

On the use of scale-dependent precision in Earth System modelling

Tobias Thornes,* Peter Düben^{ORCID} and Tim Palmer

Atmospheric, Oceanic and Planetary Physics, University of Oxford, UK

*Correspondence to: T. Thornes, Oriel College, University of Oxford, Oriel Street, Oxford OX1 4EW, UK. E-mail: tobias.thornes@oriel.ox.ac.uk

Increasing the resolution of numerical models has played a large part in improving the accuracy of weather and climate forecasts in recent years. Until now, this has required the use of ever more powerful computers, the energy costs of which are becoming increasingly problematic. It has therefore been proposed that forecasters switch to using more efficient ‘reduced precision’ hardware capable of sacrificing unnecessary numerical precision to save costs.

Here, an extended form of the Lorenz ‘96 idealized model atmosphere is used to test whether more accurate forecasts could be produced by lowering numerical precision more at smaller spatial scales in order to increase the model resolution. Both a scale-dependent mixture of single- and half-precision – where numbers are represented with fewer bits of information on smaller spatial scales – and ‘stochastic processors’ – where random ‘bit-flips’ are allowed for small-scale variables – are emulated on conventional hardware. It is found that high-resolution parametrized models with scale-selective reduced precision yield better short-term and climatological forecasts than lower resolution parametrized models with conventional precision for a relatively small increase in computational cost. This suggests that a similar approach in real-world models could lead to more accurate and efficient weather and climate forecasts.

Key Words: reduced precision; half precision; multiscale; modelling; Lorenz ‘96; predictability; stochastic parametrization

Received 6 April 2016; Revised 16 November 2016; Accepted 28 November 2016; Published online in Wiley Online Library 30 January 2017

1. Introduction

In a world threatened by rapid anthropogenic climatic change (Hansen *et al.*, 2013), the need for accurate weather and climate forecasts is more pressing than ever. For decades, improvements in numerical forecasts have arisen from using ever faster supercomputers to run models of ever greater resolution and complexity. However, transistor density, and hence the maximum frequency of calculations on a chip of a given size, can no longer be increased in standard silicon processors without incurring a prohibitive increase in power consumption to keep them cool: Dennard scaling, whereby the power density of transistors remains constant as they get smaller (Dennard *et al.*, 1974), is breaking down. Thus, computers can only be made faster by running more processors in parallel (Markov, 2014), with energy costs increasing linearly with the number of processors.

The most advanced global weather forecasting models today have an atmospheric horizontal grid-point spacing of 10–20 km, which implies a resolution of around 50 km. Every halving of the resolution requires around an eightfold increase in computer power (Skamarock, 2004), which includes a power of two for each horizontal dimension plus one for the time dimension. To globally resolve individual convective clouds, less than 1 km across, with conventional hardware would therefore require the

use of exascale supercomputers. Although these could be built within the next decade, their huge energy demands may render them unaffordable to weather and climate institutes until much later (Palmer, 2014).

Conventional computers operate in ‘double-precision’ to minimise rounding errors – which can never be completely eradicated – so that 64 bits are used to represent each number throughout all calculations. By sacrificing some of this numerical precision, forecasts could be made more cheaply. Or, the cost saving could be reinvested to increase model complexity, the number of ensemble members, or horizontal or vertical resolution. This would yield improved forecasts without needing to wait for more powerful machines, with the caveat that to gain from increased resolution other well-known supercomputing problems will need to be solved, such as the imperfect scaling of performance with the number of processors and a limit on the minimum time-step size because of hardware latency.

It has conventionally been assumed that reducing precision will increase rounding error, cancelling out any benefits. But there are good theoretical reasons to believe that, so long as reduced precision is only applied to variables that are resolved at sufficiently small spatial scales rather than to the whole model indiscriminately, the increase in error will be negligible (Palmer, 2012). In a model with finite resolution, subgrid-scale processes

introduce errors that are increasingly significant on scales approaching the truncation limit (Lorenz, 1993). This, combined with the approximate nature of parametrization schemes and data assimilation limitations, means that values of atmospheric variables are inevitably much less precisely known on smaller scales (Skamarock, 2004). All sources of error grow chaotically over time, but if the reduced precision rounding error remains smaller than the inherent uncertainties this error will not affect the overall accuracy of the forecast. Representing small-scale variables in double-precision would then be a waste of computer time, energy and memory resources (Palem, 2014). Indeed, other studies demonstrate good results when using single-precision in less important computations while solving partial differential equations (see, for example, Goddeke *et al.* (2007)), and it has been successfully applied in applications such as computational fluid dynamics (Buttari *et al.*, 2008).

In meteorology, although it has been long known that high numerical precision is not always necessary (Baede *et al.*, 1976), even single-precision is only rarely used, though it is now utilised in, for instance, the new Met Office Natural Environment Research Council (NERC) Cloud Model (Brown *et al.*, 2015). Yet in a recent study a 40% speed-up was obtained by running nearly the entire European Centre for Medium-range Weather Forecasts Integrated Forecast System (ECMWF IFS) model in single-precision with no measurable loss of accuracy (Váňa *et al.*, 2016), and other experiments show that selectively reducing precision even further or using stochastic processors could yield further benefits in idealized models (Düben *et al.*, 2014) and dynamical cores (Düben and Palmer, 2014). Those studies emulated reduced precision on double-precision hardware because although there are High Performance Computing (HPC) architectures that can load and store data in half-precision (16 bits), carrying out calculations in less than single-precision (32 bits) is only possible on special Advanced RISC Machines (ARM) processors and there are currently no HPC stochastic processors. Field Programmable Gate Arrays (FPGAs), which allow the level of precision to be customised for each operation, do exist (Gan *et al.*, 2013; Düben *et al.*, 2015), but these are difficult to program and their suitability for Earth System modelling remains unproven.

In this article we aim to test the potential of two scale-selective ‘reduced precision’ techniques: ‘mixed precision’ and ‘stochastic processors’. Mixed precision involves using fewer bits for each number so that calculations are quicker and data storage requires less memory bandwidth. In contrast to many previous studies, we go beyond single- and utilise half-precision. Stochastic processors operate in double-precision but at a reduced voltage (and energy demand), which means that one or more bits may ‘flip’ at random between 0 and 1 (Palem and Lingamneni, 2013), and that they do not always produce the same output for a given input (they are not ‘bit-reproducible’). We do not explore switching from floating- to fixed-point integer-based arithmetic to reduce cost, which we expect might lead to considerable error in a model where the dynamic range is large (Oberstar, 2007).

Both stochastic processors and half-precision arithmetic may soon be available to HPC, building on the ‘Pascal’ Graphical Processing Unit launched by NVidia in 2016, which can operate in double-, single- and half-precision, with peak speeds of 5.2, 10.6 and 21 teraflops respectively (Trader, 2016). We hope that studies such as ours will encourage the exploitation of such new hardware, especially for ensemble forecasts where reduced precision rounding error may be less than the difference between ensemble members.

Düben *et al.* (2015) have previously used the one-dimensional idealized model atmosphere formulated by Lorenz in 1996 (Lorenz, 2006a), hereafter referred to as Lorenz ’96, to demonstrate the ability of FPGAs to reduce computational costs without increasing the model error. This system exhibits chaoticity and nonlinear interactions between different spatial scales, thereby capturing these important aspects of Earth’s real atmosphere, and has thus been used to test modelling and

data analysis techniques before applying them to real-world atmospheric models.

Here, we extend the Lorenz ’96 system to encompass three rather than two scales or ‘tiers’ of variables to investigate the effects of reducing precision to different degrees on different spatial scales. We compare ‘Low Resolution’ models of the system that explicitly resolve only a single tier of variables, more computationally costly ‘High Resolution control’ models that resolve the top two tiers of variables, and ‘High Resolution reduced precision’ models that are made less costly by sacrificing precision at smaller spatial scales to the ‘truth’ system in which all three tiers are resolved.

In section 2 of this article, we introduce the three-tier extended Lorenz ’96 system for the first time (section 2.1), and describe the parametrized models that we used (section 2.2). In section 3, the methods used to emulate reduced precision are outlined. Section 4 summarises our results and in section 5 we conclude the article by discussing the implications for real-world weather and climate forecasts.

2. The three-tier Lorenz ’96 system

2.1. System set-up

The original Lorenz ’96 system comprises only two tiers of variables. There are K slowly-varying, large-scale X -variables that are analogous to a synoptic quantity such as the static instability of the atmosphere, arranged on a latitude circle. Each is coupled to J fast-varying, smaller-scale Y -variables that could represent some convective-scale quantity (Lorenz, 2006a). The coupled differential equations that govern the system’s evolution are:

$$\frac{dX_k}{dt} = X_{k-1}(X_{k+1} - X_{k-2}) - X_k + F - \frac{hc}{b} \sum_{j=1}^J Y_{j,k}, \quad (1)$$

$$\frac{dY_{j,k}}{dt} = -cbY_{j+1,k}(Y_{j+2,k} - Y_{j-1,k}) - cY_{j,k} + \frac{hc}{b} X_k. \quad (2)$$

Here, F is the large-scale forcing applied to the system and h is the strength of the coupling between spatial scales. The parameters b and c are used to prescribe the relative magnitudes and evolution speeds respectively of the two tiers of variables, which are labelled using the indices $k = 1, \dots, K$ for X and $j = 1, \dots, J$ for Y .

To enable us to compare models of the system at two different resolutions against a still higher-resolution ‘truth’ and to use scale-selective precision, we introduced for the first time a third tier of faster-varying, small-scale variables labelled Z , which might represent phenomena on the scale of individual clouds in the real atmosphere. This created a new system that evolves according to the three coupled differential equations:

$$\frac{dX_k}{dt} = X_{k-1}(X_{k+1} - X_{k-2}) - X_k + F - \frac{hc}{b} \sum_{j=1}^J Y_{j,k}, \quad (3)$$

$$\frac{dY_{j,k}}{dt} = -cbY_{j+1,k}(Y_{j+2,k} - Y_{j-1,k}) - cY_{j,k} + \frac{hc}{b} X_k - \frac{he}{d} \sum_{i=1}^I Z_{i,j,k}, \quad (4)$$

$$\frac{dZ_{i,j,k}}{dt} = edZ_{i-1,j,k}(Z_{i+1,j,k} - Z_{i-2,j,k}) - g_Z eZ_{i,j,k} + \frac{he}{d} Y_{j,k}. \quad (5)$$

There are I small-scale Z -variables coupled to each medium-scale Y -variable, labelled $i = 1, \dots, I$. The new parameters d and e represent the magnitude and evolution speed ratios respectively of the Y - and Z -variables, and g_Z (which for our present purposes will be set to 1) is a damping parameter (see below). Theoretically, any number of tiers could be added to the system in an analogous manner to facilitate more elaborate tests of multiscale behaviour. In our investigations, we used a Fortran program (File S1) to solve




Figure 1. Diagram of the three-tier Lorenz '96 system. Eight large-scale X -variables are cyclically linked and each connected to eight medium-scale Y -variables, each of which is in turn connected to eight small-scale Z -variables not explicitly resolved in this figure but represented by the outer rim.

the full system of equations at each time step of length $\Delta t = 0.005$ Model Time Units (MTUs).

The three tiers of variables are cyclically linked in space as shown in Figure 1, where $I = J = K = 8$. Although each Z -variable is one of a set of eight belonging to a given Y -variable, those at the edges of the set (with $i = 1$ or 8) interact with those at the edges of adjacent sets, which is to say that $Z_{8,j,k}$ is linked to $Z_{1,j+1,k}$ and so on, and the same is true for the Y -variables (with $j = 1$ or 8) at the edges of each X -variable set. Thus, waves can propagate independently on all three spatial scales.

To make the relationship between the small- and medium-scale variables as symmetric as possible with that between the medium- and large-scale variables, we set $b = d$ and $c = e$. Lorenz' original article and numerous studies since – see for example Arnold *et al.* (2013) and Wilks (2005) – have used $h = 1$, $b = 10$ and $c = 10$, which we follow here for simplicity. Those studies took $K = 8$ and $J = 32$, giving 264 variables overall. Here, we choose $K = J = I = 8$, yielding 584 variables, so that the system is not much more computationally complex than those two-dimensional ones and is symmetric between all three spatial scales. This also means that at successively smaller scales the number of variables increases by a factor of ~ 9 , which is close to the eightfold increase in complexity that doubling the horizontal resolution in a real-world GCM would bring about.

To assess the extent to which chaoticity affects the performance of reduced precision models, we investigated two set-ups which differed only in the large-scale forcing, with $F = 10$ and $F = 20$. The chaoticity of each is proportional to the inverse of the 'error doubling time', the time-scale on which an initial perturbation will double in size, which can be computed for each spatial scale independently using the method of Lorenz (2006a) outlined in Appendix A. Our two set-ups have error doubling times for the X -variables of 0.8 and 0.4 MTUs respectively. Comparing these values to the large-scale error doubling time of around 1.5 days that Lorenz (2006a) identified for models of the real atmosphere (in 1996), we deduce that 1 MTU ≈ 2 –4 real atmospheric days for our set-ups.

In an atmospheric model, chaos takes different forms depending on the scale-dependency of the turbulence it produces. Two forms that are important when describing the real atmosphere are 'two-dimensional' (2D) and 'three-dimensional' (3D) turbulence (Leith, 1971). 2D turbulence acts to concentrate the energy E of the system at large spatial scales. This diminishes

the importance of the smallest scales so that small errors do not rapidly propagate upwards (error doubling times are not 'scale-selective') and the error in a forecast of the system can be made arbitrarily small by minimising the uncertainty in the initial conditions. With 3D turbulence, the chaoticity is greater on smaller spatial scales and small-scale errors grow rapidly. This imposes a fundamental lower limit on the size of the error in a forecast regardless of the (inevitably non-zero) error in the initial conditions (Lorenz, 1969).

Ideally, we should like our model set-ups to mimic one of these turbulence regimes. Lorenz' two-level set-up demonstrated 3D turbulence in that the error doubling time was greater for the X - than for the Y -variables (Lorenz, 2006a). We could not use Lorenz' method to compare the doubling times of our three spatial scales because perturbing the Z -variables alone to measure the Z error doubling time knocks the system off its 'attractor' (the subset of states that it tends to occupy) which causes the error in Z to initially shrink while the system rebalances. Hence, we instead computed the leading Lyapunov exponent, a direct measure of the whole system's chaoticity, by perturbing all the variables simultaneously as outlined in Appendix B. We then compared the result for all three tiers (labelled λ_{XYZ}) with that using only the X - and Y -variables (λ_{XY}) or only the X -variables (λ_X).

For the $F = 10$ set-up we found that $\lambda_X = 2.2$, $\lambda_{XY} = 12.8$ and $\lambda_{XYZ} = 9.8$, whilst for the $F = 20$ set-up $\lambda_X = 4.5$, $\lambda_{XY} = 23.5$ and $\lambda_{XYZ} = 20.3$. These results imply that 3D (scale-selective) turbulence is simulated on large spatial scales, but introducing the Z -variables hardly changes the chaoticity, which is more consistent with 2D (non-scale-selective) turbulence on small scales. 3D turbulence on all spatial scales could be obtained by reducing the damping parameter g_Z in Eq. (5). This might be expected to make the Z -variables less sensitive to a reduction in numerical accuracy, since their behaviour would be more chaotic. However, because our reduced precision models always parametrize the Z -variables, this was found to have a negligible effect on our results, and we here retain $g_Z = 1$ for maximum consistency with the two-tier system.

2.2. Parametrization schemes

Our fully-resolved three-tier system resembles the dynamical core of a real atmospheric model in that its equations are nonlinear and describe chaotically evolving coupled quantities with limited predictability. But operational forecast models inevitably have a finite resolution, and parametrize unresolved processes as functions of the resolved variables, giving rise to inaccuracies that may lower the forecast's sensitivity to reduced precision. To make our insights into the trade-off between precision and resolution relevant to real-world models, we hence constructed limited-resolution parametrized models of the three-tier 'truth'.

In 'High Resolution' models Eq. (5) is ignored and the small-scale Z -variables in Eq. (4) are replaced by a parametrization function $U_p(Y_{j,k})$ that represents subgrid-scale behaviour, so that:

$$\frac{dY_{j,k}}{dt} = -cbY_{j+1,k}(Y_{j+2,k} - Y_{j-1,k}) - cY_{j,k} + \frac{hc}{b}X_k - U_p(Y_{j,k}). \quad (6)$$

In 'Low Resolution' models, Eqs (4) and (5) are ignored entirely and the Y -dependent term in Eq. (3) is replaced with a parametrization function $U_p(X_k)$, to give:

$$\frac{dX_k}{dt} = -X_{k-1}(X_{k-2} - X_{k+1}) - X_k + F - U_p(X_k). \quad (7)$$

Written in this way, $U_p(X_k)$ and $U_p(Y_{j,k})$ represent the true sub-grid tendencies of Low and High Resolution systems. They can each be approximated as quartic polynomial functions in a

Table 1. Best-fit coefficients a_0, \dots, a_4 we obtained by fitting a quartic polynomial to the sub-grid tendencies $U_p(X_k)$ and $U_p(Y_{jk})$ to parametrize the Y - and Z -variables respectively, for both set-ups.

Setup	Parametrized	a_0	a_1	a_2	a_3	a_4	ϕ	σ
$F = 10$	Y	0.0343	0.3557	−0.0007	−0.0029	0.0001	0.9879	1.0740
	Z	0.0310	0.3382	0.2411	−0.4320	0.2096	0.9936	0.1176
$F = 20$	Y	0.1073	0.2437	0.0002	−0.0005	0.0000	0.9625	1.4989
	Z	0.0378	0.2357	0.1191	−0.0595	0.0018	0.9875	0.1509

Also listed are the sub-grid autocorrelation ϕ and standard deviation σ we used for stochastic parametrization schemes, following Arnold *et al.* (2013).

Table 2. The number of bits in the sign, exponent and significand components and the maximum and minimum resolvable values (including sub-normal numbers) of floating point numbers when represented in double-, single- and half-precision.

Precision level	Number of bits					Resolvable range	
	Total	= Sign	+ Exponent	+ Significand		Minimum	Maximum
Double	64	= 1	+ 11	+ 52		2^{-1074}	2^{1023}
Single	32	= 1	+ 8	+ 23		2^{-149}	2^{127}
Half	16	= 1	+ 5	+ 10		2^{-24}	2^{15}

‘deterministic’ way, so that, for the High Resolution example,

$$U_p^{\text{det}}(Y_{jk}) = a_0 + a_1 Y_{jk} + a_2 Y_{jk}^2 + a_3 Y_{jk}^3 + a_4 Y_{jk}^4, \quad (8)$$

where a_0, \dots, a_4 are constants obtained by measuring the true $U_p(Y_{jk})$ as a function of Y_{jk} for a long run of the system and finding the best-fit curve (Wilks, 2005). The Low Resolution $U_p^{\text{det}}(X_k)$ is defined similarly. The coefficients we obtained in each case are listed in Table 1.

This method assumes that the effects of unresolved variables can be predicted exactly and consistently from the larger scales (Lorenz, 1969), which is unrealistic for a chaotic system. A better estimate of the effects of unresolved variables can be obtained from an ensemble of multiple forecasts, each of which uses a slightly different value for U_p that has been randomly perturbed relative to U_p^{det} . In this study, we used two such ‘stochastic’ parametrization methods, previously employed by Arnold (2013). In their ‘additive’ scheme, U_p is estimated by adding a component s_{jk} constructed from a red-noise random number, so that when the Z -variables are parametrized,

$$U_p(Y_{jk}) \approx U_p^{\text{det}}(Y_{jk}) + s_{jk}. \quad (9)$$

In the ‘multiplicative’ scheme, the deterministic value is instead multiplied by s_{jk} , so that:

$$U_p(Y_{jk}) \approx (1 + s_{jk})U_p^{\text{det}}(Y_{jk}). \quad (10)$$

3. Reduced-precision models

3.1. Emulating reduced precision

A double-precision floating-point comprises 64 bits: 1 bit for the sign (± 1); 11 bits for the ‘exponent’ E ; and 52 bits for the ‘significand’ S (a factor between 1 and 2) so that its overall value is $N = \pm S \times 2^E$, and $10^{-1022} \leq N \leq 10^{1023}$. In the Institute of Electrical and Electronics Engineers (IEEE) definition of single- and half-precision (IEEE, 2008), with fewer bits, the range and resolution of numbers that can be represented lessens as shown in Table 2. In the absence of hardware capable of half-precision arithmetic and real stochastic processors we ‘emulate’ reduced precision on conventional hardware.

In ‘mixed precision’ mode our emulator rounds numbers to the nearest value that is resolvable at the precision (double, single or half) specified for each spatial scale during calculations. It uses IEEE standards except for the special case where the first bit to

Table 3. Characteristics of the ten model types compared in this study.

Model type	Bits in X	Bits in Y	p_X	p_Y	Relative cost
D	64	–	–	–	1
DD	64	64	–	–	9
DS	64	32	–	–	5
DH	64	16	–	–	3
SH	32	16	–	–	2.5
SP1	64	64	0	0.05	5.5
SP2	64	64	0	0.1	5
SP3	64	64	0.001	0.05	5.5
SP4	64	64	0.001	0.1	5
SP5	64	64	0.005	0.01	6

Double- and mixed-precision models are labelled ‘D’ for ‘double’, ‘S’ for ‘single’ and ‘H’ for ‘half’-precision using two letters to represent the precision in the X - and Y -variables respectively. Stochastic processor models are labelled ‘SP’ and numbered in order of increasing bit-flip probability. Listed are the number of bits for each variable, the probability of a bit-flip occurring for the X -variables (p_X) and Y -variables (p_Y) and the estimated relative computational costs. Changes in p_X alone negligibly affect the cost because there are eight times as many Y -variables as X -variables.

be rounded away is 1 and the rest are 0, where the emulator does not enforce the IEEE stipulation that the final bit of the result must be 0. This rare case should not significantly affect our results. In ‘stochastic processor’ mode the emulator invokes two random numbers, r_1 and r_2 , which can each range between 0 and 1. r_1 determines whether a bit-flip occurs, and r_2 determines which of the 52 bits in the significand it applies to. If p is the probability of bit-flips, in a given calculation a flip will only take place when $r_1 < p$. Because p is always small, it is assumed that no more than one bit-flip can occur per floating-point operation. The size of the effect depends on which bit is flipped, and the sign and exponent bits have to be protected entirely to avoid crashes; such protection could be achieved in a real stochastic processor by physically isolating these bits. In such a processor, p would increase as the voltage was decreased: less reliable processors would be cheaper to use.

The types of High Resolution reduced-precision models we implemented are listed in Table 3. The three ‘mixed precision’ types are labelled using two letters representing the precision in the X - and Y -variables respectively; the five ‘stochastic processor’ types are labelled SP1–5, in order of increasing probability of bit-flips. Preliminary tests indicated that reducing the X -variables to half-precision had very deleterious results. Hence, in accordance with our multiscale hypothesis that smaller-scale variables require less precision, only the Y -variables were resolved with less than single-precision in any of our simulations. The reduced precision models were compared to High Resolution control (type DD) and Low Resolution (type D) models in double-precision.

3.2. Estimating relative computational costs

Artificially emulating reduced precision is computationally costly, so explicit measurements of our models’ relative computational costs cannot be made. However, Table 3 includes rough estimates of the costs (relative to type D) that might be expected with hypothetical hardware. The estimates are based on the ‘effective number of variables’ resolved by each model type, a function of

the actual number of variables and the effective computational cost of each variable relative to double-precision. The small costs associated with initialisation and output, which would be carried out in double-precision; with switching between different processors to exploit different levels of precision; and with the increased number of links between interconnected variables in High Resolution models, are ignored. Hence, the computational cost savings we give are overestimates of the potential overall energy savings.

In mixed precision, we assume that the cost required to compute any single X - Y - or Z -variable is independent of spatial scale but is proportional to the number of bits with which that variable is represented. Model type D explicitly resolves 8 variables in double-precision so the effective number is also 8. Type SH uses half the number of bits to resolve each of the 8 X -variables and a quarter the number to resolve each of the 64 Y -variables, so its effective number is $8/2 + 64/4 = 20$, yielding a relative cost of $20/8 = 2.5$.

For the stochastic processor models, we estimated the fractional voltage saving V for a given probability of bit-flips using figure 4 of Sloan *et al.* (2010), which is based on measurements of FPGA hardware with a stochastic ‘fault injector’. We converted this to an energy saving E through the relationship $E \propto P \propto V^2$ where P is the dynamic power consumption (Karakonstantis and Roy, 2011), a very crude approximation (Snowdon *et al.*, 2005) that gives only a rough estimate. We then computed the effective number of variables by multiplying the number of X - and Y -variables by the value of $1 - E$ appropriate to each spatial scale. For instance, Model type SP1 has a probability $p_X = 0$ of flips in the 8 X -variables and $p_Y = 0.05$ of flips in the 64 Y -variables which reduces the cost of each Y -variable by $E = 44\%$ according to the figure in Sloan *et al.* (2010), so the effective number of variables is $8 \times 1 + 64 \times 0.56 \approx 44$. Although the reduced precision methods all have higher cost estimates than type D, most are considerably cheaper than type DD.

4. Results

The accuracy of the two double-precision and eight reduced-precision model types was assessed using four metrics: the models’ ability to predict the short-term evolution of the system (‘weather’ simulations); the system’s longer-term mean state (‘climate’ simulations); its chaoticity; and its tendency to oscillate between different behavioural ‘regimes’. The models’ predictions for the large-scale X -variables were compared to the ‘true’ X -variables evaluated using all three tiers of variables; because Low Resolution models cannot resolve Y -variables, none of the models were assessed on their ability to predict these. Results for models of type DS, DH, SP2 and SP4 will not be explicitly discussed here because they behave almost identically to DD, SH, SP1 and SP3 respectively. This section will briefly discuss the methods of assessment we used and the results we obtained.

4.1. Short-term forecasting ability

Short-term forecasting ability was assessed using the absolute Error score and the Ranked Probability Skill Score (RPSS), each averaged over the 8 X -variables and over 250 independent sets of initial conditions to smooth the plots. For each initial condition, we created an ensemble of 50 forecasts that differed only in the random numbers invoked in the stochastic parametrization schemes. The absolute Error is the difference between the value forecast by the model and the ‘truth’, averaged over the 50 ensemble members. The RPSS measures the difference between the cumulative probability distribution function (pdf) of the 50 values predicted by the ensemble and the cumulative pdf of the ‘truth’. It takes into account the uncertainty associated with parametrized processes, and is a measure of a model’s ability to correctly identify the ‘risk’ of any particular weather event

occurring, reflecting the forecast’s economic value in the real world (Palmer and Richardson, 2014).

To compute the RPSS here, we created a model cumulative pdf P_l by binning the outputs predicted by the ensemble members into ten bins of possible X -variable values, labelled using the index l , which were defined so that they had equal populations when the whole time series of the ‘truth’ run was binned. P_l was normalised by dividing by the number of ensemble members. The Ranked Probability Skill (RPS) of the ensemble forecast was defined by comparing this to the ‘true’ system’s cumulative pdf T_l at this time step, which was 0 for all bins below the ‘true’ value and 1 thereafter, according to Wilks (2006):

$$RPS = \sum_{l=1}^{10} (T_l - P_l)^2. \quad (11)$$

The RPSS was obtained from this by comparing the mean RPS across all 250 sets of initial conditions, RPS , to the RPS of a reference climatology in which all the bins are equally populated, RPS_{ref} , using the equation:

$$RPSS = 1 - RPS/RPS_{\text{ref}}. \quad (12)$$

More accurate forecasts have an RPSS that is closer to 1.

Figures 2 and 3 show the absolute Error and RPSS scores of the models for the $F = 20$ and $F = 10$ set-ups respectively. The forecast skill improves when the resolution is increased (from D to DD) in all cases. But, under additive parametrization, SH performs just as well as DD (we expect any small differences would disappear with more runs). Under multiplicative parametrization, SH begins to fall in accuracy relative to DD but only after around 1 MTU in both set-ups, which suggests that a lower-cost model that begins in single-precision but switches to half-precision once the forecast reaches a lead-time τ might not exhibit this decline. Further tests confirmed that the loss of accuracy could indeed be eradicated in this way in both set-ups, so long as τ was at least 0.01 MTUs for $F = 20$ or 0.05 MTUs for $F = 10$. We attribute this to the onset of errors associated with half-precision being delayed until after these errors become very small compared to the overall model error, which increases with forecast lead-time. Implementing a similar delay in real-world weather models could be a viable way of reducing the rounding errors in low-cost, reduced-precision runs.

Figures 4 and 5 show Error and RPSS plots for the ‘stochastic processor’ models. They illustrate that using a stochastic processor is only worthwhile if the large-scale variables are protected from bit-flips. The poor performance of SP1 relative to DD shows that even a modest probability of faults in the calculations can have a significant effect. Model type SP3 performs no better than type D, whilst SP5 performs much worse: a fault rate of just 0.5% in the X -variables yields a very poor forecast. Tests revealed that the forecasting ability of SP3 significantly improved when the two most important significant bits were protected, but it is unclear how such difficult-to-isolate bits could be protected in real hardware.

4.2. Long-term forecasting ability

A model’s ability to forecast long-term mean conditions can be measured using the Kolmogorov–Smirnov (KS) statistic D_n and the Hellinger Distance H . Because of the isotropy between variables at any spatial scale these are non-local statistics that compare the cumulative pdfs of the X -variables, Ψ_{truth} and Ψ_{model} , across all time steps. The KS statistic measures the maximum difference between the ‘truth’ and model pdfs, so that (Wilks, 2005):

$$D_n = \max(|\Psi_{\text{truth}} - \Psi_{\text{model}}|). \quad (13)$$

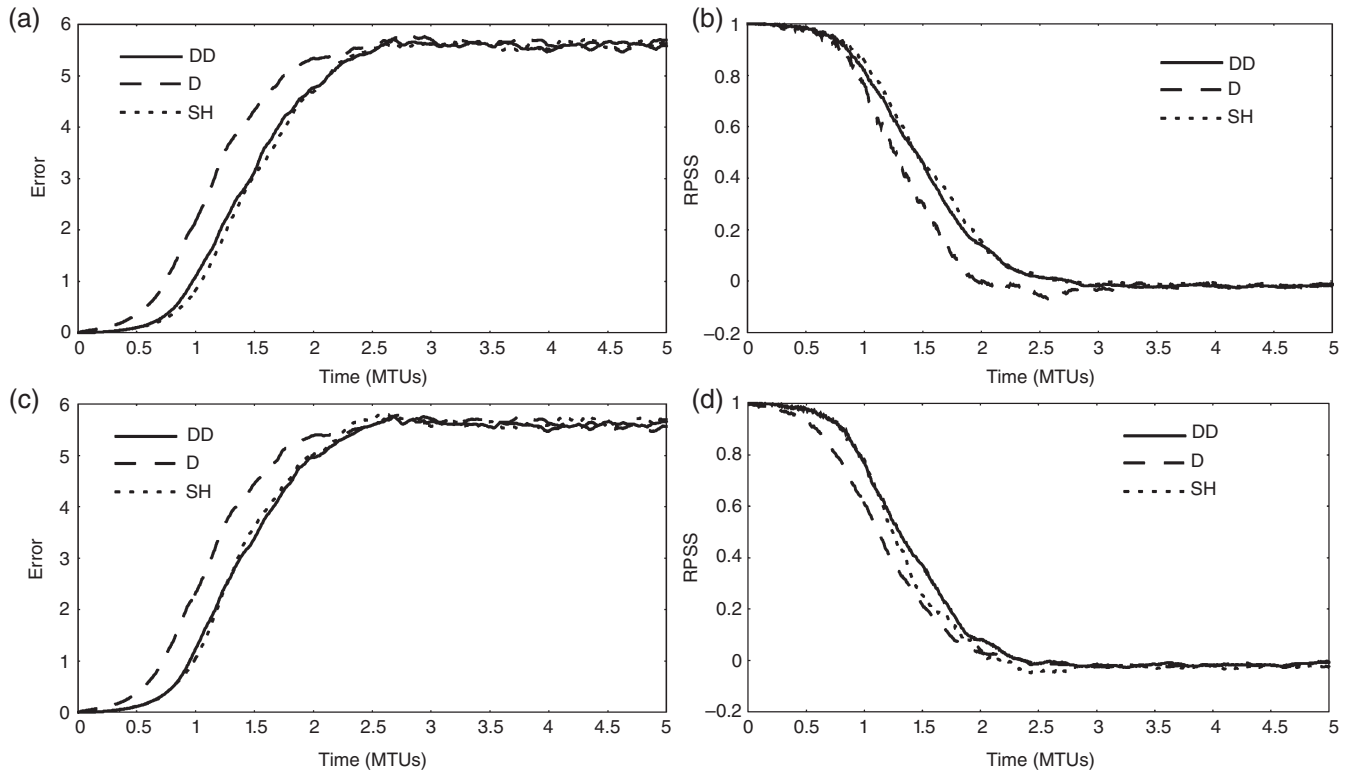


Figure 2. Error scores (a,c) and RPSSs (b,d) for models DD, D and SH of the $F=20$ system using additive (a,b) and multiplicative (c,d) stochastic parametrization schemes.

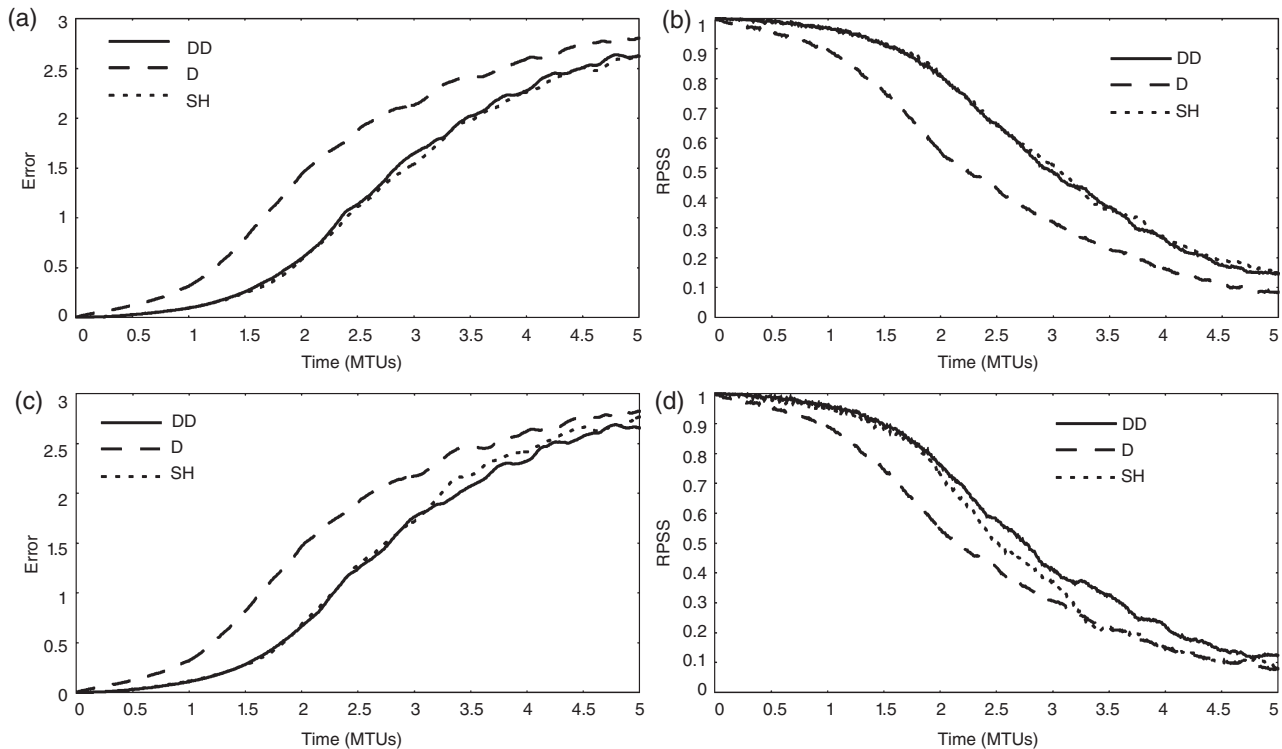


Figure 3. Error scores (a,c) and RPSSs (b,d) for models DD, D and SH of the $F=10$ system using additive (a,b) and multiplicative (c,d) stochastic parametrization schemes.

The Hellinger Distance, by contrast, measures the difference in overall shape between the two pdfs, and is defined through (Arnold *et al.*, 2013):

$$H^2 = \frac{1}{2} \int (\sqrt{\Psi_{\text{truth}}} - \sqrt{\Psi_{\text{model}}})^2 dx. \quad (14)$$

The smaller the value of D_n or H , the more accurate the forecast. Here, D_n and H were averaged across five long (50 000 MTU) integrations with independent initial conditions. The standard error in the mean (Hughes and Hase, 2010) was used to quantify the uncertainty.

Figure 6 shows the results for both set-ups, using additive stochastic and deterministic parametrization schemes to elucidate the influence of the scheme stochasticity. In both set-ups the double-precision and mixed-precision models generally lie within the error bars of one another, although reducing resolution (type D) increases the Hellinger Distance significantly when using the deterministic scheme in the $F=20$ set-up and the stochastic schemes in the less chaotic $F=10$ set-up. In reduced precision, only SP1, under additive parametrization with $F=10$, exceeds type D in accuracy by a significant amount. But most importantly,

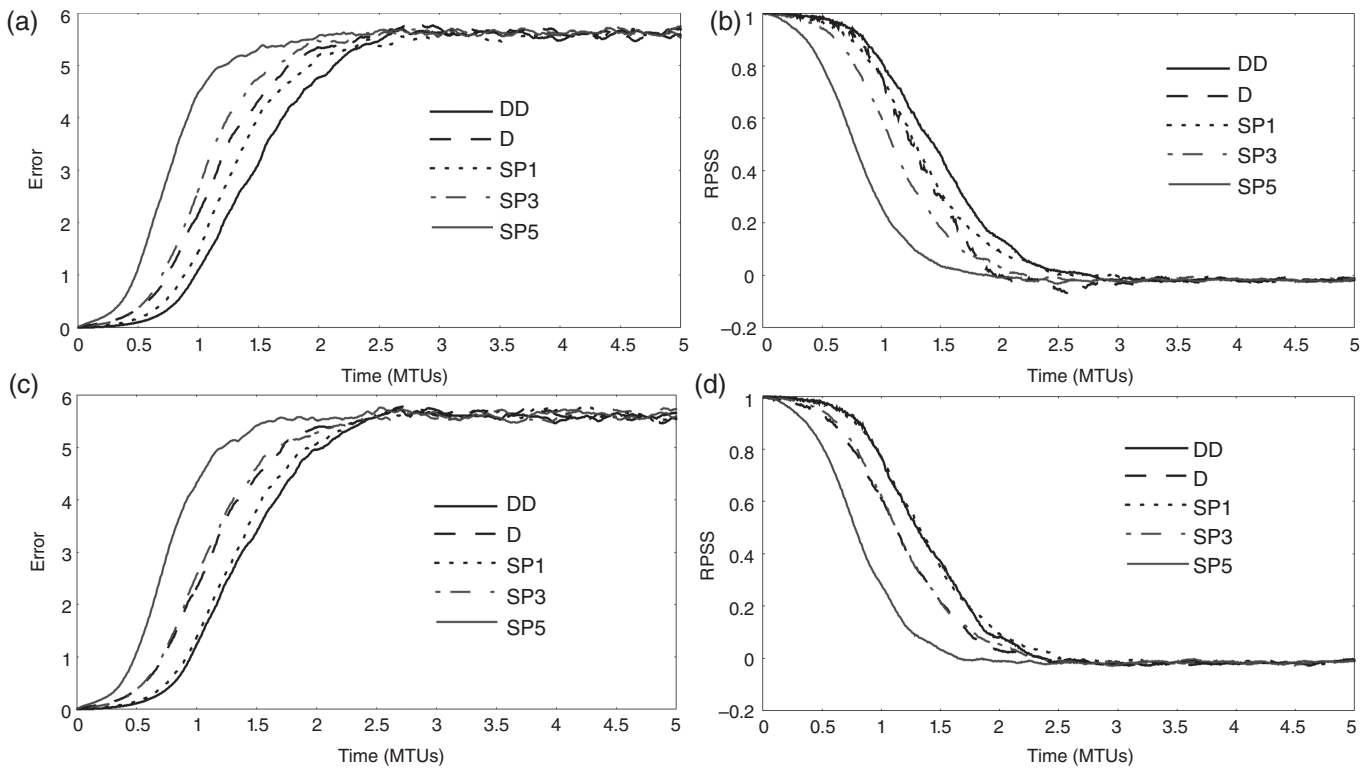


Figure 4. Error scores (a,c) and RPSSs (b,d) for stochastic processor models of the $F = 20$ system using additive (a,b) and multiplicative (c,d) stochastic parametrization schemes.

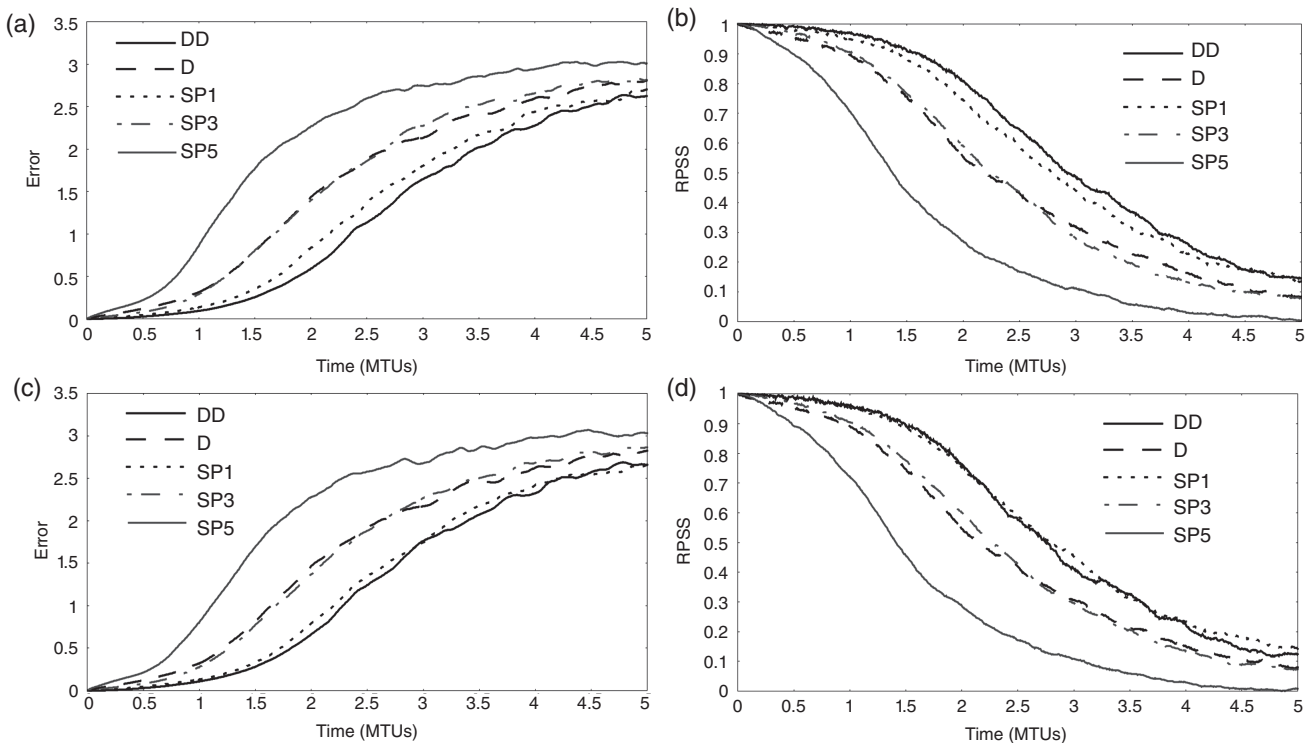


Figure 5. Error scores (a,c) and RPSSs (b,d) for stochastic processor models of the $F = 10$ system using additive (a,b) and multiplicative (c,d) stochastic parametrization schemes.

there is no significant loss of accuracy relative to DD in any of the reduced-precision cases given the error, except when bit-flips are allowed in the X -variables (SP3-5).

4.3. Atmospheric chaoticity

An accurate model of the real-Earth atmosphere should conserve certain quantities such as mass. A convenient proxy for these in our system is the chaoticity, which can be measured using the ‘Error Doubling Time’ as described in Appendix A. Figure 7 plots the error doubling times computed from the time series

for the ‘truth’ and model integrations of both set-ups. For the stochastic parametrization schemes, some of the deviation contributing to the difference between the perturbed and unperturbed runs is due not to the initial perturbation but to the stochasticity of the schemes themselves. Hence, the error doubling times are larger – because the perturbed and unperturbed runs deviate from one another less rapidly – using the deterministic parametrization scheme where all differences are attributable to the initial perturbation. Except in the deterministic case, model type D exhibits smaller error doubling times further from the truth than those of DD, implying that increasing

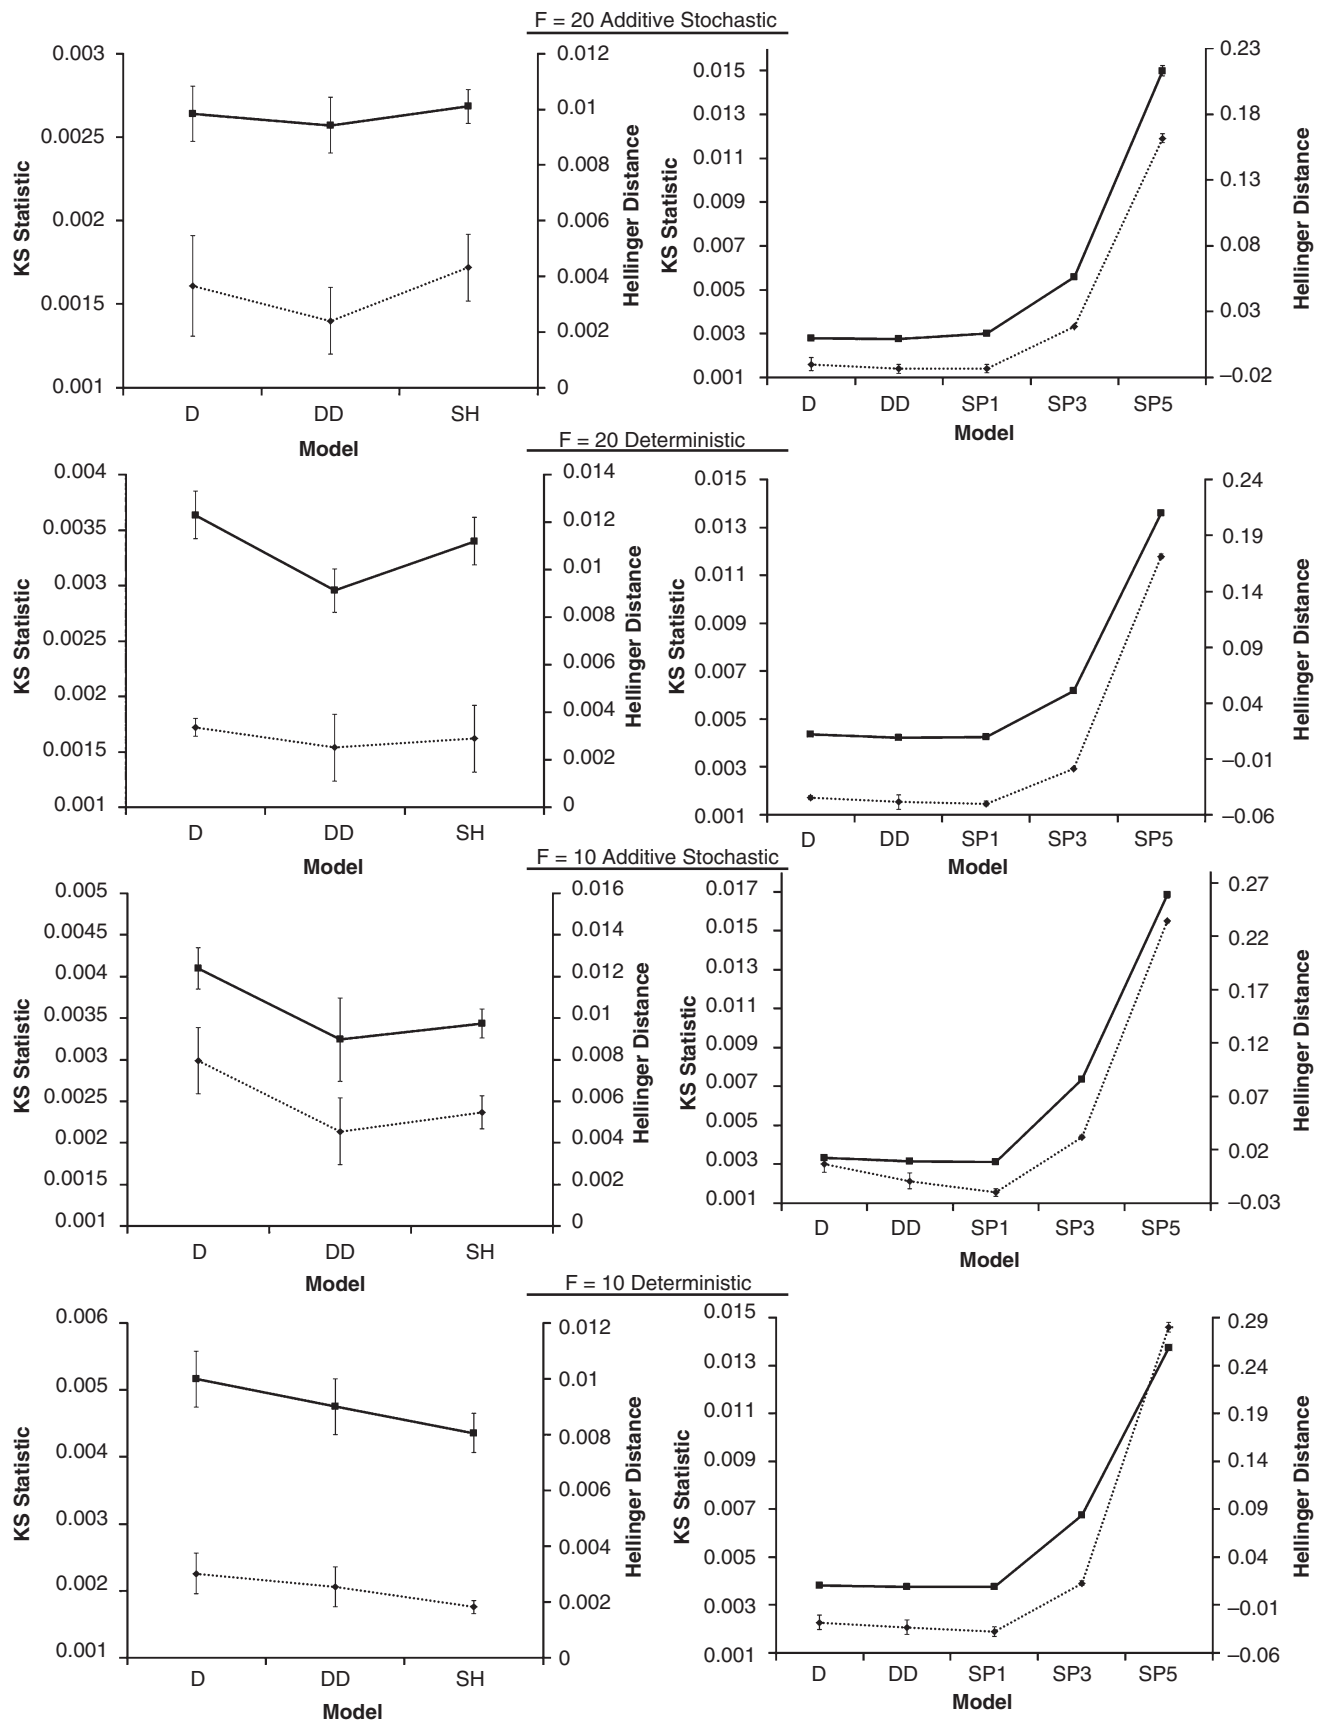


Figure 6. KS statistics (dotted lines) and Hellinger Distances (solid lines) for models D, DD and SH (left plots) or D, DD and SP1-5 (right plots) of both set-ups, using additive stochastic (top), and deterministic (bottom) parametrization schemes. Lines are drawn between points to make the differences between models clearer. Error bars (not always visible) represent the standard error in the mean.

the resolution genuinely improves the accuracy of stochastic simulations.

Model types SH and DD both perform well, matching the true doubling time for the $F=20$ set-up (0.24) exactly under additive parametrization, where model D underestimates it by 20% (0.19). SP1 sometimes simulates the true chaoticity well,

even matching DD for accuracy in, for example, the case of stochastic parametrization in the $F=10$ set-up, but in other instances performs poorly even relative to type D, indicating that even with a low probability of bit-flips stochastic processors do not reliably capture the fundamental characteristics of the simulated system.

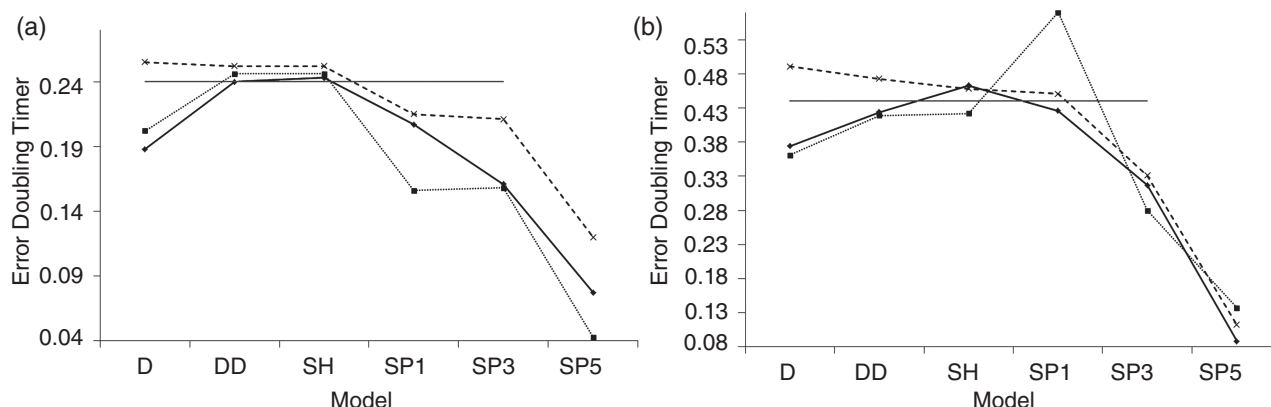


Figure 7. Error Doubling Times in MTUs for the 'truth' (grey lines), and models of type DD, D, SH and SP1–5 using additive stochastic (solid lines), multiplicative stochastic (dotted lines) and deterministic (dashed lines) parametrization schemes for the (a) $F = 20$ and (b) $F = 10$ set-ups. Lines are drawn between points to make the differences between models clearer.

4.4. Regime behaviour

A system exhibits 'regime behaviour' – akin to, for instance, El Niño/Southern Oscillation (ENSO) regimes on Earth – if it occupies distinct regions of phase space at different times and is able to switch between these regions on a time-scale that is large relative to the system's oscillations (Lorenz, 2006b). Christensen *et al.* (2015) and Dawson and Palmer (2015) have previously identified one example of dual-regime behaviour in the two-tier Lorenz '96 system whereby the spatial covariance Σ of the eight large-scale X -variables, which describes how the X -variables are linked in space, switches between positive and negative states. The regimes are also visible in the relative strengths of the first four Principal Components (PCs) of the entire mean-subtracted time-series: PCs 1 and 2 dominate when the system is in Regime A, whereas PCs 3 and 4 dominate in Regime B. Further details can be found in Christensen *et al.* (2015).

We found that our system exhibits such regimes, but only in the less chaotic, $F = 10$ case. For the 'truth' and each model, we computed PCs 1–4 at each time step for long runs of 5000 MTUs (approximately 40 atmospheric years) each. The average value was subtracted and the data were smoothed by applying a rolling average over 0.2 MTUs to make the regimes as visible as possible, then two-dimensional pdfs were plotted as a function of $|PC1, PC2|$ and $|PC3, PC4|$, which are the magnitudes of PCs 1 and 2 and PCs 3 and 4 added in quadrature respectively. The presence or absence of regimes was confirmed using 'vector plots' of the average trajectory at each point in PC space, akin to those plotted by Christensen *et al.* (2015). Any regimes are visible as the centres of the orbits of the Principal Components. These plots are shown for the entire 'truth' runs of the two set-ups in Figure 8, and for selected double-precision and reduced-precision models of the $F = 10$ set-up in Figure 9. The Principal Component pdfs themselves are provided in Figures S1 and S2.

Nearly all the models failed to capture the small Regime B peak visible in the 'truth' pdf for $F = 10$, or merged it with the much more dominant peak for Regime A, with the exceptions of model types D and SH with additive stochastic parametrization and SP1 with multiplicative stochastic parametrization, all of which are shown. It is possible that these cases possess a near-optimal level of noise, contributed to by rounding errors, bit flips and stochastic parametrization schemes, which allows the models to explore a sufficient breadth of phase space to replicate both regimes of the system (Christensen *et al.*, 2015). Increasing this noise further (as in SP2–5 and, presumably, SP1 in the additive stochastic case) may mask the regimes because random perturbations carry the system away from the true attractor too often for its structure to be accurately reproduced. For all parametrization schemes model SP3 reproduced regimes much less well than did SP1, and SP5 performed even worse. But using mixed precision (model type SH) did not significantly alter the

pdfs relative to the double-precision control (type DD), whilst lowering the resolution (type D) degraded the accuracy.

5. Conclusions and future work

Experiments were carried out to investigate the effects of applying 'mixed precision' and 'stochastic processors' in a scale-selective way to a system based on Lorenz '96 but extended to include three tiers of variables – large-scale X -, medium-scale Y - and small-scale Z -variables – for the first time, using software capable of emulating reduced precision on conventional computer hardware. The three-tier system exhibits the chaoticity and nonlinear interactions between scales that are characteristic of the real Earth atmosphere, but its low complexity and the possibility of defining the 'truth' exactly meant that the accuracy of cheap Low Resolution models and expensive High Resolution models in double-precision, and cheap High Resolution models in reduced precision could be straightforwardly compared. This enabled us to test the hypothesis that by trading numerical precision for resolution, better forecasts of the true state of the system could be made for similar computational costs.

Scale-selective mixed precision demonstrated considerable promise. For short-term forecasts, reducing the X -variables to single-precision and the Y -variables to half-precision (model type SH) had no significant effect on the accuracy when additive stochastic parametrization was used, even though the estimated computational cost was reduced from 9 (for the control) to 2.5 times that of the Low Resolution model. A loss of accuracy for this model type under multiplicative stochastic parametrization could be remedied by switching to half-precision only after a small delay, retaining most of the computational cost saving.

For forecasts of long-term mean conditions, atmospheric chaoticity and regime behaviour, mixed-precision models (SH) were as proficient as the double-precision controls (DD). KS Statistics and Hellinger Distances were consistent for DD and SH and their error doubling times sometimes exactly matched the 'truth', which the low-resolution alternative (type D) underestimated by up to 20%. The SH additive forecast exceeded the ability of DD to reproduce atmospheric regimes. At no point did the mixed-precision models crash, so the maximum half-precision exponent cannot have been exceeded anywhere even in these long runs. Overall, our results suggest that resolving atmospheric variables at scales that are currently parametrized or close to the truncation scale in half-precision and at larger scales in single-precision would improve forecast accuracy for a given computational cost.

Stochastic Processor model types SP1 and 2, where the large-scale variables are protected from bit-flips, performed well across all metrics, but they were estimated to be around twice as expensive as SH (see Table 3). Introducing low-probability bit-flips in the X -variables (types SP3–5) greatly reduced the accuracy for no cost




Figure 8. Vector plots for the (a) $F = 20$ and (b) $F = 10$ set-ups' truth time series. At each time step, the system is at a different point in Principal Component (PC) space. Across all time steps, these plots show the average direction in which all points in PC space within a radius of 0.5 Principal Component Units from each arrow's tail move in 1 time step. The size of the arrow reflects the average magnitude of the motion. The two centres of rotation in the $F = 10$ plot correspond to Regimes A (clockwise) and B (anticlockwise). At regime centres, the system's behaviour is so persistent that the arrows disappear entirely. The $F = 20$ plot does not demonstrate the presence of such clear regimes.

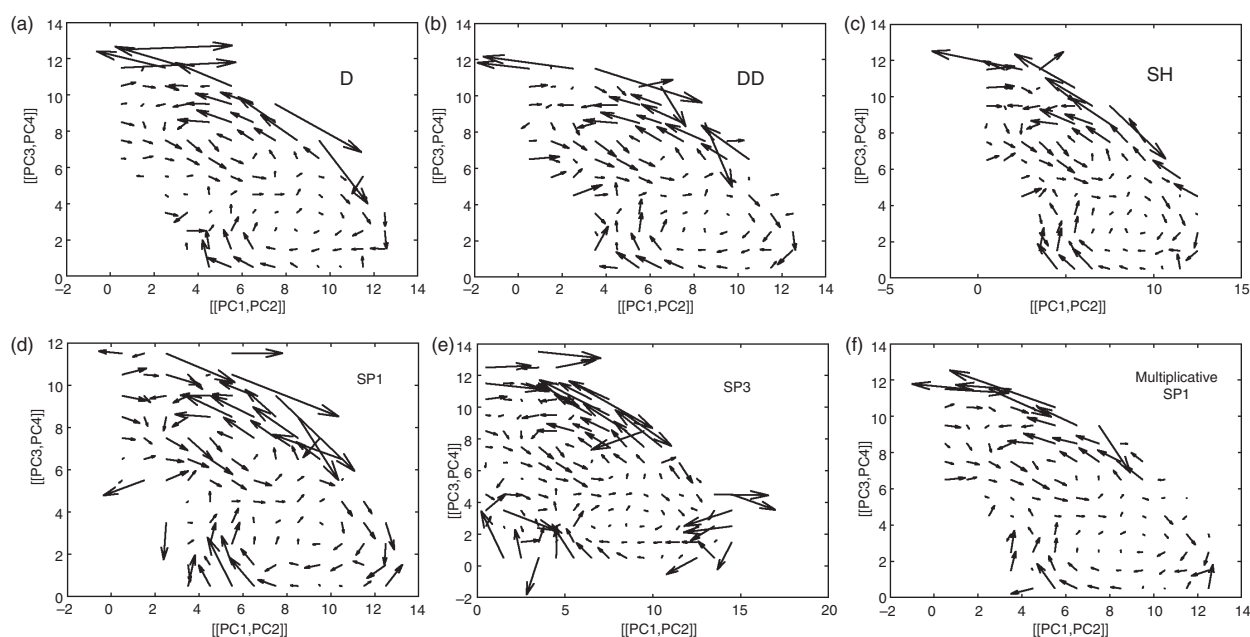


Figure 9. (a)–(f) Vector plots for selected models (labelled) of the $F = 10$ set-up under additive stochastic parametrization and model SP1 under the multiplicative scheme. Plots for all other multiplicative and deterministic models were very similar to the additive scheme equivalents and are not shown. The regimes are visible but not as dominant as the 'true' regimes in all cases shown except SP3. In the additive SP1 model, Regime B is beginning to break down: there are large arrows contrary to the anticlockwise motion around the regime centre. The contrary arrows are much smaller for the multiplicative SP1 model, but larger for the additive model SP3 where regime behaviour breaks down more completely.

saving, demonstrating the importance of a multi-scale approach. SP2 and SP4, which differ from SP1 and SP3 only in having a larger probability of bit-flips in the Y -variables, occasionally crashed in the longer runs, showing that the bit-flip probability must be kept small even for variables close to the truncation scale.

It is hoped that our results may inform future studies that scale-selectively reduce precision in more complex atmospheric models. Real-Earth models contain many more variables and may therefore be more likely to crash in reduced precision, which should hence be applied with care. But if our findings are indicative of how reduced precision affects such models, they provide an impetus for forecasters not only to more widely exploit already-existing single-precision architectures, but also to demand half-precision capability. Computational constraints mean that the resolution of operational weather forecast models is often decreased at long lead-times at the expense of increased error (Buizza, 2010); future studies might investigate whether reducing precision instead would improve medium-term forecasts. Our findings suggest that, although it by no means solves all problems on the path to cloud-resolving global forecasts, scale-selective

reduced precision could play an important part in the future of weather and climate forecasting.

Acknowledgements

The authors would like to thank Hannah Christensen (née Arnold) for her kind advice with regard to skill scores and stochastic parametrization; Andrew Dawson for his help, advice and code regarding reduced precision emulators; and Martin Schreiber and an anonymous reviewer for their constructive comments, which have helped to substantially improve the article. Tobias Thornes is funded by an NERC grant (number NE/L002612/1); Tim Palmer and Peter Düben by an ERC grant (Towards the Prototype Probabilistic Earth-System Model for Climate Prediction, project reference 291406).

Appendices

Appendix A: Computation of error doubling times

We computed error doubling times following the method of Lorenz (2006a). Here, fifty pairs of 5 MTU runs of the truth

system with random, independent initial conditions are carried out. Each pair consists of one unperturbed run and one in which the initial conditions are randomly changed (independently for each of the 8 X -variables, the 64 Y -variables or the 512 Z -variables as appropriate) by up to $\pm 0.1\%$. If the values of the X -variables, for example, at each time step n for runs with and without a perturbation are denoted X'_{kn} and X_{kn} respectively, the 'error' in each variable is:

$$e_{kn} = X'_{kn} - X_{kn}. \quad (\text{A1})$$

The square of this error is then averaged across all eight X -variables and the logarithm of this is averaged over all fifty runs to define $\log(E^2(t))$, from which $E(t)$ is calculated. The error doubling time t_d is computed from the gradient $\Delta(\log(E))/\Delta t$ of the initial exponential growth of $E(t)$, according to,

$$t_d = \log(2) \frac{\Delta t}{\Delta(\log(E))}. \quad (\text{A2})$$

A similar procedure can be followed for the error doubling times in the Y - and Z -variables.

Appendix B: Computation of Lyapunov exponents

To compute leading Lyapunov exponents, we compare a single, long run of the system to an equivalent system that is perturbed using the 'breeding vector' method (Toth and Kalnay, 1993) to ensure that the perturbed system always remains on the system's attractor. Following the initial spin-up, all variables in the system are perturbed by random amounts, with mean 0 and standard deviation 0.01; for the Y - and Z -variables the size of perturbations is reduced in proportion to their average magnitudes relative to the X -variables. The overall initial error magnitude $E(0)$ is:

$$E(0) = \frac{\sqrt{\sum_i (X_i - X'_i)^2 + \sum_{ij} (Y_{ij} - Y'_{ij})^2 + \sum_{ijk} (Z_{ijk} - Z'_{ijk})^2}}{K + J + I}, \quad (\text{B1})$$

where the non-primed and primed quantities are 'truth' and 'perturbed' values respectively.

The truth and perturbed systems are then evolved for 10 000 time steps. Every ten time steps (an interval chosen by trial and error), the perturbed system is rescaled so that the overall error magnitude, $E(10)$, matches that of the initial perturbation, $E(0)$. Except in the first 1000 time steps (which are excluded to make the results independent of the exact perturbation used), the error magnitudes are measured at the end of each ten-time-step interval prior to rescaling and their average over all intervals $E(10)$ is computed. Under the assumption that the average error $E(t)$ grows according to:

$$E(t) = E(0) \exp(\lambda t), \quad (\text{B2})$$

the Lyapunov exponent λ can be estimated by setting $t = 10$. We repeat this process five times and average the results. The procedure is the same for the one-tier and two-tier systems, except that only the X - or X - and Y -variables are included respectively.

Supporting information

The following supporting information is available as part of the online article:

File S1. The Fortran code used for our experiments and additional Principal Component plots (Figures S1 and S2).

Figure S1. Principal Component pdfs for the $F = 20$ (left) and $F = 10$ (right) set-ups' truth time series. Two regimes are clearly distinguishable only in the less chaotic $F = 10$ case.

Figure S2. Principal Component pdfs are plotted for selected models (labelled) of the $F = 10$ setup under additive stochastic and model SP1 under multiplicative stochastic parameterisation, the same cases as in Figure 9. Plots for all other multiplicative and deterministic models were very similar to the additive scheme equivalents and are not shown. Beneath each plot is shown the difference from the 'truth' pdf (Figure S1, right). The plots have consistent colour scaling to enable a fair comparison between different models, causing the SP3 difference from truth plot to saturate.

References

- Arnold HM. 2013. 'Stochastic parameterisation and model uncertainty'. DPhil thesis, University of Oxford.
- Arnold HM, Moroz IM, Palmer TN. 2013. Stochastic parametrizations and model uncertainty in the Lorenz '96 system. *Philos. Trans. R. Soc. A* **371**: 20110479, doi: 10.1098/rsta.2011.0479.
- Baede APM, Dent D, Hollingsworth A. 1976. *The Effect of Arithmetic Precision on Some Meteorological Integrations*. ECMWF: Reading, UK.
- Brown N, Weiland M, Hill A, Shipway B, Maynard C, Allen T, Reznay M. 2015. 'A highly scalable Met Office NERC cloud model.' In *Proceedings of the 3rd International Conference on Exascale Applications and Software*, 21–23 April 2015, Edinburgh, UK, pp. 132–137.
- Buizza R. 2010. The value of a variable resolution approach to numerical weather prediction. *Mon. Weather Rev.* **138**: 1026–1042, doi: 10.1175/2009MWR3077.1.
- Buttari A, Dongarra J, Kurzak J, Luszczek P, Tomov S. 2008. Using mixed precision for sparse matrix computations to enhance the performance while achieving 64-bit accuracy. *ACM Trans. Math. Softw.* **34**: 22, doi: 10.1145/1377596.1377597.
- Christensen HM, Moroz IM, Palmer TN. 2015. Simulating weather regimes: Impact of stochastic and perturbed parameter schemes in a simple atmospheric model. *Clim. Dyn.* **44**: 2195–2214, doi: 10.1007/s00382-014-2239-9.
- Dawson A, Palmer TN. 2015. Simulating weather regimes: Impact of model resolution and stochastic parameterisation. *Clim. Dyn.* **44**: 2177–2193, doi: 10.1007/s00382-014-2238-x.
- Dennard RH, Gaensslen FH, Yu H-N, Rideout VL, Bassous E, LeBlanc AR. 1974. Design of ion-implanted MOSFET's with very small physical dimensions. *IEEE J. Solid-State Circuits* **9**: 256–268, doi: 10.1109/JSSC.1974.1050511.
- Düben PD, Palmer TN. 2014. Benchmark tests for numerical weather forecasts on inexact hardware. *Mon. Weather Rev.* **142**: 3809–3829, doi: 10.1175/MWR-D-14-00110.1.
- Düben PD, McNamara H, Palmer TN. 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.
- Düben PD, Russell FP, Niu X, Luk W, Palmer TN. 2015. On the use of programmable hardware and reduced numerical precision in earth system modelling. *J. Adv. Model. Earth Syst.* **7**: 1393–1408, doi: 10.1002/2015MS000494.
- Gan L, Fu H, Luk W, Yang C, Xue W, Huang X, Zhang Y, Yang G. 2013. 'Accelerating solvers for global atmospheric equations through mixed-precision data flow engine'. In *23rd International Conference on Field Programmable Logic and Applications*, 2–4 September 2013. IEEE: Porto, Portugal, pp. 1–6.
- Goddeke D, Strzodka R, Turek S. 2007. Performance and accuracy of hardware-oriented native-, emulated- and mixed-precision solvers in FEM simulations. *Int. J. Parallel Emergent Distrib. Syst.* **22**: 221–256, doi: 10.1080/17445760601122076.
- Hansen J, Kharecha P, Sato M, Masson-Delmotte V, Ackerman F, Beerling DJ, Hearty PJ, Hoegh-Guldberg O, Hsu S-L, Parmesan C, Rockstrom J, Rohling EJ, Sachs J, Smith P, Steffen K, Van Susteren L, von Schuckmann K, Zachos JC. 2013. Assessing 'dangerous climate change': Required reduction of carbon emissions to protect young people, future generations and nature. *PLoS One* **8**: e81648, doi: 10.1371/journal.pone.0081648.
- Hughes I, Hase T. 2010. *Measurements and their Uncertainties* (1st edn). Cambridge University Press: Cambridge, UK.
- IEEE. 2008. 'IEEE standard for floating-point arithmetic'. IEEE Standard 754-2008, pp. 1–70.
- Karakonstantis G, Roy K. 2011. 'Voltage over-scaling: A cross-layer design perspective for energy efficient systems'. In *20th European Conference on Circuit Theory and Design*, 29–31 August 2011. Linköping: Sweden, pp. 548–551, doi: 10.1109/ECCTD.2011.6043592.
- Leith CE. 1971. Atmospheric predictability and two-dimensional turbulence. *J. Atmos. Sci.* **28**: 145–161, doi: 10.1175/1520-0469(1971)028<0145:APATDT>2.0.CO;2.
- Lorenz E. 1969. The predictability of a flow which possesses many scales of motion. *Tellus* **21**: 289–307, doi: 10.3402/tellusa.v21i3.10086.

- Lorenz E. 1993. *The Essence of Chaos*. University of Washington Press: Seattle, WA.
- Lorenz E. 2006a. Predictability – a problem partly solved. In *Predictability of Weather and Climate*, Palmer TN, Hagedorn R. (eds.): 40–58. Cambridge University Press: Cambridge, UK.
- Lorenz E. 2006b. Regimes in simple systems. *J. Atmos. Sci.* **63**: 2056–2073, doi: 10.1175/JAS3727.1.
- Markov IL. 2014. Limits on fundamental limits to computation. *Nature* **512**: 147–154.
- Oberstar EL. 2007. 'Fixed-point arithmetic & fractional math'. Superkits Whitepapers. <http://www.superkits.net/whitepapers/FixedPointRepresentation&FractionalMath.pdf> (accessed 1 November 2016).
- Palem K. 2014. Inexactness and a future of computing. *Philos. Trans. R. Soc. A* **372**: 20130281, doi: 10.1098/rsta.2013.0281.
- Palem K, Lingamneni A. 2013. Ten years of building broken chips: The physics and engineering of inexact computing. *ACM Trans. Embedded Comput. Syst.* **12**: 1–23, doi: 10.1145/2465787.2465789.
- Palmer TN. 2012. Towards the probabilistic earth system simulator: A vision for the future of climate and weather prediction. *Q. J. R. Meteorol. Soc.* **138**: 841–861, doi: 10.1002/qj.1923.
- Palmer TN. 2014. Build high-resolution global climate models. *Nature* **515**: 338–339.
- Palmer TN, Richardson D. 2014. *Decisions, decisions . . . !*. *ECMWF Newsl.* **141**: 12–13.
- Skamarock WC. 2004. Evaluating mesoscale NWP models using kinetic energy spectra. *Mon. Weather Rev.* **132**: 3019–3032, doi: 10.1175/MWR2830.1.
- Sloan J, Kesler D, Kumar Rakesh RA. 2010. 'A numerical optimisation-based methodology for application robustification: Transforming applications for error tolerance'. In *IEEE/IFIP International Conference on Dependable Systems and Networks*, pp. 161–170.
- Snowdon DC, Ruocco S, Heiser G. 2005. 'Power management and dynamic voltage scaling: Myths and facts'. In *Proceedings of the 2005 Workshop on Power Aware Real-time Computing*, 22 September 2005. Jersey City, NJ, Vol. 12, pp. 1–7.
- Toth Z, Kalnay E. 1993. Ensemble forecasting at NMC: The generation of perturbations. *Bull. Am. Meteorol. Soc.* **74**: 2317–2330, doi: 10.1175/1520-0477(1993)074<2317:EFANTG>2.0.CO;2.
- Trader T. 2016. 'NVIDIA unleashes monster Pascal GPU card at GTC16'. *HPC Wire*. <https://www.hpcwire.com/2016/04/05/nvidia-monster-pascal-gpu-card-gtc16/> (accessed 11 April 2016).
- Váňa F, Düben PD, Lang S, Palmer TN, Leutbecher M, Salmond D, Carver G. 2016. Single precision in weather forecasting models: An evaluation with the IFS. *Mon. Weather Rev.*, doi: 10.1175/MWR-D-16-0228.1.
- Wilks D. 2005. Effects of stochastic parameterisations in the Lorenz '96 system. *Q. J. R. Meteorol. Soc.* **131**: 389–407, doi: 10.1256/qj.04.03.
- Wilks D. 2006. *Statistical Methods in the Atmospheric Sciences* (2nd edn). Elsevier, Academic Press: Cambridge, MA.