Answer to 1st reviewer:

We would like to thank the reviewer for the work and his/her comments. We have the following remarks corresponding to the two main points he/she stated in the interactive comments:

We understood the first point the reviewer made in the following way: We did not investigate how to replace those subroutines or code parts that still require double or even quadruple precision (mentioned in Section 5) by alternative code. This is completely correct; it was due to the scope of the third-party-funded project this investigation was conducted in.

For the 2nd version of the manuscript, we now included some remarks that point this out, and mentioned it as an option for future work. This text was inserted at the beginning of Section 5 and in the conclusions.

Concerning the second point the reviewer made, we understood the following: We only validated the results of the whole model after changing the radiation to single precision. We never tested the two radiation versions (sp and dp) alone or in some kind of idealized configuration (e.g., in single atmospheric column configurations). This is also true, and we also agree that this would be a validation that would be much more exact.

The reason or answer is similar to above: Also this was due to the scope of the thirdparty-funded project this investigation was conducted in, which was very much result-oriented regarding simulation runs for coupled models.

We added remarks on this point directly in front of Subsection 6.1.1.

\_\_\_\_\_

Answer to 2nd reviewer:

We would like to thank the reviewer for the work and his/her comments. We have the following remarks corresponding to the points stated in section 1 of the interactive comments:

Reviewer's comment: Section 3.3: After reading this section several times, I still don't understand the conversion process. For example:

- line 167: "If the model output was acceptable" — how do you define "acceptable"? And what do you do if the model output is not acceptable? Do you then delete the sp version?

Answer: Here we changed the text to the following:

We then tested if the model with the sp version of the subroutine/function compiles, produces no runtime errors, and if its difference to the dp version was in an

``acceptable" range. Of course, the latter is a soft criterion, since a bit-identical result cannot be expected. Our criteria are explained below in Subsection 6.1. If the sp output was not acceptable in this sense, we marked the corresponding code part as to be treated separately, as described below in Sections 4 and 5.

Reviewer's comment:

line 167: "low\_dp, as well as the interface were redundant" — even if the model output when using low\_sp is acceptable, you still need an interface "on top of" low\_sp that allows double-precision arguments. This is the purpose of low\_dp. If you delete low\_dp, how does high\_dp call low\_sp?
It could be that the authors won't understand my questions because my thinking is so wrong. In any case, I didn't understand it and I recommend that the authors rewrite this section so the procedure is clearer, perhaps including some diagrams.

## Answer:

This was formulated in a misleading way in our first version. The two versions low\_sp and low\_dp can be deleted only at the end, when a new "wp" version of "low" is introduced, wp being a variable that could be set to "sp" or "dp". Then, also the interface becomes redundant.

Reviewer's comment: line 182: "namely -real-size 32": this depends on the compiler. I'm not even sure

GNU Fortran has an option to set the default REAL precision to 4 bytes, as this is already the language standard (as far as I'm aware). I don't think you need this paragraph at all — you can simply say that you added type declarations to all REAL variables and literals so that the type was always explicit. This is good programming practice anyway.

Answer: Paragraph was omitted.

Reviewer's comment: lines 297 - 300: I didn't understand this paragraph. For example, "we also

took a look at the minimum and maximum over all grid points" — minimum and maximum error, or minimum and maximum field values? I'm assuming the latter. If so, why does a difference in minimum and maximum indicate a bias? If single-precision has both a larger maximum and smaller minimum than doubleprecision, the mean could still be zero (meaning zero bias). I recommend either

rewriting this paragraph or just deleting it.

Answer: We removed the paragraph.

Reviewer's comment: Section 6.1.2: The equation only computes an annual mean if #months in time span = 12. In fact the period is 30 years so I think you mean "temporal mean" not "annual mean". The caption of Figure 6 even uses that name.

Answer: We added an explaining sentence after the formula.

Reviewer's comment: Figures 6, 7, 8: This could simply be GMD's formatting but I

can't read the colorbars in these figures. If the authors deliberately chose this size for the Figures,

please enlarge the colorbars.

Answer: We enlarged the whole figures to make all detailed better visible.

Reviewer's comment: line 359: so all of the shown results are for "block"? Please clarify this.

Answer: We added this at (new) line 360 (old line 359) and also before in (new) line 328.

Reviewer's comment: The terms "performance gain", "runtime reduction", "speed-up" and "acceleration"

are used interchangeably throughout the manuscript (mainly the first three) but it's not clear what they actually mean. If x and y are the wall-clock times for single and double, respectively, is the performance gain (or whatever) 1 - x/y or y/x? I recommend using the phrase "runtime reduction", meaning 1 - x/y, as much as possible, as this is what others like Vána et al. use. "Speed-up" sounds like y/xto me i.e. if x = 5 seconds and y = 10 seconds, the speed-up is 2 because single is twice as fast.

Answer: We followed the suggestion to use "(relative) runtime reduction" all the time and defined the term as 1-x/y = (dp-sp)/dp, at the beginning of subsection 6.2.

Answer to reviewer's section 2: All mentioned typos were corrected.

# Single precision arithmetic in ECHAM radiation reduces runtime and energy consumption

Alessandro Cotronei<sup>1</sup> and Thomas Slawig<sup>1</sup>

<sup>1</sup>Kiel Marine Science (KMS) – Centre for Interdisciplinary Marine Science, Dep. of Computer Science, Kiel University, 24098 Kiel, Germany.

Correspondence: Thomas Slawig, ts@informatik.uni-kiel.de

**Abstract.** We converted the radiation part of the atmospheric model ECHAM to single precision arithmetic. We analyzed different conversion strategies and finally used a step by step change of all modules, subroutines and functions. We found out that a small code portion still requires higher precision arithmetic. We generated code that can be easily changed from double to single precision and vice versa, basically using a simple switch in one module. We compared the output of the single

5 precision version in the coarse resolution with observational data and with the original double precision code. The results of both versions are comparable. We extensively tested different parallelization options with respect to the possible performance gainruntime reduction, in both coarse and low resolution. The single precision radiation itself was accelerated by about 40%, whereas the speed-up-runtime reduction for the whole ECHAM model using the converted radiation achieved 18% in the best configuration. We further measured the energy consumption, which could also be reduced.

#### 10 1 Introduction

The atmospheric model ECHAM was developed at the Max Planck Institute for Meteorology (MPI-M) in Hamburg. Its development started in 1987 as a branch of a global weather forecast model of the European Centre for Medium-Range Weather Forecasts (ECMWF), thus leading to the acronym (EC for ECMWF, HAM for Hamburg). The model is used in different Earth System Models (ESMs) as atmospheric component, e.g., in the MPI-ESM also developed at the MPI-M, see Figure 1. The current version is ECHAM 6 (Stevens et al. 2013). For a detailed list on ECHAM publications we refer to the homepage of

15 current version is ECHAM 6 (Stevens et al., 2013). For a detailed list on ECHAM publications we refer to the homepage of the institute (mpimet.mpg.de). Version 5 of the model was used in the 4th Assessment Report of the Intergovernmental Panel on Climate Change (IPCC, 2007), version 6 in the Coupled Model Intercomparison Project CMIP (World Climate Research Programme, 2019a).

Motivation for our work was the usage of ECHAM for long-time paleo-climate simulations in the German national climate

20 modeling initiative "PalMod: From the Last Interglacial to the Anthropocene – Modeling a Complete Glacial Cycle" (www. palmod.de). The aim of this initiative is to perform simulations for a complete glacial cycle, i.e. about 120'000 years, with fully coupled ESMs.



Figure 1. Schematic of the structure of the Earth System Model MPI-ESM with atmospheric component ECHAM, terrestrial vegetation model JSBACH, ocean model MPI-OM, marine biogeochemical model HAMOCC, and OASIS coupler.



Figure 2. Time consumption of the radiation part w.r.t. the whole ECHAM model in coarse (CR) and low resolution (LR) of standard PalMod experiments. The difference occurs since in these configuration the radiation part is called every 2 hours, i.e., only every fourth (in CR) or every eighth time-step (LR).

The feasibility of long-time simulation runs highly depends on the computational performance of the used models. As a consequence, one main focus in the PalMod project is to decrease the runtime of the model components and the coupled ESMs.

25

In ESMs that use ECHAM, the part of the computational time that is used by the latter is significant. It can be close to 75% in some configurations. Within ECHAM itself, the radiation takes the most important part of the computational time. As a consequence, the radiation part is not called in every time-step in the current ECHAM setting. Still, its part of the overall ECHAM runtime is relevant, see Figure 2.

30

In the PalMod project, two different strategies to improve the performance of the radiation part are investigated: One is to run the radiation in parallel on different processors, the other one is the conversion to single precision arithmetic we present in this paper. For both purposes, the radiation code was isolated from the rest of the ECHAM model. This technical procedure is not described here.

The motivation for the idea to improve the computational performance of ECHAM by a conversion to reduced arithmetic

35

5 precision was the work of Vana et al. (2017). In this paper, the authors report on the conversion of the Integrated Forecasting System (IFS) model to single precision, observing a runtime reduction of 40% in short-time runs of 12 or 13 months, and a good match of the output with observational data. Here, the terms *single* and *double precision* refer to the IEEE-754 standard for floating point numbers (IEEE Standards Association, 2019), therein also named *binary64* and *binary32*, respectively. In the IEEE standard, also an even more reduced precision, the *half precision (binary16)* format is defined. The IFS model, developed

40 also by the ECMWF, is comparable to ECHAM in some respect since it also uses a combination of spectral and grid-pointbased discretization. A similar performance gain runtime reduction of 40% was reported by Rüdisühli et al. (2014) with the COSMO model that is used by the German and Swiss weather forecast services. The authors also validated the model output by comparing it to observations and the original model version.

Recently, the usage of reduced precision arithmetic has gained interest for a variety of reasons. Besides the effect on the runtime, also a reduction of energy consumption is mentioned, see, e.g., Fagan et al. (2016), who reported a reduction by about 60%. In the growing field of machine learning, single or even more reduced precision is used to save both computational effort as well as memory, motivated by the usage of Graphical processing Units (GPUs). Dawson and Düben (2017) used reduced precision to evaluate model output uncertainty. For this purpose, the authors developed a software where a variable precision is available, but a positive effect on the model runtime was not their concern.

- The process of porting a simulation code to a different precision highly depends on the design of the code and the way how basic principles of software engineering have been followed during the implementation process. These are modularity, use of clear subroutine interfaces, way of data transfer via parameter list or global variables etc. The main problem in legacy codes with a long history (as ECHAM) is that these principles usually were not applied very strictly. This is a general problem in computational science software, not only in climate modeling, see for example Johanson and Hasselbring (2018).
- 55 Besides the desired performance gain runtime reduction, a main criterion to assess the result of the conversion to reduced precision is the validation of the results, i.e., their differences to observational data and the output of the original, double precision precision version. We carried out experiments on short time scales of 30 years with a 10 years spin-up. It has to be taken into account that after the conversion, a model tuning process (in fully coupled version) might be necessary. This will require a significant amount of work to obtain an ESM that produces a reasonable climate, see, e.g., Mauritsen et al. (2012) for
- 60 a description of the tuning of the MPI-ESM.

The structure of the paper is as follows: In the following section, we describe the situation from where our study and conversion started. In Section 3, we give an overview about possible strategies to perform a conversion to single precision, discuss their applicability, and finally the motivation for the direction we took. In Section 4, we describe changes that were necessary at some parts of the code due to certain used constructs or libraries, and in Section 5 the parts of the code that need

65 to remain in higher precision. In Section 6, we present the obtained results w.r.t. performance gainruntime reduction, output validation and energy consumption. At the end of the paper, we summarize our work and draw some conclusions.

#### 2 Used configuration of ECHAM

75

The current major version of ECHAM, version 6, is described in Stevens et al. (2013). ECHAM is a combination of a spectral and a grid-based finite difference model. It can be used in five resolutions, ranging from the *coarse resolution* (CR) or T31

70 (i.e., a truncation to 31 wave numbers in the spectral part, corresponding to a horizontal spatial resolution of 96 - × 48 points in longitude and latitude) up to XR or T255. We present results for the CR and LR (*low resolution*, T63, corresponding to 192 - × 96 points) versions. Both use 47 vertical layers and (in our setting) time steps of 30 and 15 min., respectively.

ECHAM6 is written in free format Fortran and conforms to the Fortran 95 standard (Metcalf et al., 2018). It consists of about 240'000 lines of code (including approximately 71'000 lines of the JSBACH vegetation model) and uses a number of external libraries including LAPACK, BLAS (for linear algebra), MPI (for parallelization), and NetCDF (for in- and output).

- The radiation part that we converted contains about 30'000 lines of code and uses external libraries as well. The basis ECHAM version we used is derived from the stand-alone version ECHAM-6.03.04p1. In this basis version, the radiation was modularly separated from the rest of ECHAM. This offers the option to run the radiation and the remaining part
- 80 running the ECHAM components sequentially. It was shown that the sequential version reproduces bit-identical results with the original code.

All the results presented below are evaluated with the Intel Fortran compiler 18.0 (update 4) (Intel, 2017) on the supercomputer HLRE-3 *Mistral*, located at the German Climate Computing Center (DKRZ), Hamburg. All experiments used the so-called *compute* nodes of the machine.

of the model on different processors in order to reduce the running time by parallelization, but also maintains the possibility of

#### 85 **3** Strategies for conversion to single precision

In this section we give an overview of possible strategies for the conversion of a simulation code (as the radiation part of ECHAM) to single precision arithmetic. We describe the problems that occurred while applying them to the ECHAM radiation part. At the end, we describe the strategy that finally turned out to be successful. The general target was a version that can be used in both single and double precision with as few changes to the source code as possible. Our goal was to achieve a general

90 setting of the working precision for all floating point variables at one location in one Fortran module. It has to be taken into account that some parts of the code might require double precision. This fact was already noticed in the report on conversion of the IFS model by Vana et al. (2017).

We will from now on refer to the single precision version as sp, and to the double precision version as dp version.

#### 3.1 Use of a precision switch

95 One ideal and elegant way to switch easily between different precisions of the variables of a code in Fortran is to use a specification of the kind parameter for floating point variables as showed in the following example. For reasons of flexibility, the objective of our work was to have a radiation with such precision switch.

! define variable with prescribed working precision (wp): real(kind = wp) :: x

100 The actual value of wp can then be easily switched in the following way:

105

! define	different	NOI	cking	precision	ns:				
integer,	parameter	::	sp =	4	!	single	precision	(4	byte)
integer,	parameter	::	dp =	8	!	double	precision	(8	byte)
! set wo	rking preci	sic	on:						
integer,	parameter	::	wp =	sp					

The recommendation mentioned by Metcalf et al. (2018, Section 2.6.2) is to define the different values of the kind parameter by using the selected\_real\_kind function. It sets the actually used precision via the definition of the desired number of significant decimals (i.e., mantissa length) and exponent range, depending on the options the machine and compiler offer. This reads as follows:

In fact, similar settings can be found in the ECHAM module rk\_mo\_kind, but unfortunately they are not consequently used. Instead, kind = dp is used directly in several modules. A somehow dirty workaround, namely assigning the value 4 to dp and declaring an additional precision for actual *dp* where needed, circumvents this problem. Then, compilation was possible after some modifications (concerning MPI and NetCDF libraries and the module mo\_echam\_radkernel\_cross\_messages). The compiled code was crashing at runtime because of internal bugs triggered by code in the module rk\_mo\_srtm\_solver and other parts. These issues were solved later when investigating each code part with the incremental conversion method. The cause of these bugs could not be easily tracked.

#### 125 3.2 Source code conversion of most time-consuming subroutines

As mentioned above, the conversion of the whole ECHAM model code using a simple switch was not successful. Thus, we started to identify the most time-consuming subroutines and functions and converted them by hand. This required the conversion of input and output variables in the beginning and at the end of the respective subroutines and functions. The changes in the code are schematically depicted in Figure 3.



Figure 3. Necessary code changes to convert a subroutine/function from the original double precision version (left) to single precision (right) with internal casting, modifications in red.



Figure 4. Time consumption of single and double precision taumol subroutines.

140

130 This procedure allowed an effective implementation of *sp* computations of the converted subroutines/functions. We obtained high <u>performance gain runtime reduction</u> in some code parts. But, the casting overhead destroyed the overall performance, especially if there are many variables to be converted.

For example, a time-consuming part of the subroutine gas\_optics\_lw in the module mo\_lrtm\_gas\_optics was converted in the above way. The converted part contains calls to subroutines taumol01 to taumol16, which were converted

135 to *sp*. Figure 4 shows the speed-up runtime reduction for these subroutines, which was up to 30%. But the needed casting in the calling subroutine gas\_optics\_lw doubled the overall runtime in *sp*, compared to *dp*.

The results of this evaluation lead to the following conclusion, which is not very surprising: The bigger the converted code block is with respect to the number of input and output variables, the lower the overhead for the casting will be in comparison to the gain in the calculations that are actually performed in *sp*. This was the reason for the decision to convert the whole radiation part of ECHAM, as it contains a relatively small amount of input/output variables.

#### 3.3 Incremental conversion of the radiation part

As a result of the not efficient inefficient conversion of the most time-consuming subroutines or functions only, we performed a gradual conversion of the whole radiation code. For this purpose, we started from the lowest level of its calling tree, treating each subroutine/function separately. Consider an original subroutine on a lower level,

145

155

```
subroutine low(x_dp)
real(dp) :: x dp
```

using *dp* variables. We renamed it as low\_dp and made a copy in *sp*:

```
subroutine low_sp(x_sp)
real(sp) :: x_sp
```

150 We changed the dp version such that it just calls its sp counterpart, using implicit type conversions before and after the call:

```
subroutine low_dp(x_dp)
real(dp) :: x_dp
real(sp) :: x_sp
x_sp = x_dp
call low_sp(x_sp)
x dp = x sp
```

Now we repeated the same procedure with each subroutine/function that calls the original low, e.g.,

```
subroutine high(...)
call low(x_dp)
```

160 We again renamed it as high\_dp, generated an sp copy high\_sp, and defined an interface block (Metcalf et al., 2018):

```
interface low
  module procedure low_sp
  module procedure low_dp
end interface
```

165 In both high\_dp and high\_sp, we could now call the respective version of the lower level subroutine passing either *sp* or *dp* parameters. The use of the interface simplified this procedure significantly.

If the model output was acceptable, the We then tested if the model with the *dpsp* version on the lower level, low\_dp, as well as the interface were redundant, and we deleted them. Finally, the of the subroutine/function compiles, produces no runtime errors, and if its difference to the *dp* version was in an "acceptable" range. Of course, the latter is a soft criterion, since

170 a bit-identical result cannot be expected. Our criteria are explained below in Subsection 6.1. If the sp version could be renamed

as low output was not acceptable in this sense, we marked the corresponding code part as to be treated separately, as described below in Sections 4 and 5.

This procedure was repeated up to the highest level of the calling tree. It required a lot of manual work, but it allowed the examination of each modified part of the code, as well as a validation of the output data of the whole model.

In the ideal case and if no additional code changes would be necessary for the *sp* version, this would have led <del>directly to a</del> to consistent *sp* version. Replacing then the and *dp* versions. Finally, one of them could be renamed as *splowby*, using wp in the code for a working precision that could be set to *sp* or *dp* in some central module, we would have ended up with. Then, also the interface becomes redundant. This would be a model version that has a precision switch. The next section summarizes the few parts of the code that needed extra treatment.

#### 180 4 Necessary changes in the radiation code

Changing the floating point precision in the radiation code required some modifications that are described in this section. Some of them are related to the use of external libraries, some others to an explicitly used precision-dependent implementation.

#### 4.1 Procedure

In the incremental conversion, the precision variables dp and wp that are both used in the radiation code were replaced with sp.

185 Then in the final version, sp was replaced by wp. With this modification, wp became a switch for the radiation precision. As the original radiation contained several variables lacking explicit declaration of their precision, the respective format specifiers were added throughout.

A Fortran compiler option (namely <u>real-size 32</u>) can avoid the last step by assigning *sp* to variables without format specifier. However, since the original model used both *sp* and *dp*, this option could lead to an overhead due to conversion in the

190 dp part. Although a compilation of the program with custom option was possible for the *sp* Fortran modules, this procedure led to a more complicated compilation procedure and was therefore discarded. The format specifier D0, denoting dp, was also replaced by \_wp.

#### 4.2 Changes needed for the use of the NetCDF library

In the NetCDF library, the names of subroutines and functions have different suffixes depending on the used precision. They are used in the modules

```
rk_mo_netcdf, rk_mo_cloud_optics, rk_mo_lrtm_netcdf,
rk_mo_o3clim, rk_mo_read_netcdf77, rk_mo_srtm_netcdf.
```

In *sp*, they have to be replaced by their respective counterparts to read the NetCDF data with the correct precision. The script shown in Appendix A performs these changes automatically. This solution was necessary because an implementation using an

200 interface was causing crashes for unknown reasons. It is possible that further investigation could lead to a working interface implementation for these subroutines/functions also at this point of the code.

#### 4.3 Changes needed for parallelization with MPI

Several interfaces of the module mo\_mpi were adapted to support *sp*. In particular p\_send\_real, p\_recv\_real and p\_bcast\_real were overloaded with *sp* subroutines for the needed array sizes. These modifications did not affect the calls to these interfaces. No conversions are made in this module, so no overhead is generated.

205

#### 4.4 Changes needed due to data transfer to the remaining ECHAM

In ECHAM, data communication between the radiation part and the remaining atmosphere is implemented in the module

mo\_echam\_radkernel\_cross\_messages

through subroutines using both MPI and the coupling library YAXT (DKRZ, 2013). Since it was not possible to have a mixed precision data transfer for both libraries, our solution was to double the affected subroutines to copy and send both *sp* and *dp* data. An additional variable conversion before or after their calls preserves the needed precision. Also in this case, interface blocks were used to operate with the correct precision. The changed subroutines have the following prefixes: copy\_echam\_2\_kernel, copy\_kernel\_2\_echam, send\_atm2rad and send\_rad2atm. These modifications only affect the ECHAM model when used in the parallel scheme. They have a negligible overhead.

#### 215 5 Parts still requiring higher precision

In the *sp* implementation of the radiation code, some parts are still requiring higher precision to run correctly. These parts and the reasons are presented in this section. We want to emphasize that it is desirable to determine these reasons in more detail, and to find alternatives in single precision arithmetic that still give reasonable model output. However, this was beyound the scope of the project in which our work was conducted.

#### 220 5.1 Overflow avoidance

When passing from dp to sp variables, the maximum representable number decreases from  $\approx 10^{308}$  to  $\approx 10^{38}$ . In order to avoid overflow that could lead to crashes, it is necessary to adapt the code to new thresholds. A similar problem could potentially occur for numbers which are too small (smaller than  $\approx 10^{-45}$ ).

As stated in the comments in the original code of psrad\_interface, the following exponential needs conversion if not used in *dp*:

!this is ONLY o.k. as long as wp equals dp, else conversion needed cisccp\_cldemi3d(jl,jk,krow) = 1.\_wp - exp(-1.\_wp\*cld\_tau\_lw\_vr(jl,jkb,6)) One plausible reason for this is that the exponential is too big for the range of sp. Even though this line was not executed in the used configuration, we converted the involved quantities to dp. Since the variable on the left-hand side of the assignment was

230 transferred within few steps to code parts outside the radiation, no other code inside the radiation had to be converted into dp. Also module rk\_mo\_srtm\_solver contained several parts sensitive to the precision. First of all, the following lines containing the hard coded constant 500 could cause overflow as well:

exp_minus_tau_over_m	= (:) Ou	inv_expon(MIN(tau_over_mu,	500wp),	kproma)
exp_ktau	(:) =	expon(MIN(k_tau,	500wp),	kproma)

235 Here, expon and inv\_expon calculate the exponential and inverse exponential of a vector (of length kproma in this case). The (inverse) exponential of a number close to 500 is too big (small) to be represented in *sp*. In the used configurations, this line was not executed either. Nevertheless, we replaced this value by a constant depending on the used precision, see the script in Appendix A.

#### 5.2 Numerical stability

- 240 Subroutine srtm\_reftra\_ec of module rk\_mo\_srtm\_solver, described in Meador and Weaver (1979), showed to be very sensitive to the precision conversion. In this subroutine, already a conversion to *sp* of just one of most internal variables separately was causing crashes. We introduced wrapper code for this subroutine to maintain the *dp* version. The time necessary for this overhead was in the range of 3.5 to 6 % for the complete radiation and between 0.6 and 3% for the complete ECHAM model.
- **245** In subroutine Set\_JulianDay of the module rk\_mo\_time\_base, the use of *sp* for the variable zd, defined by

zd = FLOOR(365.25\_dp\*iy)+INT(30.6001\_dp\*(im+1)) &
+ REAL(ib,dp)+1720996.5\_dp+REAL(kd,dp)+zsec

caused crashes at the beginning of some simulation years. In this case, the relative difference between the sp and the dp representation of the variable zd is close to machine precision (in sp arithmetic), i.e., the relative difference attains its maximum value. This indicates that code parts that use this variable afterwards are very sensitive to small changes in input data. The code block was kept in dp by reusing existing typecasts, without adding new ones. Thus, this did not increase the runtime. Rewriting the code inside the subroutine might improve the stability and avoid the typecasts completely.

#### 5.3 Quadruple precision

The module rk\_mo\_time\_base also contains some parts in quadruple (REAL(16)) precision in the subroutine 255 Set\_JulianCalendar, e.g.:

zb = FLOOR(0.273790700698850764E-04\_wp\*za-51.1226445443780502865715\_wp)

Here wp was set to REAL (16) in the original code. This high precision was needed to prevent roundoff errors because of the number of digits in the used constants. We did not change the precision in this subroutine. But since we used wp as indicator

for the actual working precision, we replaced wp by ap (advanced precision) to avoid conflicts with the working precision in

this subroutine. We did not need to implement any precision conversion, since all input and output variables are converted from and to integer numbers inside the subroutine anyway.

#### 6 Results

In this section, we present the results obtained with the *sp* version of the radiation part of ECHAM. We show three types of results, namely a comparison of the model output, the obtained gain in runtime and finally the gain in energy consumption.

The results presented below were obtained with the AMIP experiment (World Climate Research Programme, 2019b) by using the coarse (CR, T031L47) or low (LR, T063L47) resolutions of ECHAM. The model was configured with the cdi-pio parallel input-output option (Kleberg et al., 2017). We used the following compiler flags (Intel, 2017), which are the default ones for ECHAM:

- -03: enables aggressive optimization,

- 270 fast-transcendentals, -no-prec-sqrt, -no-prec-div: enable faster but less precise transcendental functions, square roots, and divisions,
  - -fp-model source: rounds intermediate results to source-defined precision,
  - -xHOST: generates instructions for the highest instruction set available on the compilation host processor,
  - -diag-disable 15018: disables diagnostic messages,
- 275 -assume realloc\_lhs: uses different rules (instead of those of Fortran 2003) to interpret assignments.

#### 6.1 Validation of model output

To estimate the output quality of the sp version, we compared its results with

- the results of the original, i.e., the dp version of the model
- and observational data available from several datasets.
- 280 We computed the difference between the outputs of the *sp* and *dp* versions and the differences of both versions to the observational data. We compared the values of
  - temperature (at the surface and at 2m height), using the CRU TS4.03 dataset (University of East Anglia Climatic Research Unit et al., 2019),
  - precipitation (sum of large scale and convective in ECHAM), using the GPCP data provided by the NOAA/OAR/ESRL PSD, Boulder, Colorado, USA (Adler et al., 2003).

285

- cloud radiative effect (CRE at the surface and at the top of the atmosphere, the latter split into longwave and shortwave parts), using the CERES EBAF datasets release 4.0 (Loeb and National Center for Atmospheric Research Staff, 2018).

In all results presented below, we use the monthly mean of these variables as basic data. This is motivated by the fact that we are interested in long time simulation runs and climate prediction rather than in short-term scenarios (as for weather prediction). Monthly means are directly available as output from ECHAM.

290

All computations have been performed with the use of the Climate Data Operators (CDO) (Schulzweida, 2019).

We emphasize that we present results for the output of the whole ECHAM model only. It would be also valuable to investigate the differences in the output of the two versions (*dp* and *sp*) of the radiation code alone, ideally when just using them for single atmospheric columns. This investigation was beyound the scope of our work here.

#### 295 6.1.1 Difference in RMSE between single and double version and observational data

We computed the spatial root mean square error (RMSE) of the monthly means for both sp and dp versions and the above variables. We applied the same metric for the difference between the outputs of the sp and dp versions. We computed these values over time intervals where observational data were available in the used datasets. For temperature and precipitation, these were the years 1981-2010, for CRE the years 2000-2010 or 2001-2010. In all cases, we started the computation in the year 1070, he is a specific term of the specific term of the specific term.

300 1970, having a reasonable time interval as spin-up.

Figure 5 shows the temporal behavior of the RMSE and the differences of *sp* and *dp* version, as they evolve in time. It can be seen that the RMSEs of the *sp* version are in the same magnitude as those of the *dp* version. Also the differences between both versions are of similar or even smaller magnitude. Moreover, all RMSEs and differences do not grow in time. They oscillate but stay in the same order of magnitude for the whole considered time intervals.

305

Additionally, we averaged these values over the respective time intervals. Table 1 again shows that the RMSEs of the sp version are in the same magnitude as those of the dp version. Also the differences between both versions are of similar or smaller magnitude.

Since the RMSE averages spatial differences, we also took a look at the minimum and maximum over all grid-points. This is a check whether the conversion introduced arbitrary biases for the considered variables over the runtime. This method allows

310 only a rough test, since a model producing unreliable data could anyway produce low averages, minima or maxima, even more if they are averaged annually. This comparison showed a similar magnitude as the differences obtained using two subsequent ECHAM versions. Thus, we do not show corresponding plots here.

Moreover, we compared our obtained differences with the ones between two runs of the ECHAM versions 6.3.02 and 6.3.02p1. The differences between *sp* and *dp* version are in the same magnitude as the differences between these two model versions.



**Figure 5.** Spatial RMSE of monthly means for *sp* and *dp* versions and differences between them in the same metric, for (from top left to bottom right) temperature at the surface and at 2m, precipitation, total CRE at the surface, longwave and shortwave part of CRE at top of the atmosphere.

#### 6.1.2 Spatial distribution of differences in the annual means

We also studied the spatial distribution of the differences in the annual means. Again we considered the differences between *sp* and *dp* version and of the output of both versions to the observations. Here we included the signs of the differences and no absolute values or squares. For the given time spans, this results in a function of the form

320 DIFF(grid-point) := 
$$\frac{1}{\#$$
months in time-span} \sum\_{\text{months in time-span}} (y(grid-point, month) - z(grid-point, month))

 Table 1. Spatial RMSE of monthly means for sp and dp versions and difference of both versions in the same metric, for selected model variables, averaged over the respective time spans (obs = observational data).

	Variable	ECHAM variable	Unit	Time span	sp - dp	dp - obs
	Surface temperature	169	К	1981 - 2010	2.0302	2.2862
	Temperature at 2m	167	К	1981 - 2010	2.0871	2.0585
	Precipitation	142+143	$kgm^{-2}s^{-1}$	1981 - 2010	$1.9281 \cdot 10^{-5} 1.9281 \times 10^{-5}$	$\begin{array}{ c c c c c c c c c c c c c c c c c c c$
	CRE, surface	176-185+177-186	$Wm^{-2}$	2000 - 2010	12.3661	23.4345
S	hortwave CRE, top of atmosphere	178-187	$Wm^{-2}$	2001 - 2010	14.3859	31.5827
L	ongwave CRE, top of atmosphere	179-188	$Wm^{-2}$	2001 - 2010	9.3010	13.2364

Spatial RMSE of monthly means for *sp* and *dp* versions and difference of both versions in the same metric, for selected model variables, averaged over the respective time spans (obs = observational data).

for two variables or datasets y, z of monthly data. We performed this evaluation with #months in time span set to 12 for every year in the considered interval of 30 years. This procedure can be used to see if some spatial points or areas are constantly warmer or colder over longer time ranges. It is also a first test of the model output. However, it is clearly not sufficient for validation because errors may cancel out over time. The results are shown in Figures  $\frac{??}{100}$  to  $\frac{11}{100}$ .

Additionally, we performed a statistic analysis of the annual means of the *sp* version. We checked the hypothesis that the 30years mean (in the interval 1981-2010) of the *sp* version equals the one of the original *dp* version. For this purpose, we applied a two-sided t-test, using a consistent estimator for the variance of the annual means of the *sp* version. The corresponding values are showed in the respective second rows in Figures ?? to ??6 to 11. In this test, absolute values below 2.05 are not significant at the 95 confidence level. For all considered variables, it can be seen that only very small spatial regions show higher values.

#### 330 6.2 Speed-upRuntime reduction

In this section we present the results of the obtained speed-up runtime reduction when using the modified *sp* radiation code in ECHAM. All presented values are *relative* runtime reductions, computed by the formula

 $\underbrace{\text{runtime reduction}}_{\text{runtime } dp \text{ version}} - \underbrace{\text{runtime } sp \text{ version}}_{\text{runtime } dp \text{ version}}$ 

335

Since the model is usually run on parallel hardware, there are several configuration options that might affect its performance and also the speed-up-runtime reduction when using the *sp* instead of the *dp* radiation code. We used the Mistral HPC system at DKRZ with 1 to 25 nodes, each of which has two Intel<sup>®</sup> Xeon<sup>®</sup> E5-2680v3 12C 2.5GHz ("Haswell") with 12 cores, i.e., using from 24 up to 600 cores. The options we investigated are:

- The number of used nodes.



Figure 6. Differences in temporal mean over the time interval 1981 to 2010 in temperature at surface (left) and at 2m height (right) in Kelvin. Top: difference between *sp* and *dp* version, second row: values of two-sided t-test with respect to variance of the annual *sp* output, absolute values below 2.05 are not significant at the 95 confidence level, third row: between *dp* version and observational data, bottom: between *sp* version and observational data.

- 340
- The choices *cyclic:block* and *cyclic:cyclic* (in this paper simply referred to as *block* and *cyclic*) offered by the *SLURM* batch system (SchedMD<sup>®</sup>, 2019) used on Mistral. It controls the distribution of processes across nodes and sockets inside a node. If not specified otherwise, we refer to the *block* setting, which is the default in ECHAM.



Figure 7. As Figure ?? 6, but for precipitation (left) in kgm<sup>-2</sup>s<sup>-1</sup> and surface cloud radiation effect in Wm<sup>-2</sup>temperature at 2m. For the latter difference to observations over 2000-2010.

Different values of the ECHAM parameter *nproma*, the maximum block length used for vectorization. For a detailed description see Rast (2014, Section 3.8).

We were interested in the best possible speed-up-runtime reduction when using the *sp* radiation in ECHAM. We studied the performance gain-runtime reduction achieved both for the radiation itself and for the whole ECHAM model for a variety of



Figure 8. As Figure ??.6, but longwave (left) and shortwave cloud radiation effect at the top of the atmosphere for precipitation in  $Wmkgm^{-2}s^{-1}$ . Difference to observations over 2001-2010 for both.

different settings of the mentioned options, for both CR and (with reduced variety) LR resolutions. Our focus lies on the CR version, since it is the configuration that is used in the long-time paleo runs intended in the PalMod project.

The results presented in this section have been generated with the *Scalable Performance Measurement Infrastructure for Parallel Codes* (Score-P, 2019) and the internal ECHAM timer.



Figure 9. As Figure 6, but for surface cloud radiation effect in  $Wm^{-2}$ . Here, the differences to observations are for 2000-2010.

All time measurements are based on one-year runs. The unit we use to present the results is the number of simulated years per day runtime. It can be computed by the time measurements of the one-year runs. For the results for the radiation part only, these are theoretical numbers, since the radiation is not run stand-alone for one year in ECHAM. We include them to give an impression what might be possible when more parts of ECHAM or even the whole model would be converted to *sp*. Moreover, we wanted to see if the speed-up-runtime reduction of 40% achieved with IFS model in Vana et al. (2017) could be reached.



Figure 10. As Figure 9, but for longwave cloud radiation effect at the top of the atmosphere.

355

To figure out if there are significant deviations in the runtime, we also applied a statistic analysis for 100 one-year runs. They showed that there are only very small relative deviations from the mean, see Table 2.

At the end of this section, we give some details which parts of the radiation code benefit the most from the conversion to reduced precision, and which ones not.



Figure 11. As Figure 10, but for shortwave cloud radiation effect at the top of the atmosphere.

### 6.2.1 Dependency of runtime and speed-up-runtime reduction on parameter settings

360 In order to find out the best possible speed-up highest possible runtime reduction when using the *sp* radiation code, we first analyzed the dependency of the runtime on the parameter *nproma*. For the CR resolution, we tested for 1 to 25 cores *nproma* values from 4 to 256 in steps of 4. It can be seen in the two top left pictures in Figure 12 that for 24 nodes there is no big

#### Table 2. Relative standard deviations of runtime over 100 runs.

	24 nodes block block	24 nodes eyelie eyelie	1 node <del>block <u>block</u></del>	1 node <del>cyclic cyclic</del>
Radiation dp	0.0095	0.0104	0.0081	0.0075
sp	0.0132	0.0122	0.0079	0.0084
ECHAM dp	0.0220	0.0179	0.0030	0.0023
sp	0.0158	0.0189	0.0027	0.0020



**Figure 12.** Comparison of simulated model years per day cputime for *sp* and *dp* versions in coarse resolution (CR), for radiation part (top) and whole ECHAM (bottom), 1 and 24 nodes and values of *nproma* between 4 and 256, in steps of 4.

dependency on *nproma* for the original *dp* version, when looking at radiation only. For the *sp* version, the dependency is slightly bigger, which results in a variety of the achieved speed-up runtime reduction between 25 to 35%.

365

When looking at the results for the whole ECHAM model on the two left pictures below in in-Figure 12, it can be seen that the dependency of the speed-up-runtime reduction on *nproma* becomes more significant.

Using only one node the performance for the dp version decreases with higher *nproma*, whereas the *sp* version does not show that big dependency. The effect is stronger when looking at the radiation time only than for the whole ECHAM. For very small values of *nproma*, the *sp* version was even slower than the dp version. In particular, the default parameter value (*nproma* = 12)



Figure 13. As Figure 12, but for low resolution (LR) and values of *nproma* between 8 and 512, in steps of 8.

for the sp version resulted in slower execution time than the corresponding dp version. An increased value of the parameter 370 (nproma = 48) made sp faster, even compared to the fastest nproma for dp (which was 24).

The difference between the *block* and *cyclic* options are not very significant for all experiments, even though *cyclic* was slightly faster in some cases. If not specified otherwise, we refer to the *block* setting (the default in ECHAM). The pictures for cyclic (not presented here) look quite similar.

375 Finally we note that measurements for shorter runs of only one month delivered different optimal values of *nproma*.

#### Best choice of parameter settings for CR configuration 6.2.2

Motivated by the dependency on the parameter *nproma* observed above, we computed the speed-up runtime reduction when using the fastest choice. These runs were performed depending on the number of used nodes (from 1 to 25) in the CR configuration, for both *block* and *cyclic* options. The results are shown in Figure 14. The corresponding best values of *nproma* are given in Tables 3 and 4.

380

It can be seen that for an optimal combination of number of nodes and *nproma*, the radiation could be accelerated by nearly 40%. On the other hand, a bad choice of processors (here between 16 and 23) results in no performance gain or even a loss runtime reduction or even an increase.

The speed-up-runtime reduction for the whole ECHAM model with sp radiation was about 10 to 17%, when choosing an appropriate combination of nodes and nproma. 385



Figure 14. As Figure 12, but for 1 to 25 nodes using the respective best value of *nproma*. Cooresponding optimal values can be found in Tables 3 and 4.

<b>Cable 3.</b> Best values of parameter <i>nproma</i> for different choice of nodes for radiation part.
--

# nodes	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	24	25
block sp	16	16	16	16	76	16	40	40	36	84	24	16	24	180	36	228	152
block dp	16	16	16	16	20	16	16	56	80	28	156	176	32	28	88	56	32
cyclic sp	48	16	16	24	48	16	16	96	52	188	28	16	124	164	16	252	212
cyclic dp	16	24	16	16	20	16	16	152	172	136	148	32	40	124	16	20	52

#### 6.2.3 Parts of radiation code with highest biggest and lowest speed-upsmallest runtime reduction

We identified some subroutines and functions with a very high big and some with a very low performance gain small runtime redecution by the conversion to spsp. They are shown in Tables 5 and 6.

A cause that no even higher speed-up was achieved is that We could not achieve a significant runtime reduction in some cases 390 because several time-consuming parts (as in rk\_mo\_random\_numbers) use expensive calculation with integer numbers,

Table 4. Best values of parameter nproma for different choice of nodes for whole ECHAM.

# nodes	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	24	25
block sp	48	48	32	116	72	120	136	44	36	152	120	56	24	236	52	212	48
block dp	24	24	32	68	148	16	88	220	84	232	84	124	252	36	228	240	228
cyclic sp	48	48	32	124	208	72	192	100	220	152	184	100	92	136	68	208	32
cyclic dp	24	24	16	24	200	40	40	160	72	128	144	132	28	60	84	52	204

Table 5. Parts of radiation code with highest speed-up-runtime reduction by conversion to sp.

Module <del>name</del>	Subroutine/Functionname	Time $dp$ (s) nproma=24	Time $sp$ (s) nproma=48	Speed-up_Runtime_reduction (%)
rk_mo_srtm_solver	delta_scale_2d	960.27	413.35	56.95
rk_mo_echam_convect_tables	lookup_ua_spline	17.42	7.67	55.97
rk_mo_rrtm_coeffs	lrtm_coeffs	78.59	37.06	52.84
rk_mo_lrtm_solver	lrtm_solver	9815.12	4663.04	52.09
rk_mo_srtm_solver	srtm_solver_tr	5005.70	2425.09	51.55
rk_mo_radiation	gas_profile	27.69	13.57	50.99
rk_mo_rad_fastmath	tautrans	3455.26	1790.45	50.53
rk_mo_rad_fastmath	transmit	2837.84	1503.41	47.02
rk_mo_o3clim	o3clim	87.10	47.32	45.67
rk_mo_aero_kinne	set_aop_kinne	233.65	127.88	45.27

taking over 30% of the total ECHAM time consumption in some cases, e.g., in rk\_mo\_random\_numbers. Therefore, these parts are not affected by the *sp* conversion.

#### 6.3 Energy Consumption

We also carried out energy consumption measurements. We used the *IPMI (Intelligent Platform Management Interface)* of the *SLURM* workload manager *ADD* (SchedMD<sup>®</sup>, 2019). It is enabled with the experiment option

```
#SBATCH --monitoring=power=5
```

Here we used one node with the corresponding fast configuration for *nproma* and the option *cyclic*. Simulations were repeated 10 times with a simulation interval of one year.

As Table 7 shows, the obtained energy reduction was 13% and 17% in blade and cpu power consumption, respectively. We consider these measurements only as a rough estimate. A deeper investigation of energy saving was not the focus of our work.

Table 6. Parts of radiation code with lowest speed-up-runtime reduction by conversion to spsp.

Module <del>name</del>	Subroutine/Functionname	Time $dp$ (s) nproma=24	Time $sp$ (s) nproma=48	Speed-up_Runtime reduction (%)
rk_mo_lrtm_gas_optics	gas_optics_lw	6517.89	5647.92	13.15
rk_mo_lrtm_solver	find_secdiff	232.42	209.15	10.01
rk_mo_random_numbers	m	$\underbrace{1.94\cdot10^4}_{2.024}\times10^4$	$\underbrace{1.83 \cdot 10^4}_{1.83 \times 10^4}$	5.67
rk_mo_random_numbers	kissvec	$\frac{8.22 \cdot 10^4}{8.22 \times 10^4}$	$7.79 \cdot 10^{4} $	5.23
rk_mo_lrtm_driver	planckfunction	2169.69	2070.20	4.59
rk_mo_srtm_gas_optics	gpt_taumol	4117.74	3931.92	4.51
rk_mo_random_numbers	low_byte	$\frac{1.36 \cdot 10^4}{1.36 \times 10^4}$	$\underbrace{1.33 \cdot 10^4 1.33 \times 10^4}_{1.33 \times 10^4}$	2.20

Table 7. Energy reduction when using sp radiation in ECHAM.

Energy Consumption	ECHAM with <i>dp</i> radiation (J)	ECHAM with sp radiation (J)	Saved energy (%)		
Blade power	803368	698095	13.1		
CPU power	545762	452408	17.1		

#### 7 Conclusions

We have successfully converted the radiation part of ECHAM to single precision arithmetic. All relevant part of the code can now be switched from double to single precision by setting a Fortran kind parameter named wp either to dp or sp. There is one exception where a renaming of subroutines has to be performed. This can be easily done using a (provided) shell script before the compilation of the code.

We described our incremental conversion process in detail and compared it to other, in this case unsuccessful methods. <u>Some</u> small parts of the code had to remain in higher precision. Here, it would be desirable to further investigate how these parts may be replaced by alternative code or algorithms.

We tested the output for the single precision version and found a good agreement with measurement data. The deviations 410 over decadal runs are comparable to the ones of the double precision versions. The difference between the two version versions lie in the same range.

We achieved an improvement in runtime in coarse and low and resolution of up to 40% for the radiation itself, and about 10 to 17% for the whole ECHAM. In this respect, we could support results obtained for the IFS model by Vana et al. (2017), where the whole model was converted. We also measured an energy saving of about 13 to 17%.

415

405

Moreover, we investigated the parts of the code that are sensitive to reduced precision, and those parts which showed comparably high and low runtime reduction.

The information we provide may guide other people to convert even more parts of ECHAM to single precision. Moreover, they may also motivate to consider reduced precision arithmetic in other simulation codes.

As a next step, the converted model part will be used in coupled ESM simulation runs over longer time horizons.

420 *Code and data availability.* The code is available in the DKRZ git repository under the URL https://gitlab.dkrz.de/PalMod/echam6-PalMod/ ~/network/mixed\_precision\_new2 upon request. The conversion scripts (see Appendix A) and the output data for the single and double precision runs that were used to generate the output plots are available as NetCDF files under https://doi.org/10.5281/zenodo.3560536.

#### **Appendix A: Conversion Script**

The following shell script converts the source code of ECHAM from double to single precision. It renames subroutines and function from the NetCDF library, changes a constant in the code to avoid overflow (in a part that was not executed in the used setting), and sets the constant wp that is used as Fortran kind attribute to the current working precision, either dp or sp. A script that reverts the changes is analogous. After the use of one of the two scripts, the model has to be re-compiled. These scripts cannot be used on the standard ECHAM version, but on the one mentioned in the code availability section.

#!/bin/bash # script rad dp to sp.sh 430 # To be executed from the root folder of ECHAM before compilation for i in ./src/rad src/rk mo netcdf.f90 ./src/rad src/rk mo srtm netcdf.f90 ./src/rad src/rk mo read netcdf77.f90 435 ./src/rad src/rk mo o3clim.f90 ./src/rad\_src/rk\_mo\_cloud\_optics.f90; do sed -i 's/ double/ real/g' \$i sed -i 's/ DOUBLE/ REAL/q' \$i 440 done sed -i 's/numthresh = 500.\_wp/numthresh = 75.\_wp /g' ./src/rad\_src/rk\_mo\_srtm\_solver.f90 sed -i 's/INTEGER, PARAMETER :: wp = dp/INTEGER, PARAMETER :: wp = sp/q' ./src/rad src/rk mo kind.f90 echo "ECHAM radiation code converted to single precision." 445

*Author contributions.* The first author performed all experiments, generated a part of the plots and tables and wrote parts of the manuscript. The second author generated the other part of the plots and tables and wrote main parts of the manuscript.

Competing interests. The authors declare that no competing interests are present.

Acknowledgements. This work was supported by the German Federal Ministry of Education and Research (BMBF) as a Research for

- 450 Sustainability initiative (FONA) through the project PalMod (FKZ: 01LP1515B), Working Group 4, Work Package 4.4. The authors wish to thank Mohammad Reza Heidari, Hendryk Bockelmann and Jörg Behrens from the German Climate Computing Center DKRZ in Hamburg, Uwe Mikolajewicz and Sebastian Rast from the Max-Planck Institute for Meteorology in Hamburg, Peter Düben from the ECMWF in Reading, Robert Pincus from Colorado University, Zhaoyang Song from the Helmholtz Centre for Ocean Research Kiel GEOMAR and Gerrit Lohmann from AWI Bremerhaven and Bremen University for their valuable help and suggestions. Moreover, the authors would like
- 455 to thank the anonymous reviewers for their corrections and helpful comments.

#### References

- Adler, R., Huffman;, G., Chang, A., Ferraro, R., Xie, P., Janowiak, J., Rudolf, B., Schneider, U., Curtis, S., Bolvin, D., Gruber, A., Susskind, J., and Arkin, P.: The Version 2 Global Precipitation Climatology Project (GPCP) Monthly Precipitation Analysis (1979-Present), J. Hydrometeor., 4, 1147–1167., https://www.esrl.noaa.gov/psd, 2003.
- 460 Dawson, A. and Düben, P.: rpe v5: an emulator for reduced floating-point precision in large numerical simulations, Geoscientific Model Development, 10(6), 2221–2230, 2017.
  - DKRZ: Yet Another exChange Tool, https://www.dkrz.de/redmine/projects/yaxt/wiki/Documentation, 2013.
  - Fagan, M., Schlachter, J., Yoshii, K., Leyffer, S., Palem, K., Snir, M., Wild, S., and Enz, C.: Overcoming the Power Wall by Exploiting Inexactness and Emerging COTS Architectural Features, Preprint ANL/MCS-P6040-0816, http://www.mcs.anl.gov/publications, 2016.
- 465 IEEE Standards Association: IEEE Standard for Floating-Point Arithmetic, Tech. rep., https://ieeexplore.ieee.org/servlet/opac?punumber= 8766227, 2019.
  - Intel: Intel<sup>®</sup> Fortran Compiler 18.0 for Linux\* Release Notes for Intel<sup>®</sup> Parallel Studio XE 2018, https://software.intel.com/en-us/articles/ intel-fortran-compiler-180-for-linux-release-notes-for-intel-parallel-studio-xe-2018, 2017.
  - IPCC: Climate Change 2007: Synthesis Report. Contribution of Working Groups I, II and III to the Fourth Assessment Report of the
- Intergovernmental Panel on Climate Change, IPCC, Geneva, Switzerland, 2007.
   Johanson, A. and Hasselbring, W.: Software Engineering for Computational Science: Past, Present, Future, Computing in Science & Engineering, 20, 90–109, https://doi.org/10.1109/MCSE.2018.021651343, 2018.
  - Kleberg, D., Schulzweida, U., Jahns, T., and Kornblueh, L.: CDI-pio, CDI with parallel I/O, https://code.mpimet.mpg.de/attachments/ download/13746/pio\_docu.pdf, 2017.
- 475 Loeb, N. and National Center for Atmospheric Research Staff: The Climate Data Guide: CERES EBAF: Clouds and Earth's Radiant Energy Systems (CERES) Energy Balanced and Filled (EBAF), https://climatedataguide.ucar.edu/climate-data/ ceres-ebaf-clouds-and-earths-radiant-energy-systems-ceres-energy-balanced-and-filled, 2018.
  - Mauritsen, T., Stevens, B., Roeckner, E., Crueger, T., Esch, M., Giorgetta, M., Haak, H., Jungclaus, J., Klocke, D., Matei, D., Mikolajewicz, U., Notz, D., Pincus, R., Schmidt, H., and Tomassini, L.: Tuning the climate of a global model, Journal of Advances in Modeling Earth
- 480 Systems, 4, https://doi.org/10.1029/2012MS000154, 2012.
  - Meador, W. E. and Weaver, W. R.: Two-Stream Approximations to Radiative Transfer in Planetary Atmosphere: A Unified Description of Existing Methods and a New improvement, Journal of the Atmospheric Sciences, 37, 630–643, 1979.
  - Metcalf, M., Reid, J., and Cohen, M.: Modern Fortran Explained, Numerical Mathematics and Scientific Computation, Oxford University Press, 2018.
- 485 Rast, S.: Using and programming ECHAM6 a first introduction, https://www.mpimet.mpg.de/fileadmin/staff/rastsebastian/echam\_6.1\_ lecture.pdf, 2014.
  - Rüdisühli, S., Walser, A., and Fuhrer, O.: COSMO in Single Precision, in: Cosmo Newsletter, 14, pp. 70–87, Swiss Federal Office of Meteorology and Climatology MeteoSwiss, 2014.

SchedMD<sup>®</sup>: Slurm workload manager, https://slurm.schedmd.com, 2019.

490 Schulzweida, U.: Climate Data Operators User Guide, Max-Planck Institute for Meteorology, Hamburg, Germany https://code.mpimet.mpg. de/projects/cdo/embedded/cdo.pdf, 1.9.8 edn., 2019. Score-P: Scalable Performance Measurement Infrastructure for Parallel Codes, http://scorepci.pages.jsc.fz-juelich.de/scorep-pipelines/docs/ scorep-4.1/html/index.html, 2019.

Stevens, B., Giorgetta, M., Esch, M., Mauritsen, T., Crueger, T., Rast, S., Salzmann, M., Schmidt, H., Bader, J., Block, K., Brokopf, R.,

- 495 Fast, I., Kinne, S., Kornblueh, L., Lohmann, U., Pincus, R., Reichler, T., and Roeckner, E.: Atmospheric component of the MPI-M Earth System Model: ECHAM6, Journal of Advances in Modeling Earth Systems, 5, 146–172, https://doi.org/10.1002/jame.20015, 2013.
  - University of East Anglia Climatic Research Unit, Jones, P., and Harris, I.: Climatic Research Unit (CRU): Time-series (TS) datasets of variations in climate with variations in other phenomena v3. NCAS British Atmospheric Data Centre, 2019.
- Vana, F., Düben, P., Lang, S., Palmer, T., Leutbecher, M., Salmond, D., and Carver, G.: Single Precision in Weather Forecasting Models: An
  Evaluation with the IFS, Monthly Weather Review, 145, 495–502, 2017.
  - World Climate Research Programme: Coupled Model Intercomparison Project (CMIP) https://www.wcrp-climate.org/wgcm-cmip, 2019a.
     World Climate Research Programme: Atmospheric Model Intercomparison Project (AMIP), https://www.wcrp-climate.org/modelling-wgcm-mips-2/240-modelling-wgcm-catalogue-amip, 2019b.