

# Single Precision in Weather Forecasting Models: An Evaluation with the IFS

FILIP VÁŇA

*European Centre for Medium-Range Weather Forecasts, Reading, United Kingdom*

PETER DÜBEN

*Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford, Oxford, United Kingdom*

SIMON LANG

*European Centre for Medium-Range Weather Forecasts, Reading, United Kingdom*

TIM PALMER

*Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford, Oxford, United Kingdom*

MARTIN LEUTBECHER, DEBORAH SALMOND, AND GLENN CARVER

*European Centre for Medium-Range Weather Forecasts, Reading, United Kingdom*

(Manuscript received 20 June 2016, in final form 21 November 2016)

## ABSTRACT


Earth's climate is a nonlinear dynamical system with scale-dependent Lyapunov exponents. As such, an important theoretical question for modeling weather and climate is how much real information is carried in a model's physical variables as a function of scale and variable type. Answering this question is of crucial practical importance given that the development of weather and climate models is strongly constrained by available supercomputer power. As a starting point for answering this question, the impact of limiting almost all real-number variables in the forecasting mode of ECMWF Integrated Forecast System (IFS) from 64 to 32 bits is investigated. Results for annual integrations and medium-range ensemble forecasts indicate no noticeable reduction in accuracy, and an average gain in computational efficiency by approximately 40%. This study provides the motivation for more scale-selective reductions in numerical precision.

## 1. Introduction

A question of both theoretical and practical importance is how many bits of real information are carried in the hundreds of millions of variables in a weather or climate prediction model. Related to it is the need of minimal numerical precision of weather models still sufficient to grant a simulation of the desirable quality. For example, since the small-scale variables in a dynamical core are strongly influenced by the irreducible stochastic

uncertainty in subgrid parameterization schemes, one can expect that fewer bits will be needed to encode the real information content contained in small-scale variables, compared with large-scale variables (Palmer 2014). Studies that have tried to reduce precision to a level that can be justified theoretically using a scale-dependent approach for an idealized atmospheric application and a spectral dynamical core suggest that precision can be reduced significantly (see Düben et al. 2014; Düben and Palmer 2014; Thornes et al. 2017).

Naturally since the early days of numerical weather prediction the issue of appropriate numerical precision has been the subject of considerable interest (see Baede et al. 1976). Some precision-sensitive operations, such as matrix inversions, often require so-called double-precision (64 bit) representation of real numbers. But the

 Denotes content that is immediately available upon publication as open access.

Corresponding author e-mail: Filip Váňa, filip.vana@ecmwf.int

common practice to keep double precision as the exclusive representation of all real numbers in the models remains questionable. Reducing the number of bits used to represent weather and climate variables to levels that can be justified theoretically could have a substantial impact on reducing data movement and processing time, resulting in savings that could be reinvested to increase model resolution and hence to improve forecasts.

There are many opportunities to reduce the numerical precision in models. For example, future computer hardware will offer a choice of double, single (32 bit), and additionally half (16 bit) floating-point precision. Furthermore, programmable hardware may allow variables with fully variable levels of precision [see [Düben et al. \(2015\)](#)].

As a motivation to undertake such an activity, we study the impact in a global operational weather prediction model of reducing the precision of real number variables from double to single precision. Although we do have spectral representation of some of the model fields there is no scale-selective approach in the Integrated Forecast System (IFS). Most of the computational expense occurs in the gridpoint space where all scales are treated jointly. Still, we think it is of value to investigate the impact of changing from double to single precision on forecast accuracy. There are few examples of realistic models that run with single precision in the literature, such as the U.S. Navy NOGAPS ([Hogan and Rosmond 1991](#)) or the COSMO model ([Rüdisühli et al. 2013](#)). No systematic study exists, however, comparing double- and single-precision versions of a comprehensive, operational, global weather forecast model. Here we present an analysis of the ECMWF IFS spectral model ([ECMWF 2016b](#)).

In [section 2](#) we describe the model setup and outline the necessary code changes required to obtain a model with reduced precision yet practically identical forecast skill. Results from extensive qualitative and quantitative comparisons are presented in [section 3](#), with further discussion in [section 4](#).

## 2. Code modifications allowing the use of single precision

The version of IFS used for the model simulations was based on model cycle CY41R2 (operational cycle at ECMWF since March 2016). The presented simulations are not coupled to an ocean model, though coupling to the operational wave model is included.

Technically the switch between single- and double-precision arithmetic is realized through the FORTRAN KIND parameter of all real variables in the code. This KIND parameter is defined in one module to modify the precision of all real variables. Additional code changes

are required to allow both single- and double-precision execution with the same code. Those rather technical changes were as follows:

- Replace hard coded thresholds and security constants by parameters defined with intrinsic functions of precision [e.g., to replace code such as “ $x = 10.E+100$ ” by an intrinsic FORTRAN function “ $x = \text{huge}(x)$ ”].
- Extend code interfaces with I/O, MPI, and mathematical libraries (LAPACK/BLAS) to handle both double- and single-precision data.
- Ensure system functions (such as the system clock) are independent of the requested model precision.
- Make sure binary input files are read with the appropriate precision level.

To avoid run-time problems (i.e., floating-point exceptions or under/overflow), the model code had to be further analyzed to ensure all of the computations remain within the floating-point interval available to single precision (between  $2^{-126}$  and  $2^{127}$ ). These changes were mostly seen as desirable in any case in terms of securing the general performance of the model, regardless of its actual precision. Typical changes in this respect were as follows:

- Increase security by avoiding divisions by floating-point numbers that may become zero.
- Rewrite a few formulas to make sure that all intermediate results stay within the single-precision range [e.g., rewrite  $A^4/B^3$  into  $A(A/B)^3$ ].

Finally, to ensure that the single-precision model delivers meteorologically similar results to the double-precision reference, three parts in the model code had to be further modified.

- The cost of the radiation scheme is reduced by recomputing the fluxes on a longer time step than the model step. The contribution to a standard time step has to be upscaled from this coarser temporal sampling. This procedure was found responsible for introducing of an error for low solar zenith angles (below  $86^\circ$ ) in the shortwave radiation code. The double-precision version of the code using a lower threshold to truncate numerical noise was in this case responsible for amplifying a higher amount of noise. This led to the difference of results from the two precision options. By rewriting the code to an equivalent but simpler formula, substituting the implicit and zero-division protected double division of the same variable by single multiplication (not requiring a protection), equal and improved performance for both precision levels was achieved.
- The setup for Legendre transformations, mainly the calculations of the roots and the associated polynomials,

is very sensitive to round-off errors due to their recurrence relations. Before the Swarztrauber algorithm (Swarztrauber and Spitz 2000) in the Legendre transforms was implemented in the IFS, this part of the model had to be computed with quadruple precision for resolutions higher than  $\sim 40$  km. When this code is run in single-precision arithmetic, the model results for higher resolution are substantially degraded. The solution adopted was to ensure that this code is always evaluated in double precision during the model setup (first time step). The Legendre transformation during the standard time step could then be evaluated in single precision.

- Use of single precision in the vertical finite element scheme (VFE) (Untch and Hortal 2004) was found responsible for amplifying model error through the pressure gradient term evaluation. This resulted in unacceptably poor mass conservation. The adopted solution, which almost entirely solved this issue, was to precompute the VFE integral operator with double-precision arithmetic. The resulting operator was truncated to a chosen precision.

The code changes outlined above have been regarded as beneficial with respect to code quality, an improved computational efficiency, and stability of the model. The exercise of exploring use of single precision exposed various problematic places in the code that could potentially also degrade the forecast when using double precision.

### 3. Results

In this section, we compare results of the experiments in single and double precision of long-term simulations and medium-range ensemble forecasts. The former test serves to assess a realism of the model-generated climate when comparing it with the reality (represented by ERA-40 and observations), while the latter test gives a quantitative comparison of the two configurations. Both of them represent typical experiments used for relative quick evaluation of scientific or technical modifications of the model.

#### a. Long-term simulations

To assess the model climate and compare against observations, a 4-member ensemble started from initial time shifted by 36 h to mitigate an effect of diurnal cycle was integrated for 13 months. All integrations are run in uncoupled mode, with prescribed sea surface temperatures, where the first month of the integrations was discarded. The 12-month integrations for 2000–01 are compared with 35 observational datasets and model reanalysis. A pressure mass fixer is used to maintain mass conservativeness for long simulations (Rasch and Williamson 1990). Experiments use T399 ( $\sim 50$  km) horizontal resolution and 137 vertical levels. Examples from

this diagnostic tool could be found on the ECMWF website under model climate evaluation (ECMWF 2016a).

The differences between the two experiments in single and double precision are generally small relative to the magnitude of the systematic forecast errors. The main differences were observed in midlatitudes and polar areas, typically above continents. This corresponds with the high flow variability in these parts of Earth. The maximum difference of annual zonal averages of the two sets of climate simulations did not exceed 2% in relative humidity, 0.5 K in temperature, and  $1 \text{ m s}^{-1}$  in flow speed. In terms of scores computed against observations it was very hard to choose the winner from the two tests. Table 1 summarizes root-mean-square errors from annual means of various model parameters as they are compared against the different datasets. It is quite evident that the single-precision arithmetic is capable of offering very competitive performance to the double-precision reference. Figure 1 shows the spatial distribution of the difference between the averaged precipitation and the GPCP2.2 dataset (Figs. 1a and 1b) and the difference of the TOA shortwave flux compared against CERES-EBAF data (Figs. 1c and 1d). The double-precision IFS outperforms the single-precision IFS for precipitation while shortwave flux favors the opposite conclusion. In either case the error patterns from the two experiments are very similar.

#### b. Medium-range ensemble forecasts

Ensemble forecasts were run 15-day forecasts for 46 start dates that were distributed evenly over an entire year with 50 ensemble members each, starting from 4 December 2013 to 29 November 2014. The resolution used is T399 with 91 vertical levels.

Figure 2 shows the continuous ranked probability score (CRPS; Hersbach 2000) for the experiment run in single precision (dashed red) and double precision (blue) for geopotential and temperature.

The results of the single- and the double-precision experiments are very similar in the extratropics and the small visible differences might still be influenced by the limited sample size (Figs. 2a and 2b). For 500 hPa in the tropics, there is a more striking difference between the experiments (Fig. 2c). These differences stem from a known problem of the IFS. It does not conserve mass, rather it gains or loses mass depending on the combination of time step and resolution, which then results in a forecast bias. This bias then can strengthen or compensate other biases in the model. Hence, it can appear in the scores as an improvement or degradation. Especially in the tropics the scores are quite sensitive to biases, namely geopotential at 500 hPa, due to the smaller variability in comparison to the extratropics.

TABLE 1. Root-mean-square errors from annual simulations compared to various datasets as obtained with single-precision and default double-precision IFS models.

Model quantity	Verifying dataset	Single-precision RMSE	Double-precision RMSE	Unit
500-hPa geopotential	ERA-Interim	111.6	119.8	gpm
Surface winds	QuikSCAT	1.56	1.60	$\text{m s}^{-1}$
Surface winds	SSM/I V6	1.31	1.30	$\text{m s}^{-1}$
High-level cloud cover	<i>CALIPSO</i>	7.12	7.01	%
Midlevel cloud cover	<i>CALIPSO</i>	8.33	8.31	%
Low-level cloud cover	<i>CALIPSO</i>	9.99	10.00	%
Total cloud cover	<i>CALIPSO</i>	9.20	9.28	%
Total cloud cover	ISCCP	8.72	8.73	%
Total precipitation	Xie–Arkin	1.13	1.13	$\text{mm day}^{-1}$
Total precipitation	GPCP 2.2	1.02	1.00	$\text{mm day}^{-1}$
Total precipitation	HOAPS 3.1	1.42	1.39	$\text{mm day}^{-1}$
Total precipitation	SSM/I V6	1.61	1.59	$\text{mm day}^{-1}$
Total precipitation	TRMM_3B43	1.35	1.35	$\text{mm day}^{-1}$
Total precipitation	TRMM/TMI (RSS)	3.17	3.17	$\text{mm day}^{-1}$
Ice water path	<i>CloudSat</i>	13.6	13.7	$\text{g m}^{-2}$
Ice water path	<i>NOAA-10</i>	14.7	14.7	$\text{g m}^{-2}$
Liquid water path	SSM/I V6	35.3	35.1	$\text{g m}^{-2}$
Liquid water path	TRMM/TMI (RSS)	31.9	31.8	$\text{g m}^{-2}$
Water vapor path	HOAPS 3.1	2.37	2.42	$\text{kg m}^{-2}$
Water vapor path	SSM/I V6	2.70	2.73	$\text{kg m}^{-2}$
Water vapor path	TRMM/TMI (RSS)	2.93	2.96	$\text{kg m}^{-2}$
TOA longwave	CERES <i>Aqua</i>	7.00	7.26	$\text{W m}^{-2}$
TOA longwave	CERES EBAF	6.16	6.50	$\text{W m}^{-2}$
TOA longwave	NOAA OLR	13.2	13.6	$\text{W m}^{-2}$
TOA longwave (total – clear sky)	CERES <i>Aqua</i>	7.16	7.40	$\text{W m}^{-2}$
TOA longwave (total – clear sky)	CERES EBAF	6.88	7.11	$\text{W m}^{-2}$
TOA shortwave	CERES <i>Aqua</i>	12.6	12.5	$\text{W m}^{-2}$
TOA shortwave	CERES EBAF	9.11	9.21	$\text{W m}^{-2}$
TOA shortwave (total – clear sky)	CERES <i>Aqua</i>	11.5	11.6	$\text{W m}^{-2}$
TOA shortwave (total – clear sky)	CERES EBAF	9.36	9.46	$\text{W m}^{-2}$
2-m dewpoint temperature	ERA-Interim	1.148	1.126	K
2-m temperature	ERA-Interim	0.939	0.935	K
Snow depth	ERA-Interim	23.46	23.44	cm
Surface latent heat flux	HOAPS 3.1	24.8	25.2	$\text{W m}^{-2}$
Surface sensible heat flux	HOAPS 3.1	13.8	13.9	$\text{W m}^{-2}$
Surface solar radiation downward	Meteosat	12.1	11.9	$\text{W m}^{-2}$
Wind in tropics at 200 hPa	ERA-Interim	1.2945	1.3187	$\text{m s}^{-1}$
Wind in tropics at 700 hPa	ERA-Interim	0.787 43	0.771 94	$\text{m s}^{-1}$
Wind in tropics at 925 hPa	ERA-Interim	0.643 09	0.606 55	$\text{m s}^{-1}$

To test further the sensitivity of the single-precision results, two additional experiments are run with the mass fixer activated so the model conserves global mass. With the fixer enabled, the differences in the extratropics between the single- and double-precision experiment become even smaller and the differences in the tropics are also very small as it could be seen with nearly indistinguishable CRPS scores (Fig. 3).

### c. Computing cost in single and double precision

To evaluate the performance gain brought by the single precision, the two sets of ensemble described in section 3b were evaluated on the ECMWF Cray XC30. The system at that time consists of two Cray XC30 compute clusters, each having some 3500 compute

nodes with Intel E5-2697v2 (Ivy Bridge) processors. Each node has in total 24 physical cores running at 2.7-GHz clock frequency. All experiments were submitted with the same topology having 96 MPI tasks and 8 OpenMP threads running on 768 cores.

To mitigate the variability in job execution time when running on a nondedicated cluster caused by contention on shared resources, MPI network, and file system, all perturbed ensemble members (i.e., excluding the control forecast running with slightly different setting) were considered. The minimum time for the model time step was used as a reliable measure of the undisturbed setup. One experiment as explained in section 3 then represents 2300 forecast jobs. Every 15-day forecast requires 720 time steps to be completed making altogether

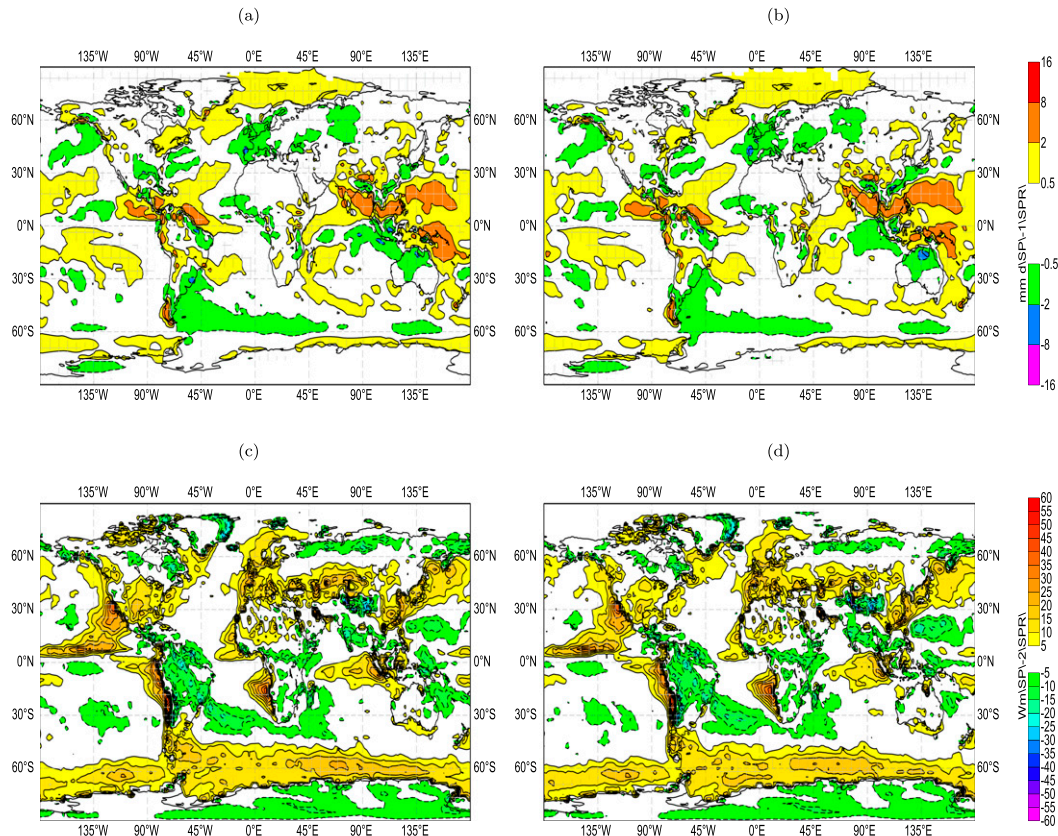


FIG. 1. (a),(b) Difference in total precipitation distribution between the long-term simulations and the GPCP2.2 dataset and (c),(d) difference in TOA shortwave flux between the long-term simulations and the CERES-EBAF dataset obtained with the (left) double-precision reference and (right) single-precision experiments. For precipitation, the double-precision experiments have a mean error of 0.274 and an rmse of 1.00, while the single-precision experiments have a mean error of 0.272 and an rmse of 1.02. For the TOA shortwave flux, the double-precision experiments have a mean error of 1.67 and an rmse of 9.21, while the single-precision experiments have a mean error of 1.3 and an rmse of 9.11  $\text{W m}^{-2}$ . The scores are computed for latitudes between 50°N and 50°S.

1 656 000 time steps to be investigated for the minimum wallclock time per model time step.

The performance of the single-precision executable for the T399 resolution IFS took 37% less run time with respect to the reference double precision. A further gain was achieved after considering the difference in cache utilization between the single-precision and double-precision versions and hence increasing the length of the inner “DO loops” in the model. Nearly all inner loops in gridpoint space are controlled by a parameter set through the namelist. This concept originally designed for vector computers has been useful in controlling cache optimization on scalar machines. The default value of this parameter (i.e., 16) is optimized for double precision. Another ensemble experiment was completed with this parameter doubled to 32. This experiment using the single-precision code with longer inner loops was 40.7% more efficient than the double-precision experiment. This translates to the average reduction

of over 45% of the total wall-clock time for the single-precision jobs, to give an estimate of how those absolute numbers impact the real computation in the concurrent environment of the standard load at the ECMWF computer.

#### 4. Discussion and outlook

We have shown the potential of extensive use of single precision in the ECMWF IFS. We have been able to reproduce very similar results with IFS in terms of model seasonal variability and ensemble forecasting system for a reduced run time of approximately 40%.

The presented simulations exclusively focused on the atmospheric component of the IFS. This represents only one part of the entire forecast system that also includes other major components such as data assimilation, data handling, and ocean model. It is important to further explore a reduction in precision for the fully

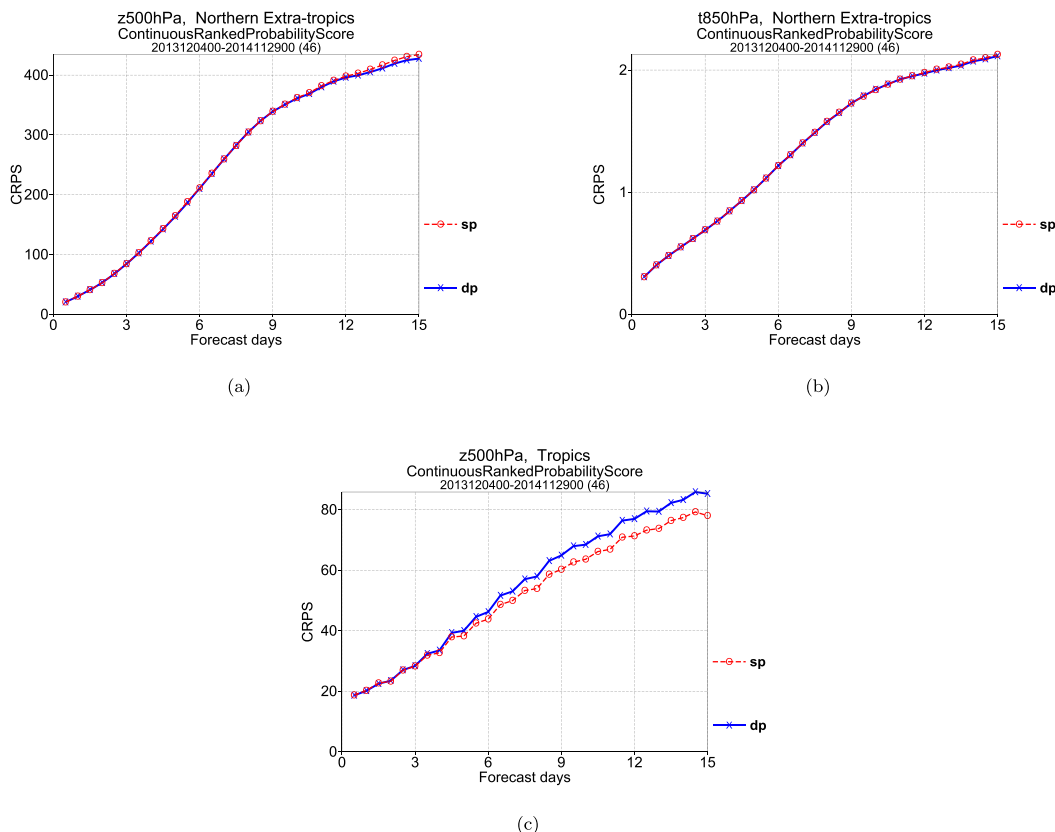


FIG. 2. CRPS of (a) Northern Hemisphere geopotential height at 500 hPa, (b) Northern Hemisphere temperature at 850 hPa, and (c) geopotential height at 500 hPa in the tropics; experiment with single precision (sp, dashed red) and double precision (dp, blue).

coupled system. Further work to verify the model performance and skill at higher resolutions or for climate modeling is also required. This current work will also be the starting point for further studies on the use of reduced precision beyond single precision in both the physical parameterization schemes and the spectral core of the IFS, and in particular, a reduction of numerical precision with decreasing spatial scale following the work of Düben et al. (2014) and Düben and Palmer (2014).

This study is part of a larger scientific and technical program to develop models that only retain those bits of information that can be justified, either theoretically or practically. Therefore, we want to remind the reader of the benefit and impact that an increase in performance by 40% could potentially bring to the scientific community for both weather and climate simulations such as a significant increase in vertical resolution, an increase of ensemble size for ensemble forecasts, or an increase of horizontal resolution at the same computational cost. To illustrate the potential of single precision one can do a simple experiment comparing the single-precision

full ensemble with 50 members with a reduced ensemble being composed from 34 members of a double-precision ensemble, both representing roughly equal computational cost. The attached Fig. 4 with CRPS scores demonstrates the potential gain in terms of forecast quality one could get without requiring more computing resources.

In addition to the reduced computational cost, the process of testing single precision through the code of the IFS has been found generally beneficial in terms of detecting and fixing badly conditioned code. In the future we would like to review more configurations with the single-precision alternative (such as the tangent-linear code) to reveal potentially problematic code there.

The single-precision version of IFS documented in this paper yields overall meteorologically equivalent results at the resolutions examined here (up to 50 km) while significantly reducing the computational resources. These experiments do not demonstrate whether this low precision approach will still work at any higher resolution, though they do not imply any evidence to the

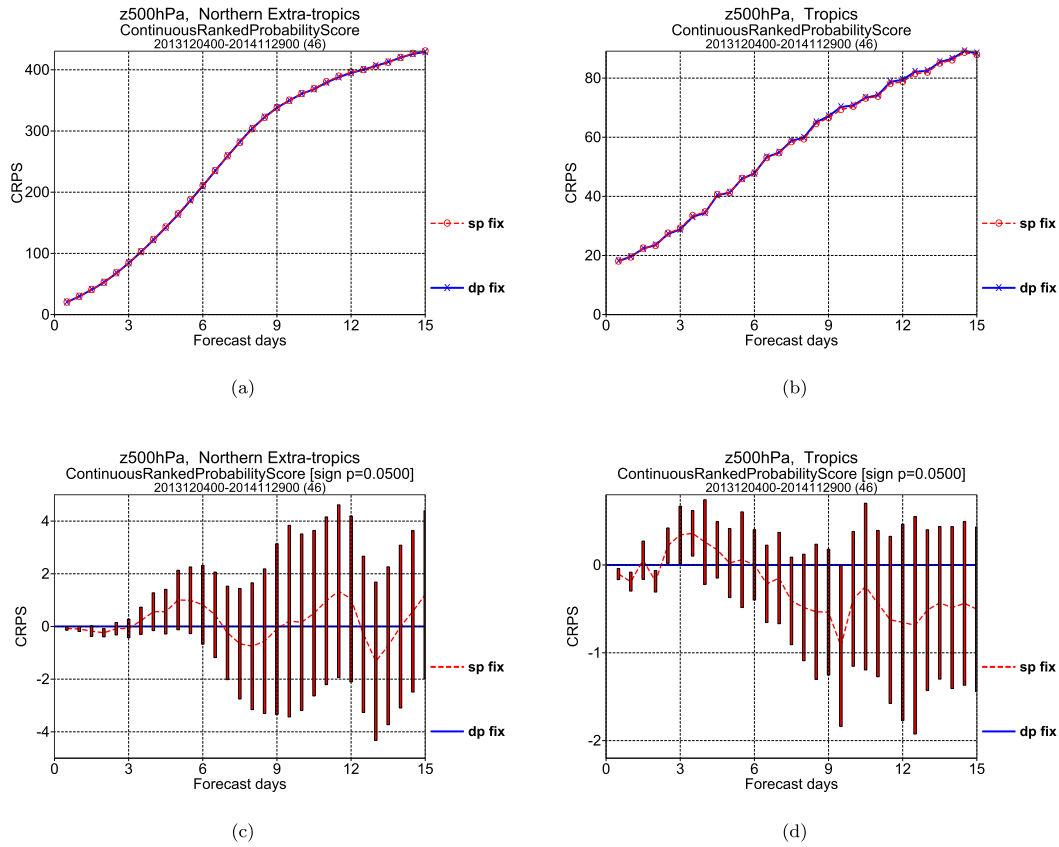


FIG. 3. CRPS of geopotential height at 500 hPa in (a) the Northern Hemisphere and (b) in the tropics; experiment with single precision (sp fix, dashed red) and double precision (dp fix, blue). Both experiments are run with the mass fixer activated. (c),(d) The difference between the two experiments with error bars.

contrary. These conclusions show the potential to save electric power consumption and computational time by investing in more sophisticated codes that can handle reduced precision arithmetic. This may guide the design of future weather forecasting and climate models. We

see this as a first step in the path toward the development of energy efficient high-resolution weather and climate models, and crucial to the goal of developing global cloud resolving models at the earliest opportunity [see also Palmer (2015)].

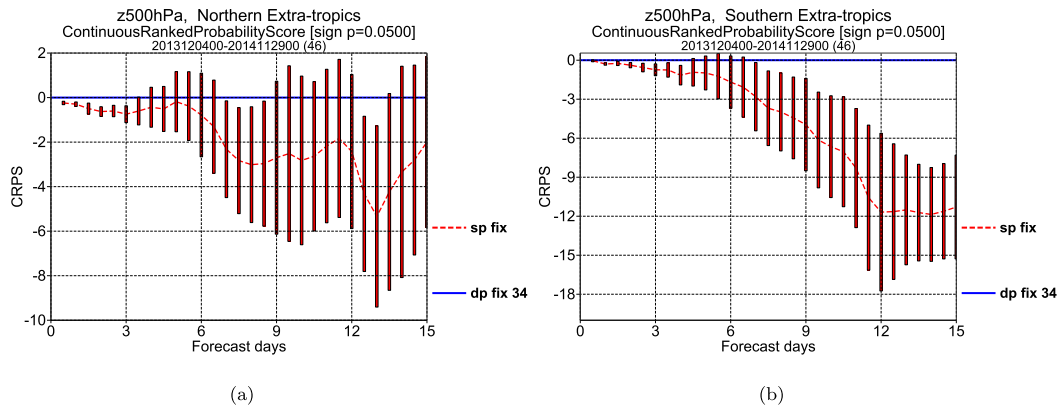


FIG. 4. Difference in CRPS of geopotential height at 500 hPa in (a) the Northern Hemisphere and (b) the Southern Hemisphere; single-precision experiment with full 50 members ensemble with mass fixer (sp fix, dashed red) relative to the double-precision ensemble reduced to 34 members (dp fix 34, blue zero line).

*Acknowledgments.* We want to thank M. Hamrud, N. Wedi, R. Hogan, P. Towers, S. Saarinen, P. Janssen, and M. Diamantakis for their useful comments and suggestions during the single-precision implementation process to the IFS. Special thanks to P. Bechtold for his continuous development and maintenance of the model climate diagnostics package at ECMWF. Tim Palmer and Peter Düben received funding from an ERC Grant (Towards the Prototype Probabilistic Earth-System Model for Climate Prediction, Project Reference 291406).

## REFERENCES

- Baede, A., D. Dent, and A. Hollingsworth, 1976: The effect of arithmetic precision on some meteorological integrations. Tech. Rep. 2, ECMWF, 22 pp. [Available online at <http://www.ecmwf.int/en/elibrary/7869-effect-arithmetic-precision-some-meteorological-integrations>.]
- Düben, P. D., and T. N. Palmer, 2014: Benchmark tests for numerical forecasts on inexact hardware. *Mon. Wea. Rev.*, **142**, 3809–3829, doi:[10.1175/MWR-D-14-00110.1](https://doi.org/10.1175/MWR-D-14-00110.1).
- , H. McNamara, and T. N. Palmer, 2014: The use of imprecise processing to improve accuracy in weather and climate prediction. *J. Comput. Phys.*, **271**, 2–18, doi:[10.1016/j.jcp.2013.10.042](https://doi.org/10.1016/j.jcp.2013.10.042).
- , F. P. Russell, X. Niu, W. Luk, and T. N. Palmer, 2015: On the use of programmable hardware and reduced numerical precision in earth-system modeling. *J. Adv. Model. Earth Syst.*, **7**, 1393–1408, doi:[10.1002/2015MS000494](https://doi.org/10.1002/2015MS000494).
- ECMWF, 2016a: Forecast evaluation. ECMWF, accessed October 2016. [Available online at <http://www.ecmwf.int/en/research/modelling-and-prediction/forecast-evaluation>.]
- , 2016b: IFS documentation. CY41R2, Parts III–VI, ECMWF. [Available online at <http://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model/ifs-documentation>.]
- Hersbach, H., 2000: Decomposition of the continuous ranked probability score for ensemble prediction systems. *Wea. Forecasting*, **15**, 559–570, doi:[10.1175/1520-0434\(2000\)015<0559:DOTCRP>2.0.CO;2](https://doi.org/10.1175/1520-0434(2000)015<0559:DOTCRP>2.0.CO;2).
- Hogan, T. F., and T. E. Rosmond, 1991: The description of the Navy Operational Global Atmospheric Prediction System's spectral forecast model. *Mon. Wea. Rev.*, **119**, 1786–1815, doi:[10.1175/1520-0493\(1991\)119<1786:TDOTNO>2.0.CO;2](https://doi.org/10.1175/1520-0493(1991)119<1786:TDOTNO>2.0.CO;2).
- Palmer, T. N., 2014: More reliable forecasts with less precise computations: A fast-track route to cloud-resolved weather and climate simulators? *Philos. Trans. Roy. Soc.*, **A372** (2018), doi: [10.1098/rsta.2013.0391](https://doi.org/10.1098/rsta.2013.0391).
- , 2015: Modelling: Build imprecise supercomputers. *Nature*, **526**, 32–33, doi:[10.1038/526032a](https://doi.org/10.1038/526032a).
- Rasch, P., and D. Williamson, 1990: Computational aspects of moisture transport in global models of the atmosphere. *Quart. J. Roy. Meteor. Soc.*, **116**, 1071–1090, doi:[10.1002/qj.49711649504](https://doi.org/10.1002/qj.49711649504).
- Rüdisühli, S., A. Walser, and O. Fuhrer, 2013: COSMO in single precision. *COSMO Newsletter*, No. 14, 5.1. [Available online at [http://www.cosmo-model.org/content/model/documentation/newsLetters/newsLetter14/cnl14\\_09.pdf](http://www.cosmo-model.org/content/model/documentation/newsLetters/newsLetter14/cnl14_09.pdf).]
- Swarztrauber, P. N., and W. S. Spitz, 2000: Generalized discrete spherical harmonic transform. *J. Comput. Phys.*, **159**, 213–230, doi:[10.1006/jcph.2000.6431](https://doi.org/10.1006/jcph.2000.6431).
- Thornes, T., P. D. Düben, and T. N. Palmer, 2017: On the use of scale-dependent precision in earth system modelling. *Quart. J. Roy. Meteor. Soc.*, doi:[10.1002/qj.2974](https://doi.org/10.1002/qj.2974), in press.
- Untch, A., and M. Hortal, 2004: A finite-element scheme for the vertical discretization of the semi-Lagrangian version of the ECMWF forecast model. *Quart. J. Roy. Meteor. Soc.*, **130**, 1505–1530, doi:[10.1256/qj.03.173](https://doi.org/10.1256/qj.03.173).