Reduced precision computing for weather and climate models

Leo Saffin, Jan Ackmann, Matthew Chantry, Sam Hatfield, Milan Klöwer, Andrew McRae

University of Oxford

26 September 2018

Floating-point numbers

$$x = \pm 1.s_1 s_2 \ldots s_{\text{bits}} \times 2^{\text{exponent}-\text{bias}}$$

	Half	Single	Double
Decimal Accuracy $log_{10}(2^{N_{\text{sbits}}+1})$	3.3	7.2	16.0

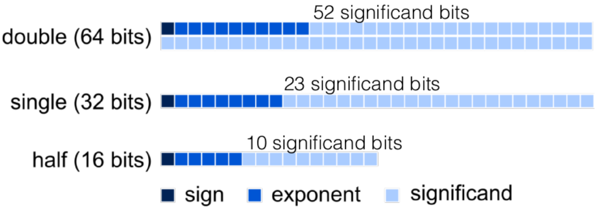

We should use no more precision than is necessary.

If the errors due to rounding is negligible compared to the inherent model uncertainty then the precision is sufficient.

Model uncertainty is explicitly represented in ensemble forecasts

The gridpoint physics tendencies are modified by a random field r between -1 and 1
SPPT: $P = (1 + \mu r) \times \sum P_i$

Single-precision models

Until recently, models typically just use double precision.
Recently single-precision versions of models have emerged.

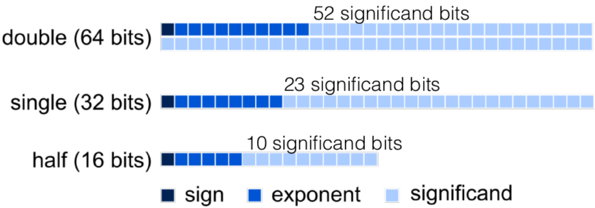

- MeteoSwiss - COSMO
  - Full forecast model ($\approx$ 40% reduction in runtime)
- ECMWF - IFS
  - Full forecast model ($\approx$ 40% reduction in runtime)
- Met Office UM
  - Pressure solve
    - Requires more iterations to converge
    - Net improvement in runtime due to efficiency of single precision
  - Large-scale precipitation scheme (15-50% reduction in runtime)

Beyond single precision

Published work on reduced precision in meteorological models

- Reduced precision in simple and intermediate complexity models - Peter Düben
- Superparametrization in OpenIFS single-column model - Peter Düben
- Scale selective reduced precision - Tobias Thornes
- Data assimilation. Sam Hatfield
- Reduced precision models implemented on FPGAs. Lorenz63 (Jeffress et al. 2017), Lorenz96 (Düben at al. 2015)

This talk

- Physics schemes - Leo Saffin (Me)
- Spectral computations - Matthew Chantry
- Legendre transforms - Sam Hatfield & Matthew Chantry
- Preconditioner - Jan Ackmann
- Adjoint model - Andrew McRae
- Posits - Milan Klöwer

Reduced-precision emulator (Dawson and Düben 2017)

Defines new Fortran type
- **rpe_var**(*double* value, *int* sbits)

and function
- **apply_truncation**(*double* value, *int* sbits)

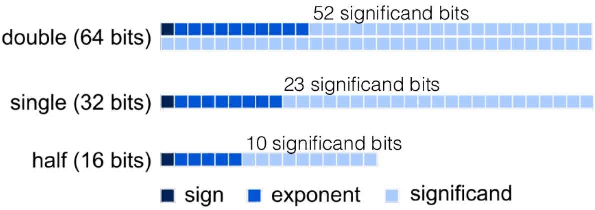

- The integer allows precision to be set on a variable-by-variable basis
  - Alternatively the parameter **rpe_default_sbits** is used when sbits is unset
- The double precision number provides the underlying representation of the reduced precision number
  - Each operation on an **rpe_var** is applied normally to the double-precision number
  - After each operation **apply_truncation** is used on the double-precision number

SPEEDY

- SPEEDY. T30L8 spectral model with simplified physics
- Reduce precision in each physics scheme individually
- Compare "forecasts" with a double-precision reference
- Nonzero forecast error at any precision
- Range of precisions with interchangeable forecast errors

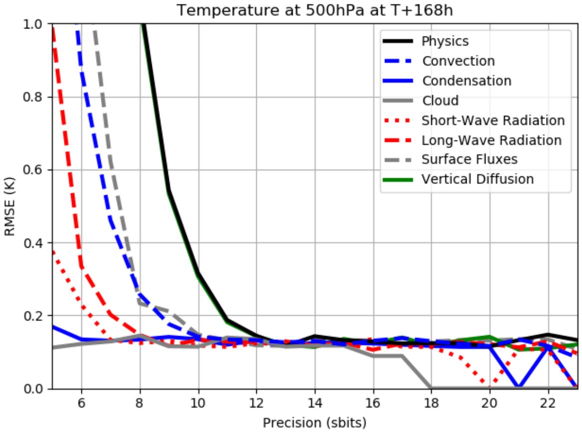

SPEEDY SPPT

SPPT off

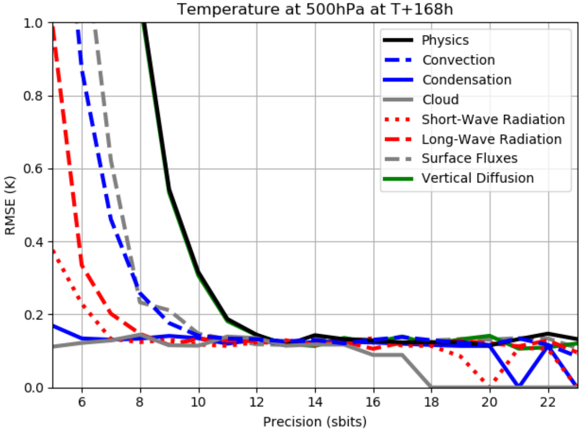

SPPT on $(P = (1 + \mu r) \sum P_i)$

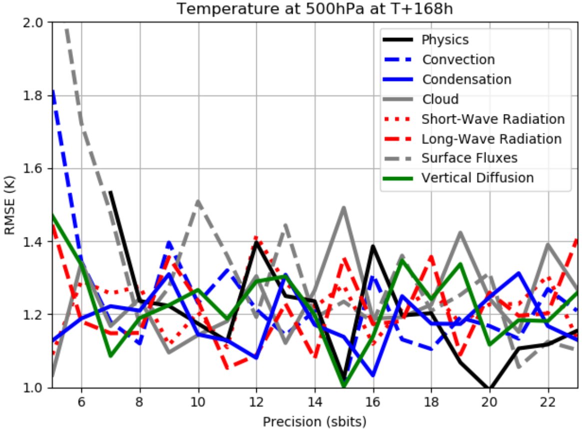

Error growth due to stochastic physics is much greater than those introduced by rounding errors at intermediate precisions and also masks errors that may have seemed unacceptable from the "deterministic" runs

OpenIFS

SPEEDY

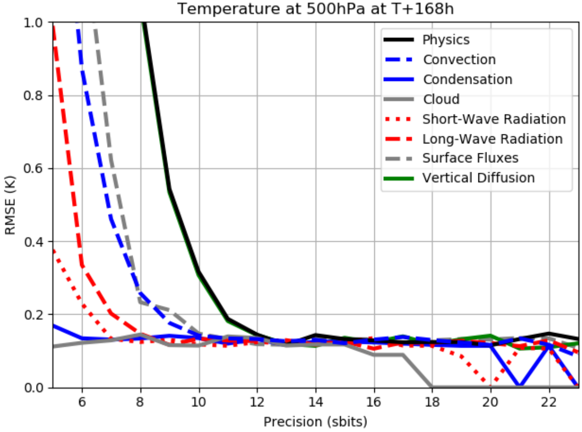

OpenIFS (T21L60)

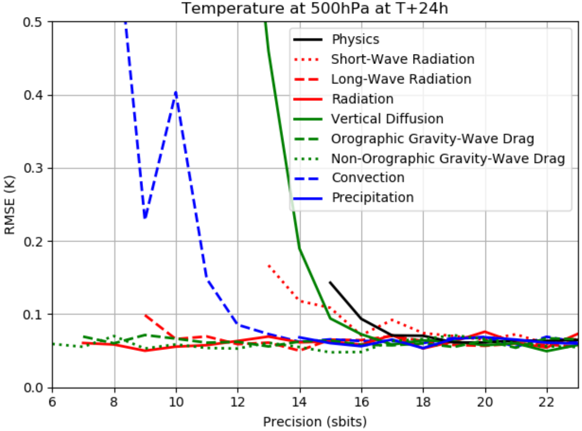

OpenIFS has much more complex physics schemes than SPEEDY which gives a wider variety in lowest acceptable precisions

Precision errors as noise

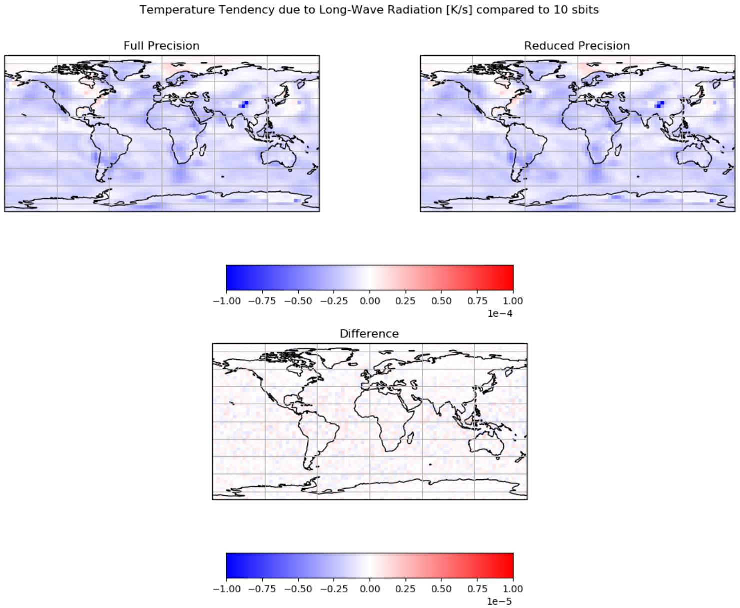

Precision errors as noise

SPPT: $P = (1 + \mu r) \sum P_i$

$$P_{lw}(hp) - P_{lw}(dp)$$

$$r = \frac{P_{lw}(hp)}{P_{lw}(dp)} - 1$$

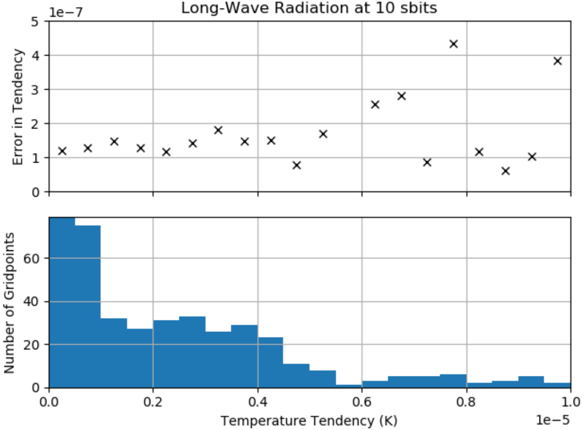

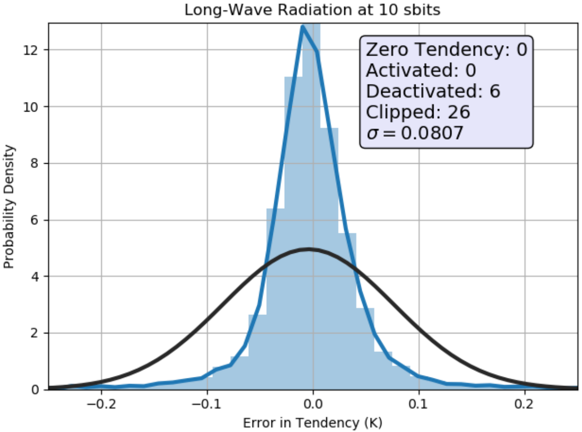

Rounding errors in long-wave radiation do not act like SPPT. Error is not proportional to tendency and relative error is not Gaussian.

Precision errors as switch

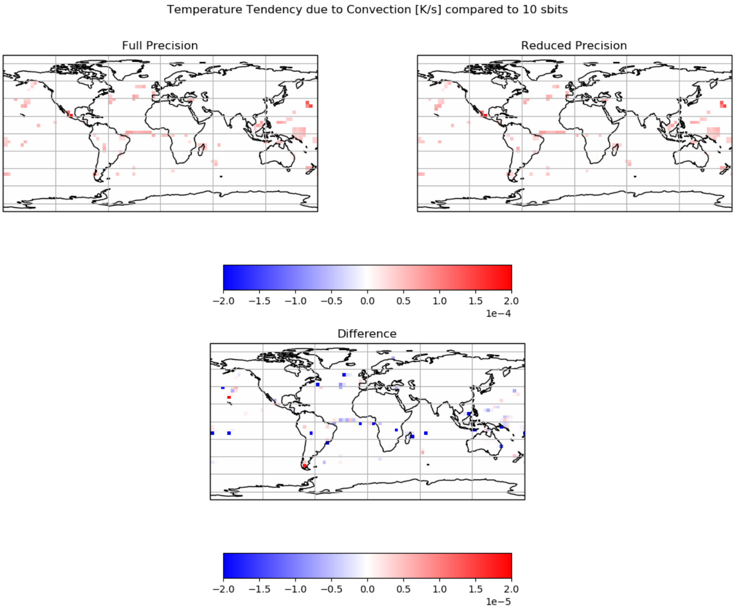

Precision errors as switch

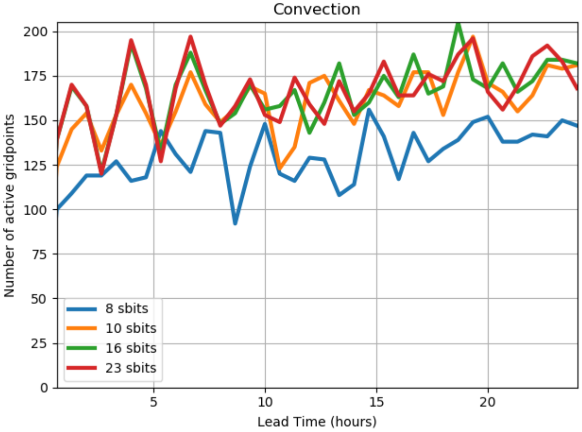

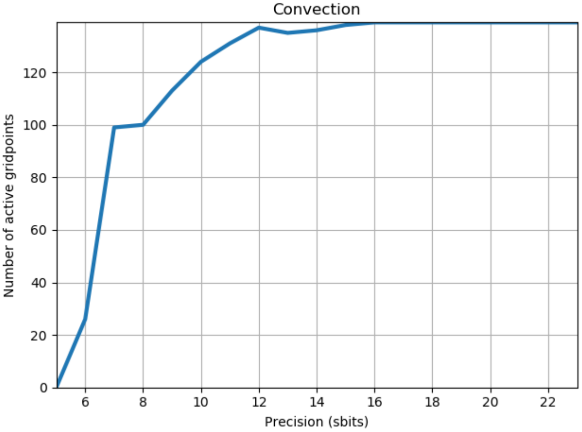

Rounding errors can modify the triggering of convection and convection is less likely to trigger as precision is reduced.

Spectral dynamical core schematic

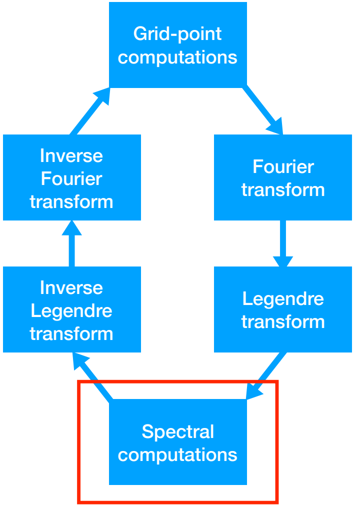

What we've done

- Reduced precision calculations in spectral-space only.
- Spectral transforms and grid-point calculations in double precision.

Will ...

- introduce rounding errors to prognostic variables: vorticity, temperature etc.

Won't ...

- cover all algorithmic error propagation

Why spectral space?

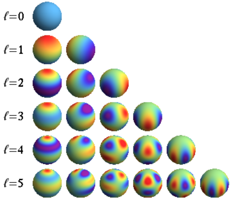

- Spectral models represent fields as a sum of modes representing different lengthscales.

- Can we reduce precision when calculating the small scales?

- This is appealing due to the high inherent uncertainty in small scale dynamics (parametrisation, viscosity, data-assimilation,...).

L2-norm — Global precision

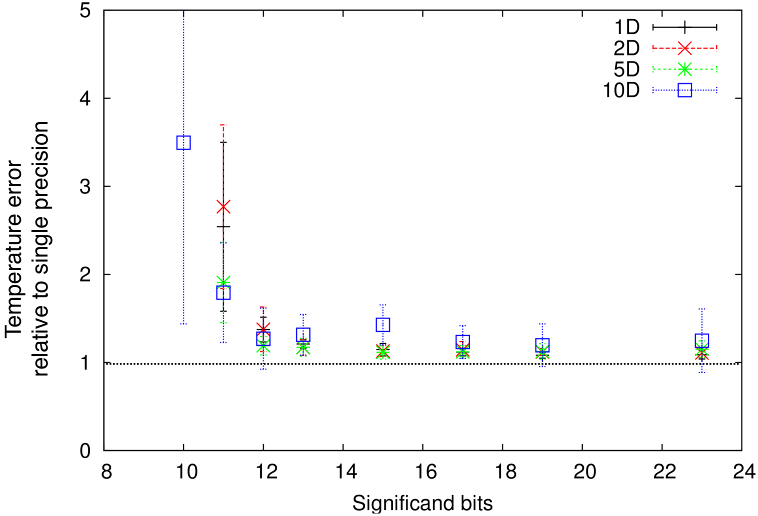

Half-precision from which wavenumber?

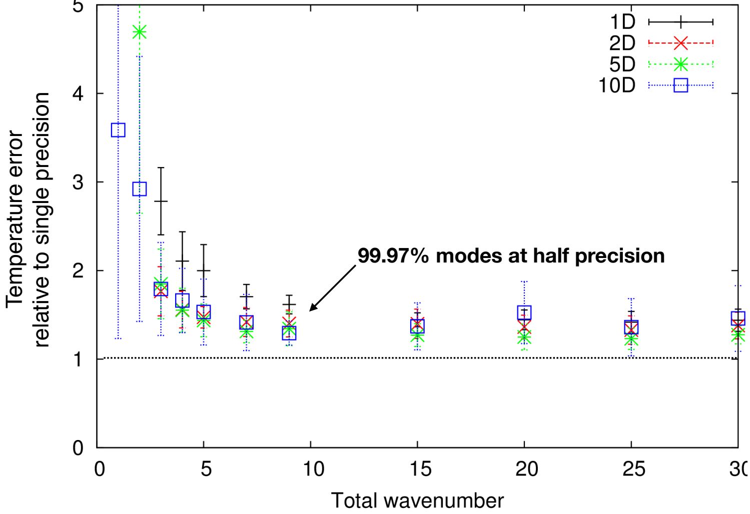

RMSE of several 10 day T511 OpenIFS forecasts, verified against a double-precision control.

- half_trans_10: Half-precision transforms with first 10 modes in double precision.
- NVIDIA's tensor cores allow mixed half/single precision GEMMs and they claim a speed-up factor of 16x.

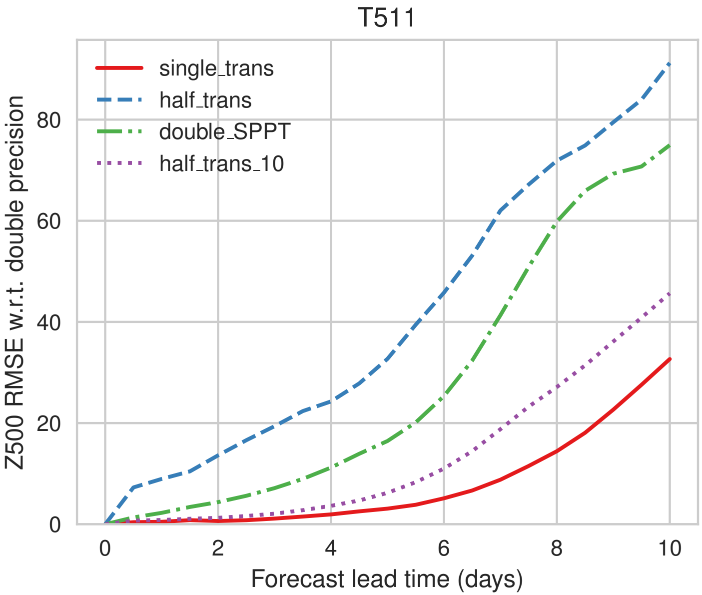

Problem setup (MIT GCM)

- Not 4DVar, but similar: find time-independent surface heat flux Q^{netm} to make model climatology $\bar{T}$ consistent with dataset $\bar{T}^{\text{obs}}$. Cost function

$$J = \lambda_1 \cdot \underbrace{\frac{1}{N_1} \sum_{i=1}^{N_1} \left[\frac{\bar{T}_i - \bar{T}_i^{\text{obs}}}{\sigma_i^T} \right]^2}_{\text{mismatch to observations}} + \lambda_2 \cdot \underbrace{\frac{1}{N_2} \sum_{i=1}^{N_2} \left[\frac{Q_i^{\text{netm}}}{\sigma_i^Q} \right]^2}_{\text{magnitude of heat flux}}$$

- $\bar{T}_i^{\text{obs}}$, σ_i^T, σ_i^Q based on observational data set
- $4^\circ \times 4^\circ$ resolution, 15 layers (2315 Q^{netm} dofs – surface only), run for one year
- Quasi-Newton optimisation (M1QN3)
- Precision:
  - Forward model: full
  - Adjoint model: reduced**
  - Optimisation: full

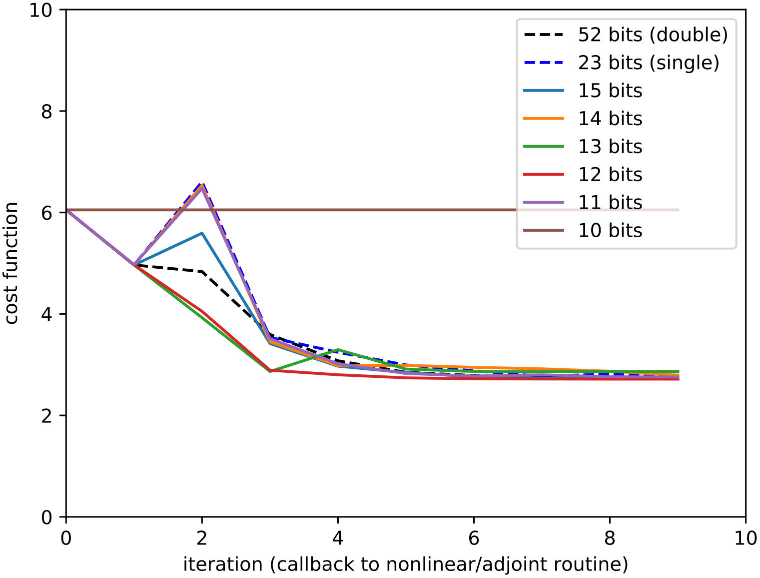
10
8
6
4
2
0
cost function
0
2
4
6
8
10
iteration (callback to nonlinear/adjoint routine)
52 bits (double)
23 bits (single)
15 bits
14 bits
13 bits
12 bits
11 bits
10 bits

Solver Iterations for 0.65°

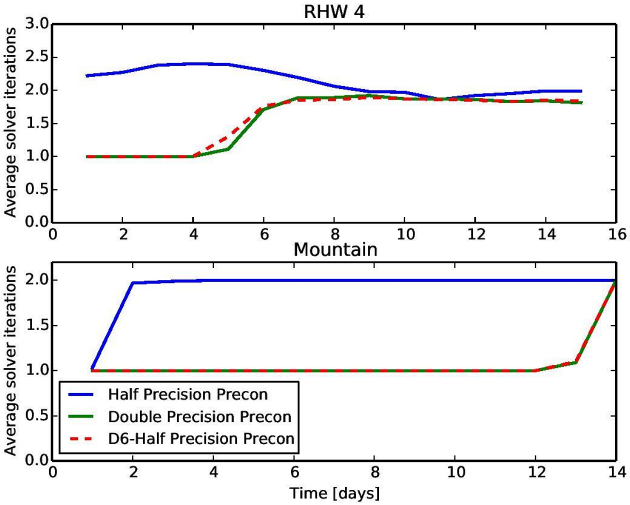

Bit-wise information content in geostrophic turbulence

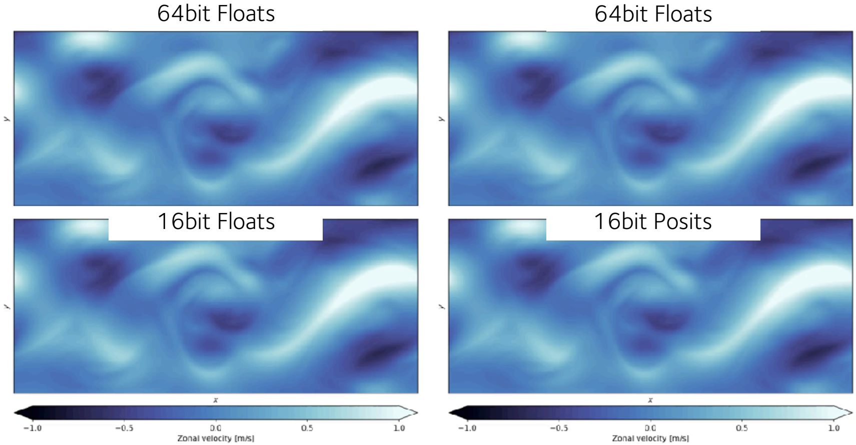

Worst-case decimal precision

Half precision (16bit)

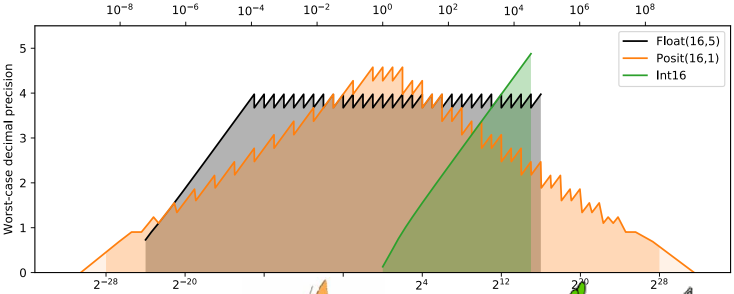

Bart Simpson

vs

Son Goku

vs

Johnny Bravo

Bit-wise information content in Lorenz 63

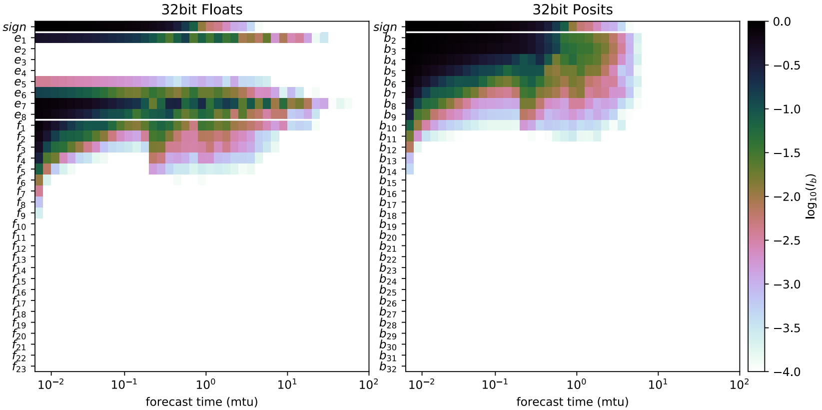

Summary

- Some operations need to be kept at higher precision but most operations can still be reduced
- Precision can be reduced further if model uncertainty is taken into account
- The precision needed for spectral calculations is dependent on scale
- Errors in physics schemes are more complex than random noise
- Rate of convergence and number of iterations are more important for optimization-type problems
- Posit is a more efficient use of bits than floating point