some improvements in the skill of weather forecasts by calling the radiation scheme more frequently for the same overall computational cost.

Another proposal is based on the reorganization of the radiation component within the atmospheric models. In this approach, the radiative transfer is calculated in separate tasks in parallel with the rest of the model, in pursuit of improved scalability and

55 performance. In the classical atmospheric modeling, all physical processes including solar radiation are resolved sequentially

with respect to each other; thus creating a prohibitively long latency in the overall runtime. Reorganizing the radiation component as separate parallel tasks allows for the simultaneous calculation of radiative transfer along with the other atmospheric processes and removes the long response latency from the component. Mozdzynski and Morcrette (2014) reports a reorganization of the radiation calculations in IFS at ECMWF and demonstrates the radiation-in-parallel configuration, in which

60

the model. In a similar effort but in a broader sense, Balaji et al. (2016) proposes the *Coarse-grained Component Concurrency* (CCC) to increase the level of concurrency within ESMs. This approach suggests reorganizing more lower- and higher-level components

calculating radiative transfer is performed on separate MPI (Message Passing Interface) processes in parallel with the rest of

in parallel with each other and exploiting fine-grained parallelism within each component individually. Additionally, the report 65 demonstrates the result of applying this approach to the radiation component of the Geophysical Fluid Dynamics Laboratory (GFDL) Flexible Modeling System (Balaji, 2004). In this use case, the atmospheric radiative transfer is configured to run in parallel with the atmospheric dynamics and all other atmospheric physics. This technique uses shared-memory parallelism and divides the available threads in an OpenMP (Open Multi-Processing) region between the parallel components.

This paper concentrates on the performance optimization of the radiation scheme in ECHAM6, which is the sixth generation

- of the atmospheric general circulation ECHAM (Stevens et al., 2013). The model was developed at the Max Plank Institute 70 for Meteorology (MPI-M) in Hamburg. It is the traditional atmospheric component of the coupled Earth System Model MPI-ESM, as described by Giorgetta et al. (2013). ECHAM6 benefits from spectral and finite difference methods in five different grid resolutions, ranging from the coarse (CR) and low resolution (LR) to the very high resolution (XR). The CR or T31 is a truncation to 31 wave numbers in the spectral part and corresponds to a horizontal spatial resolution of 96×48 points in
- longitude and latitude. The LR or T63, however, corresponds to 192×96 points. Its special prominence in this research is due 75 to its application in the German climate modeling initiative (PalMod), which aims at simulation of a complete glacial cycle (i.e., about 120000 years) from the last interglacial to the Anthropocene (https://www.palmod.de). There, however, remains a serious caveat as to the feasibility of such an ambitious project which should be acknowledged in advance. In particular, a major concern has been raised over the poor performance of ECHAM6 suffering from the limited number of grid points at
- 80 setups used in palaeoclimate simulation. For this reason, the performance optimization of the model is instrumental in ensuring the viability of such long-time simulations.

Experiments reveal that the radiation component is one of the most expensive computational parts in ECHAM6, at least for palaeoclimate simulation. Endeavors to adopt higher optimized radiative calculations will, therefore, be opportune for PalMod experiments. In response, two solutions have been investigated in parallel within the PalMod project to alleviate

- 85 the computational burden of the radiative transfer on the atmospheric simulations in ECHAM6: single-precision arithmetic and *concurrent radiation scheme*. The technique of mixed-precision calculations is a practice of reducing time-to-solution of scientific algorithms whenever lower accuracy is permitted. Empirical studies have revealed that radiation calculations can benefit from reduced-precision arithmetic. Cotronei and Slawig (2020) discuss the results of applying single-precision arithmetic to the radiation calculations in ECHAM6. They indicate that this mathematical treatment accelerates the component
- by about 40% while the overall runtime of the model reduces by 18%. 90

This paper, on the other hand, presents a report on the concurrent radiation scheme applied to the atmospheric model ECHAM6 and provides a thorough analysis on the performance and stability of the model. Calculating radiative transfer in parallel with other atmospheric processes can potentially affect the model's accuracy since the radiation fields will always lag one more radiation time step behind in comparison with the classical scheme. This lag may have negative impacts on physical

95 processes that benefit from a tighter coupling in time with radiation. The boundary-layer clouds, particularly stratocumulus, are a good example. They are maintained by longwave cooling at cloud tops once they are formed. This could explain why (Hogan and Bozzo, 2018) found that calling radiation more frequently led to more skillful forecasts of near-surface temperature and low cloud cover.

In contrast to the OpenMP approach used by Balaji et al. (2016), the concurrent radiation scheme opts for the MPI parallelization in order to fully exploit the potential of higher concurrency in the model. In addition, encapsulating the radiation

- calculations in a distinct component and name space realizes the idea of separation of concerns (SoC), which results in more degrees of freedom. An immediate benefit includes the independent development and optimization of the radiation component from the main model pursuing higher throughput which is essential for the ambitious long simulation runs of the PalMod project. This architectural merit also enables the potential of combining the virtues of the concurrent radiation scheme with
 other appropriate optimized solutions such as "single-precision arithmetic in ECHAM radiation" (Cotronei and Slawig, 2020)
- 105

115

in the future.

The paper is organized as follows: Section 2 describes the classical approach to the radiation calculations in ECHAM6. In Section 3, we introduce the solution of a new radiation scheme in the model. Finally, the performance analysis and the scientific evaluation of the new scheme are presented in section 4 and 5, respectively.

110 2 The Classical Radiation Scheme in ECHAM6

Radiative transfer in ECHAM6 is represented with PSrad/RRTMG (a postscript to the Rapid Radiative Transfer Model for GCMs (Pincus and Stevens, 2013)) for both shortwave and longwave parts of the electromagnetic spectrum (Stevens et al., 2013). The radiation component is one of the most expensive components within atmospheric physics. As a result, this component is stepped forward at a slower rate than the rest of the atmospheric physics in ECAHM6 as well as most of the climate models around the world (Balaji et al., 2016). In its flagship configuration, ECHAM6 updates the optical properties of radiation every two hours, except for very high-resolution (T255) simulations where the radiation is calculated hourly. Figure 1 shows the organization of the classical radiation scheme in ECHAM6. As it is apparent in this figure, the radiative transfer is generally calculated once in every n normal atmospheric time steps, i.e. $\Delta t_{rad} = n * \Delta t_{atm}$.

In this scheme, the radiation results are used beyond the state of the input tracers, which may be as much as $\Delta t_{rad} - \Delta t_{atm}$ (approximately two hours) behind. Figure 1 schematically gives a clear account of what happens. As it shows, there are multiple normal atmospheric time steps between two consecutive radiation time steps. The atmospheric calculations are provided with old feedback from the radiation component within the normal time steps. On non-radiation time steps in between an update of the optical properties, longwave irradiance is rescaled based on the surface temperature while shortwave irradiance is rescaled



Figure 1. The organization of the classical radiation scheme in ECHAM6: the radiative transfer is resolved sequentially with respect to the other atmospheric physics and dynamics and it is stepped forward at a slower rate than the other atmospheric processes.



Figure 2. The relative time contribution of radiation calculations in ECHAM6 (using the classical radiation scheme) in the CR simulations.

by the zenith angle (Stevens et al., 2013). Infrequent calculations of the radiative heating may result in numerical instability
in climate models, as described by Pauluis and Emanuel (2004). As pointed out by Balaji et al. (2016), the use of the lagged state can be viewed as a potential source of discrepancy between the cloud field and the "cloud shadow field" seen by the radiation component. This, therefore, introduces numerical errors in atmospheric models and becomes considerably worse at higher resolutions (Xu and Randall, 1995; Balaji et al., 2016). Although the choice of larger radiation time steps (Δt_{rad} > Δt_{atm}) evidently reduces the overall runtime of simulations, the radiative portion is yet considered relatively high for some configurations in ECHAM6. As shown in Figure 2, the radiative calculations take up from almost 40% to 58% of the total simulation time at the CR resolution (which is also one of the PalMOD settings), depending on the number of MPI processes

assigned to the model.

The entire crux of the problem in the classical radiation scheme can be attributed to the sequential organization of the components inside the model. In this architecture, atmospheric processes are stepped forward one at a time in every time step and the

- 135 computation time of each component directly contributes to the overall simulation time. In particular, the radiation component significantly delays the following calculation of other atmospheric physics and dynamics during the entire course of simulation. As it is apparent in Figure 1, this architecture prolongs the radiation time step in proportion to the high computational cost of the radiation component. It will be shown in the next section that this long response time of the radiation calculations is not, however, inevitable and can be avoided by reorganizing the component inside the model.
- 140 Moreover, the sequential organization of components in the classical ECHAM6 creates another obstacle that hinders the optimization of the model. In fact, ECHAM6 traditionally benefits from MPI and implements domain decomposition parallelism to expedite the computations. As Figure 3 shows, the radiation calculations display a higher scalability than the main model in this framework. This can be attributed to the columnwise organization of the atmospheric physics and the embarrassingly parallel nature of the workload. In other words, since the individual columns (iterating the k index in an (i, j, k) discretiza-
- 145 tion) have no cross-dependency in (i, j), it allows for fine-grained parallelism (Balaji et al., 2016). This is the reason why the radiation component can intrinsically scale better beyond the limitations of the main model. Such a higher scalability is instrumental in reducing the high computational cost of the radiation calculations. However, the sequential architecture of the classical ECHAM6 restrains the benefit by forcing the radiation component to use the same computational resources as the rest of the model. As a consequence, the component is hindered by the limited scalability of the whole model. In the next section,
- 150 however, it will be shown that re-organization of the components in ECHAM6 is essential to improve the overall performance of the model.



Figure 3. The scaling curve of the radiation component vs. the atmospheric model ECHAM6 shows that the radiative transfer has a higher scalability and keeps scaling beyond the limitations of the main model.

The Concurrent Radiation Scheme 3

It was discussed in the previous section that, even in light of larger radiation time steps, radiation calculations vet impose a daunting cost on the atmospheric simulation at coarse and low resolution in ECHAM6. It was also shown that the sequential treatment of resolving atmospheric processes is implicated in the long response time of the radiation component in every radiation time step and restricts the benefit of the higher scalability of the radiation calculations.

The concurrent radiation scheme, on the other hand, puts forward a feasible solution to the problem. It implements an additional level of parallelism inside the model by applying coarse-grained component concurrency to the radiation calculations. This approach eliminates the high response latency from the radiation time step and paves the way for a higher scalable model.

- In contrast to the classical scheme, the concurrent radiation scheme starts resolving the radiative transfer much earlier be-160 fore the next radiation time step arrives. As a result, the main model receives feedback from the radiation component much faster upon the request. This technique minimizes the response latency of the component and reduces the overall simulation time. In this approach, the radiative transfer is calculated concurrently with other atmospheric processes along the course of normal time steps. The coupling fields are also exchanged between the radiation component and the main model within the
- 165 radiation time steps but without the typical delay experienced in the classical scheme. Figure 4 describes a method for casting the radiative transfer as a concurrent component using distributed memory computing. This technique organizes the radiation component and the main model on separate MPI processes and enables the concurrent calculation of the radiative transfer and other atmospheric processes.

Due to the close time-dependency of radiation on evolving model fields, the concurrency between the atmosphere and ra-

- 170 diation component is only enacted between consecutive radiation time steps. In addition, the synchronization between the concurrent components takes place during each radiation time step. As it is shown in the next section, the interprocess synchronization overhead is, however, negligible compared to the cost of the radiation calculations. The data exchange between the radiation and the atmospheric model benefits from the communication library YAXT (Behrens et al., 2014), which was developed at German Climate Computing Center (DKRZ) in Hamburg. YAXT simplifies the formulation of the communica-
- tion problem and generates suitable communication objects to efficiently execute the data exchange. The library is specially 175 suited for bulk communication as given in this use case. This is due to the automatic generation of MPI datatypes that enable direct access to model data without requiring additional data copies or pack/unpack overhead to create messages. YAXT is built on top of MPI and takes high level descriptions of arbitrary domain decomposition and automatically derives an efficient collective data exchange. Figure 5 shows the use of the YAXT library for coupling two concurrent components with different 180
- domain decomposition.

155

In addition, running the concurrent components on separate MPI processes is a prerequisite to applying arbitrary domain decomposition to the radiation calculations. It offers a way forward to scale the component beyond the limits traditionally imposed by the atmospheric model. Hence, the concurrent radiation scheme prepares the ground for improving the physical consistency between the radiative and physicochemical atmospheric states. This approach can accordingly minimize the dis-

crepancy between Δt_{rad} and Δt_{atm} (Pauluis and Emanuel, 2004; Xu and Randall, 1995; Balaji et al., 2016). Furthermore, the 185



Figure 4. The re-organization of the radiation component in parallel with the rest of atmospheric physics and dynamics in ECHAM6. In the first (radiation) time step, ATM (the main model) sends the input data to RAD (the radiation component) but it has to wait long until it receives the results from RAD. In the following radiation time steps, however, the data exchange takes place immediately one after the other (i.e. ATM first receives the results of radiation calculations and immediately provides the input data to RAD for the next radiation calculations.). This way, ATM is supposed to experience a minimum idle time when it interacts with RAD.



Figure 5. YAXT library facilitates MPI communication between concurrent components with different domain decomposition layouts.

independent resource allocation feature is essential for an efficient load balancing between concurrent components and parallel efficiency of the model. The next section presents these merits in more detail with some concrete examples.

4 Results I: Performance analysis

190

This section presents the performance evaluation of the concurrent radiation scheme in ECHAM6 in comparison with the classical approach. It should be emphasized that the new scheme utilizes almost the same original implementation of radiation calculations with a radically different orchestration. The new organization, therefore, exerts a major impact on the overall simulation time rather than the pure computational performance of the radiation component. In consequence, the performance

evaluation presented in this section explicitly aims at the assessment of the whole model and will not be limited to the radiation component. For the purpose of this study, a new version of ECHAM6 (based on ECHAM-6.3.05p2) is deployed with both

195 classical and concurrent radiation schemes which can be configured to calculate the radiative transfer with or without separate MPI processes. Experiments are performed on the *Mistral* supercomputer at DKRZ on a machine configuration with Intel Haswell processors (E5-2680v3 12C 2.5GHz) and a Mellanox FDR Infiniband high speed interconnect. All runs allocate a layout of one MPI process per CPU core on the computing nodes equipped with two processors which are exclusively dedicated to the experiments. The performance results presented in this paper are obtained from the Atmospheric Model Intercomparison 200 Project (AMIP) experiments performing simulations at the CR from 1976 to 1981.

To understand the impact of the concurrent radiation scheme on the overall performance, it is useful to extract the scaling curves of the model with both radiation schemes and study the gained speedup via the concurrent radiation scheme. For this purpose, the frequency at which the radiative transfer is calculated is by default every two hours in all runs, i.e. $\Delta t_{rad} = 8 * \Delta t_{atm}$ with the normal time step of 15 minutes. Additionally, once the model is configured to use the concurrent radiation scheme, an equal number of MPI processes, and thus identical domain decomposition, is assigned to the radiation component and the

main model.

Figure 6 compares two scaling curves which reflect the performance of the model with the classical (the blue curve) and concurrent radiation schemes (the red curve). The horizontal axis shows the total number of MPI processes allocated by the model. It is worth emphasizing that ECHAM6 (using the classical radiation scheme) uses the same MPI processes to calculate

- 210 radiative transfer and other atmospheric processes. However, when the model is configured to use the concurrent radiation scheme, half of the allocated MPI processes are exclusively dedicated to the radiation calculations and the other half to the rest of the atmospheric physics and dynamics. The Y-axis, on the other hand, reflects the throughput of the model in terms of the number of simulated years per day (SYPD). As it can be inferred from Figure 6, ECHAM6 can achieve only a maximum performance of 450 SYPD at 576 MPI processes using the classical radiation scheme at the CR resolution. However, it yields
- 215 a significant improvement using the concurrent radiation scheme and reaches a maximum performance of 734 SYPD at 1152 MPI processes. It is noteworthy that, due to the limited number of grid points at the CR resolution, running the classical model at higher domain decomposition is not justified theoretically. Needless to say, it does not attain any significant performance improvement in practice either, as asserted in Figure 3 where the scaling curve of ECHAM6 tends to flatten towards the end. This should explain why the blue curve in Figure 6 stops at 576 MPI processes, as opposed to the red curve scaling beyond.
- 220 On this account, the concurrent radiation scheme acquires a new significance as it becomes conducive to higher scalability of the model.



Figure 6. The scaling curves of ECHAM6 using the classical and concurrent radiation schemes.

The red curve in Figure 7 displays the methodical speedup of the model using the concurrent radiation scheme. Here, the methodical speedup means the improved runtime of the model by making use of the proposed concurrency, in contrast to the classical definition of speedup, where additional resources are used for the same computation. The methodical speedup is therefore the ratio of the overall performance of the model using the concurrent radiation scheme (using *N* resources) divided by the performance of the model using the classical radiation scheme (using *N* resources). The methodical speedup **S** is defined

 $SYPD_{concurrent}$ = The number of simulated years per day with the concurrent radiation scheme (using 2N MPI processes)

230 $SYPD_{classical}$ = The number of simulated years per day with the classical radiation scheme (using N MPI processes) (2)

Methodical Speedup =
$$S = \frac{SYPD_{concurrent}}{SYPD_{classical}}$$
 (3)

On this account, for each point on the speedup curve(s), the number of resources assigned to the model using the classical radiation scheme is half of the resources allocated by the model using the concurrent radiation scheme. Hence, the X-axis indicates only the total number of allocated MPI processes to the model if the concurrent radiation scheme is used by the model. However, the model allocates half of the MPI processes shown at the X-axis when it adopts the classical radiation scheme. The red curve shows that that the model achieves an actual speedup ranging from a minimum 1.6x to over 1.9x when it benefits from the concurrent radiation scheme. The blue curve, however, shows the asymptotic speedup which the model would achieve if there were no cross-dependency, and thus no communication latency, between the radiation component and

the main model.

225

as below:



Figure 7. The red curve shows the methodical speedup of ECHAM6 using the concurrent radiation scheme. The blue curve shows the methodical speedup that the model would achieve asymptotically.

Moreover, the coupling fields between the radiation component and the rest of the model are exchanged at every radiation time step and potentially contribute to the overall simulation time. There is nevertheless mounting evidence in Figure 8 indicating that the communication overhead compared to both radiation and total runtime of the model is negligible (less than 1.4%) and it has, therefore, little impact on the total simulation time. Consequently, the performance of the model (using concurrent radiation scheme) is mainly affected by the relative cost of the radiation calculations, which has a strong dependency on the number of allocated computing resources.

245

Figure 9 shows the efficiency of the resource utilization when the model runs the concurrent radiation scheme. The parallel efficiency of the concurrent radiation scheme is defined as the ratio of the methodical speedup S to the relative number of the allocated resources R, as shown below.

250 Ratio of allocated resources =
$$R = \frac{\text{(resources assigned to the concurrent radiation scheme)}}{\text{(resources assigned to the classical radiation scheme)}} = \frac{2N}{N} = 2$$
 (4)

Parallel Efficiency =
$$E = \frac{S}{R} = \frac{S}{2}$$
 (5)

As visible in the figure, the model achieves a parallel efficiency of 80% or more across the scaling curve. However, attaining the maximum parallel efficiency requires an optimal distribution of workload among MPI processes. A close investigation reveals a load imbalance between the concurrent components inside the model. This problem starts appearing when the radiation and other atmospheric calculations are configured to use identical domain decomposition and allocate an equal number of MPI processes. The measured parallel efficiency in Figure 9 suggests that the MPI processes assigned to the radiation calculations experience an idle time during the course of simulation.



Figure 8. Comparing the communication costs of data exchange between the radiation component and the main model to the computation time of the radiative transfer (the red curve) and other atmospheric processes (the blue curve) and to the overall simulation time (the green curve).



Figure 9. The parallel efficiency of ECHAM6 (using the concurrent radiation scheme).

260

Figure 10 depicts the length of time that the concurrent components have to remain idle, waiting for the slower component to catch up and become ready to exchange the radiation results. The red and blue curves show the respective idle times that the radiation component (denoted by RAD) and other atmospheric processes handled by the main model (denoted by ATM) experience. The idle times appear when the radiative transfer is resolved faster or slower than the other atmospheric calculations. At lower numbers of MPI processes, as shown in Figure 10, calculating the radiative transfer takes longer and it, thus, forces ATM to wait relatively long in each radiation time step for feedback from RAD. At higher numbers of MPI

265

processes, however, RAD scales better and finishes the calculations faster than ATM. It therefore has to wait for the arrival of the next radiation time step so that the radiation results can be transferred to the main model. It should be noted that Figure 10 shows the maximum length of time that an MPI process (assigned to RAD or ATM) has to wait for its peer MPI process (assigned to the other component) until it catches up. From Figure 10, it can reasonably be inferred that the total idle time experienced by ATM and RAD becomes minimum at 384 MPI processes. This can also explain why the parallel efficiency in 270 Figure 9 has an extremum at this point. It is also apparent in Figure 10 that the radiation component experiences a longer idle time as the number of MPI processes increases. This behavior accounts for the higher scalability of the radiation component, which was already reflected in Figure 3.



Figure 10. The red curve (RAD) shows the increasing idle time of the MPI processes responsible for resolving radiative transfer. It suggests that RAD is a dominant computation in all the configurations before 576 MPI processes and does not wait for ATM. The blue curve (ATM), on the other hand, shows that resolving other atmospheric physics and dynamics is not a dominant computation for all the configurations before 576 MPI processes and ATM therefore experiences a long idle time. However, ATM becomes dominant towards the end, inflicting a long idle time on RAD. Yet the idle time of ATM is not lifted completely. This can be attributed mainly to the unavoidable waiting time of ATM in the first (radiation) time step (as reflected in Figure 4) and some infrequent (slightly) longer radiation time steps.

The load imbalance between the concurrent processes is directly affected by the number of MPI processes assigned to the radiation component as well as the frequency at which the radiative transfer is calculated. Figure 11 schematically illustrates a contrived configuration in which the radiation component is forced to remain idle almost half of the total runtime. To remove such an idle time, Figure 12 provides an inspiring example. It takes advantage of the higher scalability of the radiation com-



Figure 11. A contrived model configuration in which the same domain decomposition is assigned to the radiation calculations and the main model. Here, this setup imposes an idle time on the radiation component, leading to an inefficient resource usage.

ponent and reduces the radiation time step to a half. This new configuration eliminates the idle time from the MPI processes assigned to the radiation component and hence removes the load imbalance between the concurrent components successfully. As a consequence, the resource utilization of the model improves significantly and attains a parallel efficiency close to 100% without affecting the achieved speedup. This example presents a viable solution to decrease the gap between Δt_{rad} and Δt_{atm} in pursuit of a more consistent atmospheric model.

280

285

290

The concurrent radiation scheme, however, puts forward a general solution to remove the load imbalance between the radiation component and the main model. This solution provides a remedy for the idle time imposed on the main model at some configurations (such as 48, 96, 192, 288 or 384 MPI processes as shown in Figure 10) which exhibit a suboptimal resource efficiency due to the slow calculation of radiative transfer. In this approach, the radiation component is enabled to adopt finer domain decomposition and allocates a higher number of resources (in comparison to the main model) in order to catch up with the fast calculation of other atmospheric processes. By the same token, Figure 13 suggests a configuration in which the radiation component is also a remedy to remove the load imbalance at the configurations (such as 576, 768 and 1024 MPI processes as shown in Figure 10) in which the radiation component experiences a long idle time due to the slow calculation of other atmospheric processes.

In addition, the concurrent radiation scheme offers an opportunity for coupling the radiation component to the other atmospheric processes at every normal time step (i.e. $\Delta t_{rad} = \Delta t_{atm}$). This feature can ultimately bring the model to the physical consistency between the radiative and physicochemical atmospheric states, albeit probably with a negligible impact on the model's accuracy. It is notable that the current implementation of the concurrent radiation scheme in ECHAM6 already provides the technical support for the adoption of finer or coarser domain decomposition for the radiation calculations. In particular, the

295



Figure 12. A model configuration for assigning identical domain decomposition to the radiation component and the main model but reducing the radiation time step to a half ($\Delta t_{rad} = n/2 * \Delta t_{atm}$), hence improving the parallel efficiency of the model.

YAXT library simplifies the data exchange between the concurrent components with disparate domain decomposition. The scientific viability of these schemes, however, requires further investigations and the results will be presented in a follow-up paper.



Figure 13. A model configuration for assigning coarser domain decomposition to the radiation calculations. This setup allows for allocating a lower number of MPI processes to the radiation component and hence improves the parallel efficiency of the model.

Results II: Scientific evaluation 300 5

We evaluate the concurrent radiation scheme in ECHAM 6.3.05 version at the CR and LR configurations, hereafter termed as CR_{CRR} and LR_{CRR}, which differ in their horizontal grid spacings. Both CR_{CRR} and LR_{CRR} share the vertical resolution of 47 levels. The classical radiation scheme has been tuned to optimize the simulated climate (Mauritsen et al., 2019). However, the concurrent radiation scheme at these two configurations have not been individually tuned. The parameters of the convection

- 305 scheme, e.g., the convection conversion rate for cloud water to rain in the concurrent radiation scheme are the same as in the classical scheme. The evaluation and documentation of the concurrent radiation scheme in ECHAM6.3 is based on the Atmospheric Model Intercomparison Project (AMIP) historical experiments. Experiments were performed according to the AMIP protocol (Eyring et al., 2016). The historical forcings include the emissions of short-lived species and long-lived greenhouse gases, solar forcing, stratospheric aerosol forcing, monthly sea surface temperatures and sea ice concentrations. Experiments
- 310
- span from 1960 to 2013. The monthly output from 1980 to 2013 are taken for analysis. AMIP experiments performed with the classical ECHAM6.3.05, as well as the observations and reanalysis, serves as a reference. Observational and reanalysis data used in this paper are listed in Table 1.

Variable	Validation data	Period	Reference
Surface land Temperature	HADCRU4	1980-2013	Jones et al. (2012)
Total Precipitation	GPCP	1980-2013	Adler et al. (2003)
Surface LW radiation	CERES-EBAF	2001-2013	Loeb et al. (2012)
Surface SW radiation	CERES-EBAF	2001-2013	Loeb et al. (2012)
Zonal mean temperature	ERA-interim	1980-2013	Dee et al. (2011)
Zonal mean zonal wind	ERA-interim	1980-2013	Dee et al. (2011)
Air pressure at sea level	ERA-interim	1980-2013	Dee et al. (2011)

Table 1. Observation and reanalysis data for model evaluation.

Mean State 5.1

315

We first evaluate the simulated mean climate by the concurrent radiation scheme. CRUTEM4 observations are interpolated to the model grids for the comparison against simulated surface-air temperature (SAT). Across all configurations, global annual bias in SAT over land exhibits similar spatial patterns (Figure 14a, 14c and S1a, S1c). General warm bias occurs over eastern Siberia, central Asia, Tibetan Plateau, North America and central Australia, while sporadic cold biases exist in south Asia, Africa and northern South America. For a rough assessment, global root-mean-square-error (RMSE) is calculated for SAT. Horizontal resolution exhibits larger impact on the SAT bias than the radiation scheme. Both radiation scheme shows smaller SAT bias in LR relative to CR (Figure 14a, 14c and S1a, S1c). The concurrent radiation scheme illustrates a somewhat ambiguous reduction in the annual SAT bias, albeit insignificant at the 95% confidence interval (Figure 14e).





Figure 14. Annual bias in surface-air temperature (SAT) over land (K) and total precipitation (mm/day) in the (a, b) concurrent and (c, d) classical radiation experiments relative to Research Unit time series 4.01 data set (CRU TS; Harris et al., 2014) and the Global Precipitation Climatology Project data set v2.3 (GPCP; Adler et al., 2003) for the period 1980–2013, respectively. Differences in the (e) SAT over land and (f) total precipitation between the concurrent and classical radiation scheme for the period 1980-2013. Hatching indicates significance at the 95% confidence interval using Student's t-test.

All experiments suffer from severe precipitation bias over the tropics as compared to the GPCP satellite measurements (Figure 14b, 14d and S1b, S1d). Specifically, excessive precipitation is simulated over the tropical Indian Ocean, western tropical Pacific and western tropical Atlantic Ocean, while deficient precipitation occurs over the east tropical Pacific Ocean,

325

330

South America and east tropical Atlantic Ocean. Consistent with the response of SAT bias, the global RMSE decreases along with a refinement of horizontal resolution, whereas the concurrent radiation scheme leads to larger bias in the precipitation. Particularly, the precipitation bias over the equatorial region of South America and western equatorial Africa increases in the concurrent radiation (Figure 14f and S1f), which is likely linked to the intensification of local SAT bias (Figure 14e and S1e). Clouds play a complex and crucial role in the earth's radiation budget, which significantly affects the simulated surface climate. To understand the biases in SAT, we evaluate the errors in cloud feedback by estimate cloud-radiative-forcing (CRF) following (Cess et al., 1990).

$$CRF = F_{all-sky} - F_{clear-sky} \tag{6}$$

where F_{all-sky} refers to radiative flux at the top-of-atmosphere (TOA) and F_{clear-sky} represents the radiative flux assuming the absence of clouds. CRF is positive when clouds warm the surface atmosphere, and vice versa. The radiative effect of
clouds consists of two competing components in the radiation budget: warming impact on the surface through the emission of longwave (LW) radiation and cooling impact by shading the shortwave (SW) radiation. The sign and spatial structure of CRF bias is largely determined by the SW CRF bias, which are partially compensated by the LW CRF bias (Figure 15 and S2). CRF exhibits larger bias in CR relative to LR with both radiation scheme (Figure 15a, 15b, 15d, 15e, and S2a, S2b, S2d, S2e). In CR_{CRR} and CR_{SEQ}, strong negative SW CRF biases exist over regions of upwelling and the western coast of Australia, whereas
strong positive CRF biases occur over the Arctic, the Southern Ocean, the eastern tropical Indian Ocean and the northern equatorial Pacific and Atlantic Ocean. A refinement of horizontal resolution largely alleviates negative CRF biases over the

upwelling regions, whereas the CRF bias over other regions changes barely. Over the tropics, the concurrent radiation scheme shows a slight increase of the bias compared to the classical radiation scheme (Figure 15g, 15h and S2g, S2h).

Increased CRF bias in the concurrent radiation scheme is largely attributed to the response of SW CRF errors (Figure 15g and S2g), with negligible contribution from the errors in LW CRF (Figure 15h and S2h). Across all experiments, there are widespread discrepancies between the spatial distribution of biases in the CRF and SAT except over the Tibetan Plateau, where positive CRF bias agrees well with the warm SAT bias. To further explore such inconsistency, biases in the net SW radiation at TOA are estimated (Figure 15c, 15f and S2c, S2f). SAT biases over the North America, the Tibetan Plateau and central Australia are associated with the excessive net SW radiation at TOA, resulting from the negative biases of the surface albedo (not shown). In contrast, warm SAT biases over eastern Siberia and tropical South America exist along with deficient SW

radiation biases, which suggests a dynamical cause for the biased SAT. A misrepresentation of the atmospheric circulation in ECHAM6.3.05 may be responsible for this discrepancy. The concurrent radiation scheme exhibits strong negative biases in the net SW radiation at TOA in the tropics relative to the classical radiation scheme (Figure 15i and S2i). This implies that model parameters concerning the cloud formation should be carefully adjusted for the concurrent radiation scheme. Specifically, the



Figure 15. Annual bias in shortwave cloud radiative forcing (CRF), longwave CRF and net short wave radiation fluxes (W/m²) at the top-ofatmosphere (TOA) in the (a, b, c) concurrent and (d, e, f) classical radiation experiments relative to data from Clouds and the Earth's Radiant Energy System Energy Balanced and Filled product (CERES-EBAF) surface fluxes edition 4.0 for the period 2001-2013. Differences in the (g) shortwave CRF, (h) longwave CRF and (i) net shortwave radiation at TOA between the concurrent and classical radiation scheme for the period 2001-2013. Hatching indicates the significance at the 95% confidence interval using Student's t-test.

relative humidity threshold for cloud formation in the upper troposphere and the lowest model level should be changed to 355 improve the match with observational records (Mauritsen and Roeckner, 2020).

The zonal mean temperature biases in the troposphere are smaller than in the stratosphere for all experiments (Figure 16 and S3). In the lower troposphere (between 850 and 1000 hPa), a cold bias amounting to -4 K occurs over the Antarctic for all experiments. Thermal biases in the mid troposphere are relatively small. Moderate warm biases up to 2 K extend from the tropics to the Arctic. In the lower stratosphere (between 250 hPa and 100 hPa), all experiments suffer from prominent warm biases over the tropics and mid-latitudes, and severe cold biases over the high-latitudes in both hemispheres. Such biases are significantly reduced in LR (Figure S3a and S3c) relative to CR (Figure 16a and 16c) experiments, consistent with the notion by (Stevens et al., 2013) that zonal mean temperature bias can be significantly reduced by enhancing the horizontal resolution. The concurrent radiation scheme affects the biases in zonal mean temperature in CR simulations barely, yet slightly reduces

365 the biases in the lower stratosphere between 40°S and 90°S (Figure 16e and S3e).

The patterns of zonal mean westerly wind biases are common to all configurations (Figure 16b and 16d and S3b, S3d) and reflect the meridional structure of temperature biases (Figure 16a and 16c and S3a, S3c). In CR_{CRR} and CR_{SEO}, temperature biases drive a northward shift of westerly winds between 30°S and 60°S. LR_{CRR} and LR_{SEO} exhibit alleviated westerly wind biases than their CR counterpart due to reduced temperature bias. In the tropics, westerly wind biases are characterized by

overly strong easterly in the low- and mid- troposphere, and large westerly biases in the upper troposphere. The concurrent 370 radiation scheme intensifies the easterly wind biases over the tropics in CR and LR (Figure 16f and and S3f). As suggested by (Stevens et al., 2013), the tropical bias pattern is an indication of excessive heating associated with the deep convection.

5.2 **Climate Variability**

360

To explore the simulated climate variability by the concurrent radiation, we present an analysis of the ENSO teleconnection 375 and interannual variability of in the extra tropics, in the form of the northern and southern annular modes.

5.2.1 **ENSO feedbacks and teleconnections**

The El Niño/Southern Oscillation (ENSO) is the leading mode of interannual variability in the tropical Pacific. ENSO is mainly characterized by the variations of sea surface temperature (SST) in the eastern and central equatorial Pacific. Multiple negative and positive coupled atmosphere-ocean processes that either favor or suppress the growth of SST anomalies governs the 380 ENSO behavior (Philander, 1989; Neelin, 1998; Jin et al., 2006). Equatorial Pacific SST anomalies associated with ENSO can affect the tropical convection and results in zonal shifts of the Walker Circulation (Philander, 1989; Bayr et al., 2014). Further, the changes in the convection stimulate Rossby waves that propagates to the mid- and high-latitudes through the atmospheric bridge. Many studies (Guilyardi et al., 2004; Guilyardi, 2006; Guilyardi et al., 2009) have shown that the atmospheric model dominates the simulated ENSO properties in the coupled climate models. Therefore, it is crucial to evaluate the ENSO

385 teleconnections that project the influence of ENSO globally.



Figure 16. Mean bias in zonal-mean temperature (K) and zonal wind (m/s) in the (a, b) concurrent and (c, d) classical radiation experiments relative to ERA-Interim for the period 1980–2013 (shading) and the climatological mean from ERA-Interim (contours). Differences in (e) zonal-mean temperature and (f) zonal wind between the concurrent and classical radiation for the period 1980-2013. Hatching indicates the significance at the 95% confidence interval using Student's t-test.



Figure 17. Covariance (cov(mm/day, mm/day)) of DJF (December-January-February) precipitation anomalies with the normalized time series of precipitation anomaly in the Niño4 region for (a) GPCP observations, (b) concurrent and (c) classical radiation experiments. Covariance (cov(m, mm/day)) of DJF 500 hPa geopotential height anomalies in the Northern Hemisphere with the normalized time series of precipitation anomaly in the Niño4 region for (d) ERA-Interim, (e) concurrent and (f) classical radiation experiments.

The covariance between the global DJF (December-January-February) precipitation anomalies and the DJF precipitation anomalies over the Niño 4 region is shown in Figure 17 ans S4. CR_{SEQ} and LR_{SEQ} typically capture the pattern over the tropical Pacific, tropical Atlantic and North America (Figure 17c and S4c) as indicated by the GPCP observations. The concurrent radiation scheme exhibits similar teleconnection patterns to the classical radiation scheme, yet influences the magnitude of the response to ENSO. CR_{CRR} underestimates the magnitude of atmospheric teleconnection over the Maritime Continent (Figure 17b), whereas LR_{CRR} overestimates the teleconnection over this region (Figure S4b). Additionally, CR_{CRR} and LR_{CRR} exhibits artificially positive response in the southern tropical Indian Ocean relative to their counterpart using the classical radiation scheme.

395

Covariance of DJF geopotential height anomalies at 500 hPa (Z500) and normalized DJF precipitation anomalies over the Niño 4 region is calculated to investigate the diabatic forcing of the tropical Pacific on the boreal winter atmospheric circulation in the Northern Hemisphere (Figure 17 and S4). The ERA-interim depicts positive covariance over Canada and southern Greenland, and negative relations over northeast Pacific, north Atlantic and Siberia. CR_{SEO} is able to simulate the pattern over northeast Pacific, Canada and Greenland. However, it fails to capture the component over Siberia and the signal over the North Atlantic is shifted northeastward. Increasing the horizontal resolution to LR shows no improvement for the

400 teleconnection pattern with the classical radiation scheme. CR_{CCR} retains the main features of ENSO teleconnection as shown by CR_{SEQ} , whereas LR_{CCR} exhibits weak response of atmospheric circulation to the diabatic forcing over the tropical Pacific compared to the ERA-interim (Figure 17d). This implies that a tuning for the convection scheme in the concurrent radiation scheme is required.

5.2.2 Northern Annular Mode

- 405 The Northern and Southern Hemisphere Annular mode (NAM and SAM, also known as the Arctic and Antarctic Oscillation) are the leading modes of variability of the extratropical circulation in both hemispheres. Both annular modes explain the month-to-month and year-to-year (especially the cold-season) variability of the atmospheric circulation, which exhibit pronounced impacts on the climate in the mid-latitudes and the polar region (McAfee and Russell, 2008; Kidston et al., 2009). The NAM is defined as the leading empirical orthogonal function (EOF) of the DJF sea level pressure (SLP) in the Northern Hemisphere.
- 410 As shown by the ERA-interim reanalysis data, the NAM is characterized by zonally symmetric structures (Figure 18a). There are large and positive loadings over the polar cap region, surrounded by zonally ring-shaped negative loadings centered over the northeast Pacific and the North Atlantic. The leading mode explains 30.4% of total variance. All experiments show larger explained variance of the NAM than the ERA-interim (Figure 18 and S5). Horizontal resolution does not show strong influence on the NAM pattern, except that CR_{CRR} and CR_{SEO} exhibit slightly higher explained variance relative to their LR counterparts.
- 415 The concurrent radiation scheme captures the centers of action better than the classical scheme, especially over the North Atlantic sector. The center of action over North Atlantic is shifted eastward in LR_{SEQ} compared to the ERA-interim and LR_{CRR} .

The atmospheric teleconnection associated with the NAM is calculated by regressing the DJF SAT anomalies on the time series of principal components (PC1) corresponding to the leading EOF (Figure 18b, 18d, 18f and S5b, S5d, S5f). Associated

- 420 with the positive phase of the NAM, warm SAT anomalies occur over Europe, Siberia and the western United States, while cold temperature anomalies exist over western Alaska, far eastern Russia, Greenland and eastern Canada (Figure 18b). CR_{SEQ} significantly underestimate the response over Greenland, Baffin Bay and Siberia, but overestimate the response over western America. Additionally, an artificially negative response occurs over western Canada (Figure 18b). CR_{CCR} share similarities with CR_{SEQ}, with minor differences in the magnitude of the response, which is also seen in the LR experiments. The atmospheric
- teleconnection simulated by two LR experiments generally agree better with the ERA-interim than the CR simulations, which reproduces the widespread negative SAT anomalies over Europe and Siberia (Figure S5d and S5f).

5.2.3 Southern Annular Mode

Similarly, the JJA (June-July-August) SAM is defined as the leading EOF of JJA SLP anomalies in the Southern Hemisphere. Similar to the NAM, the SAM depicts large loadings over the Antarctic and three surrounding centers of action over the



Figure 18. The leading empirical orthogonal functions (EOF) of December-January-February (DJF) sea level pressure (SLP) anomalies calculated from (a) ERA-interim, (c) concurrent (e) classical radiation experiments for the period 1980-2013. Linear regressions of the DJF 2m-air-temerpature on the NAM index (normalized PC1) from (b) ERA-interim, (d) concurrent and (f) classical radiation experiments for the period 1980-2013.

- southern Pacific, Indian and Atlantic Ocean at approximately 45°S (Figure 19a). CR_{CCR} and CR_{SEO} simulate slightly westward 430 shifted centers of action over the southern Indian (Figure 19c and 19e). The explained variances by the leading EOFs are overestimated in these two experiments. A refinement of horizontal resolution improves the simulation of the SAM pattern. Among all experiments, LR_{CCR} shows the best agreement with the ERA-interim (Figure 19a and S6c). This may be linked to small biases in the mean atmospheric circulation (Figure S3b).
- 435 The SAT response to positive phase of SAM is computed by regressing the SAT anomalies on the PC1 of SLP in the Southern Hemisphere (Figure 19 and S6). There are warm temperature anomalies over most of the Antarctic and Ross Sea, while cold anomalies exist over the Weddell. LR_{CCR} and LR_{SEO} underestimate the SAT response over the Weddell Sea whereas CR_{CCR} and CR_{SEO} fails to reproduce the SAT response over both the Ross Sea and Weddell Sea (Figure 19d, 19f and S6d, S6f). Overall, an increase in the horizontal resolution improves the simulated climate variability and its atmospheric teleconnections both in 440 the Northern and Southern Hemisphere, however, the concurrent radiation scheme changes these features barely.

Conclusions 6

This paper presents the implementation of the concurrent radiation scheme in the atmospheric model ECHAM6 and demonstrates its impact on the performance and stability of the model. A detailed analysis shows that the radiative transfer is a relatively expensive component of the model especially for the CR and LR resolution, which are also used in paleoclimate

- simulation. Although the component exhibits a higher scalability profile, the study reveals that it cannot freely scale to its full 445 potential due to the sequential architecture of the classical ECHAM6. The concurrent radiation scheme, on the other hand, organizes the radiation component in parallel with the rest of the atmospheric physics and dynamics and hides its long computation time. The experiments asserted explicitly a noticeable model speedup across the scaling curve with a strong dependency on the relative computational cost of the radiative transfer to the other atmospheric processes. Unlike the classical scheme,
- 450 this approach enables the radiation component to adopt any viable domain decomposition arbitrarily which may differ from the main model's configuration. The component can accordingly follow a different scaling scheme and benefit from the higher scalability of radiation calculations. This salient feature can eventually decrease the discrepancy between the radiation time step $\Delta t_{\rm rad}$ and normal atmospheric time step $\Delta t_{\rm atm}$ with the objective of creating more physical consistency in the model.
- The simulated mean climate and internal climate variability by the classical and concurrent radiation scheme have been eval-455 uated. A suite of AMIP experiments have been performed on CR and LR configurations using the two radiation schemes. In terms of long-term mean state biases, e.g., biases in land surface temperature, precipitation, cloud radiative forcings, and zonal mean temperature and wind, a refinement of horizontal resolution exhibits better agreement with the observations. The concurrent radiation scheme generally yields similar results with the classical radiation, except some minor improvements in the mean atmospheric circulation in the Southern Hemisphere. Regarding to the climate variability and associated atmospheric telecon-
- 460 nections, LR simulations agree better with the ERA-interim than their CR counterparts. The concurrent radiation scheme on LR improves the atmospheric teleconnection to the SAM, which is likely linked to the alleviated bias in the mean circulation. On the other hand, the classical radiation scheme on LR shows better atmospheric teleconnection to ENSO than the concurrent



Figure 19. The leading empirical orthogonal functions (EOF) of June-July-August (JJA) sea level pressure (SLP) anomalies calculated from (a) ERA-interim, (c) concurrent and (e) classical radiation experiments for the period 1980-2013. Linear regressions of the JJA 2m-air-temerpature on the SAM index (normalized PC1) from (b) ERA-interim, (d) concurrent and (f) classical radiation experiments for the period 1980-2013.

radiation. One possible reason is that the classical radiation scheme in ECHAM6 has been properly tuned towards the observations. To conclude, the concurrent radiation scheme presented in this study substantially improves the scalability of ECHAM6,

465 with the major features in the mean climate and internal variability retained.

Code and data availability.

1. The source code of the atmospheric model ECHAM6 adopted to the project PalMod (for the concurrent execution of radiative transfer) and used for generating the plots presented in this paper is available under https://doi.org/10.35089/WDCC/SC_PalMod_ECHAM6 (MPI-M and DKRZ, 2021).

2. The output data for generating the plots presented in this paper is available under https://doi.org/10.5281/zenodo.4589140 (Heidari et al., 2021).

Author contributions.

1. MH wrote Abstract, Sections 1, 2, 3, 4 and contributed to Conclusions. He conducted the performance analysis and performed the required AMIP experiments.

2. ZS wrote Section 5 and contributed to Conclusions. He performed the AMIP experiments required in Section 5 and analyzed the output.3. ED contributed to the AMIP experiment required for the performance analysis.

4. JB wrote the excerpt on the YAXT library and provided detailed advice on the performance analysis.

5. HB had the major role in the funding acquisition and supervision of the project leading to this publication. He also provided detailed advice on the direction of the paper and the performance analysis.

480 All co-authors contributed to the research and the text in this manuscript.

Competing interests. The authors declare no competing interests.

Acknowledgements.

1. The authors are grateful to the colleagues at the Max Planck Institute for Meteorology for sharing their results and ideas on implementing a similar approach in the ICON model.

485 2. The authors are also grateful to Prof. Dr. Thomas Ludwig and his Scientific Computing group at Universität Hamburg for their kind support.

3. The authors would like to thank the reviewers for their corrections and helpful comments.

Financial support. This research has been funded by the Federal Ministry of Education and Science (BMBF) under grant number FKZ:01LP1515A.

References

- 490 Balaji, V.: FMS: the GFDL Flexible Modeling System, Princeton, NJ, 2004.
 - Balaji, V., Benson, R., Wyman, B., and Held, I.: Coarse-grained component concurrency in Earth System modeling, 2016.
 - Bayr, T., Dommenget, D., Martin, T., and Power, S. B.: The eastward shift of the Walker Circulation in response to global warming and its relationship to ENSO variability, Climate dynamics, 43, 2747–2763, 2014.
 - Behrens, J., Hanke, M., and Jahns, T.: How to use MPI communication in highly parallel climate simulations more easily and more efficiently.,

495 EGUGA, p. 13569, 2014.

505

- Cess, R. D., Potter, G., Blanchet, J., Boer, G., Del Genio, A., Deque, M., Dymnikov, V., Galin, V., Gates, W., Ghan, S., et al.: Intercomparison and interpretation of climate feedback processes in 19 atmospheric general circulation models, Journal of Geophysical Research: Atmospheres, 95, 16 601–16 615, 1990.
- Cotronei, A. and Slawig, T.: Single precision arithmetic in ECHAM radiation reduces runtime and energy consumption, arXiv preprint arXiv:2001.01214, 2020.
 - Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., and Taylor, K. E.: Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization, Geoscientific Model Development, 9, 1937–1958, 2016.
 - Giorgetta, M. A., Jungclaus, J., Reick, C. H., Legutke, S., Bader, J., Böttinger, M., Brovkin, V., Crueger, T., Esch, M., Fieg, K., et al.: Climate and carbon cycle changes from 1850 to 2100 in MPI-ESM simulations for the Coupled Model Intercomparison Project phase 5, Journal of Advances in Modeling Earth Systems, 5, 572–597, 2013.
 - Guilyardi, E.: El Niño-mean state-seasonal cycle interactions in a multi-model ensemble, Climate Dynamics, 26, 329-348, 2006.
 - Guilyardi, E., Gualdi, S., Slingo, J., Navarra, A., Delécluse, P., Cole, J., Madec, G., Roberts, M., Latif, M., and Terray, L.: Representing El Nino in coupled ocean-atmosphere GCMs: the dominant role of the atmospheric component, Journal of Climate, 17, 4623–4629, 2004.

Guilyardi, E., Braconnot, P., Jin, F.-F., Kim, S. T., Kolasinski, M., Li, T., and Musat, I.: Atmosphere feedbacks during ENSO in a coupled

- GCM with a modified atmospheric convection scheme, Journal of Climate, 22, 5698–5718, 2009.
 Heidari, M. R., Song, Z., Degregori, E., Behrens, J., and Bockelmann, H.: Outputs of AMIP experiments using ECHAM-6.3.05p2 with the concurrent radiation scheme, https://doi.org/10.5281/zenodo.4589140, https://doi.org/10.5281/zenodo.4589140, 2021.
 - Hogan, R. J. and Bozzo, A.: Mitigating errors in surface temperature forecasts using approximate radiation updates, Journal of Advances in Modeling Earth Systems, 7, 836–853, 2015.
- 515 Hogan, R. J. and Bozzo, A.: A flexible and efficient radiation scheme for the ECMWF model, Journal of Advances in Modeling Earth Systems, 10, 1990–2008, 2018.
 - Hogan, R. J. and Hirahara, S.: Effect of solar zenith angle specification in models on mean shortwave fluxes and stratospheric temperatures, Geophysical Research Letters, 43, 482–488, 2016.

Jin, F.-F., Kim, S. T., and Bejarano, L.: A coupled-stability index for ENSO, Geophysical Research Letters, 33, 2006.

- 520 Kidston, J., Renwick, J., and McGregor, J.: Hemispheric-scale seasonality of the Southern Annular Mode and impacts on the climate of New Zealand, Journal of Climate, 22, 4759–4770, 2009.
 - Mauritsen, T. and Roeckner, E.: Tuning the MPI-ESM1. 2 global climate model to improve the match with instrumental record warming by lowering its climate sensitivity, Journal of advances in modeling earth systems, 12, e2019MS002 037, 2020.

Mauritsen, T., Bader, J., Becker, T., Behrens, J., Bittner, M., Brokopf, R., Brovkin, V., Claussen, M., Crueger, T., Esch, M., et al.: Devel-

- 525 opments in the MPI-M Earth System Model version 1.2 (MPI-ESM1. 2) and its response to increasing CO2, Journal of Advances in Modeling Earth Systems, 11, 998–1038, 2019.
 - McAfee, S. A. and Russell, J. L.: Northern annular mode impact on spring climate in the western United States, Geophysical Research Letters, 35, 2008.
- Morcrette, J.-J.: On the effects of the temporal and spatial sampling of radiation fields on the ECMWF forecasts and analyses, Monthly weather review, 128, 876–887, 2000.
- Morcrette, J.-J., Mozdzynski, G., and Leutbecher, M.: A reduced radiation grid for the ECMWF Integrated Forecasting System, Monthly

Mozdzynski, G. and Morcrette, J.-J.: Reorganization of the radiation transfer calculations in the ECMWF IFS, p. 20, https://doi.org/10.21957/pxjpl93ov, https://www.ecmwf.int/node/11300, 2014.

- 535 MPI-M and DKRZ: Source code of the Max Planck Institute atmospheric general circulation model (ECHAM6) adopted to the project PalMod, https://doi.org/10.35089/WDCC/SC_PalMod_ECHAM6, https://doi.org/10.35089/WDCC/SC_PalMod_ECHAM6, 2021.
 - Neelin, Zebiak, E. J. A. S. B. J. C. H. F. B. D. S. H. A. C. Z. S. E.: ENSO theory, Journal of Geophysical Research, 103, 14261–14290, 1998.

Pauluis, O. and Emanuel, K.: Numerical instability resulting from infrequent calculation of radiative heating, Monthly weather review, 132,

540 673–686, 2004.

Philander, S. G.: El Niño, La Niña, and the southern oscillation, International geophysics series, 46, X–289, 1989.

Pincus, R. and Stevens, B.: Paths to accuracy for radiation parameterizations in atmospheric models, Journal of Advances in Modeling Earth Systems, 5, 225–233, 2013.

Rasp, S.: Statistical methods and machine learning in weather and climate modeling, Ph.D. thesis, Imu, 2019.

545 Salby, M. L.: Fundamentals of atmospheric physics, Elsevier, 1996.

Stevens, B., Giorgetta, M., Esch, M., Mauritsen, T., Crueger, T., Rast, S., Salzmann, M., Schmidt, H., Bader, J., Block, K., et al.: Atmospheric component of the MPI-M Earth system model: ECHAM6, Journal of Advances in Modeling Earth Systems, 5, 146–172, 2013.
Wallace, J. M. and Hobbs, P. V.: Atmospheric science: an introductory survey, vol. 92, Elsevier, 2006.
Xu, K.-M. and Randall, D. A.: Impact of interactive radiative transfer on the macroscopic behavior of cumulus ensembles. Part I: Radiation

parameterization and sensitivity tests, Journal of the atmospheric sciences, 52, 785–799, 1995.

weather review, 136, 4760–4772, 2008.