REQUEST FOR A SPECIAL PROJECT 2015–2017

MEMBER STATE:	United Kingdom …….…………………………………………...…
Principal Investigator[1]:	Tim Palmer
Affiliation:	University of Oxford,
Address:	Atmospheric, Oceanic and Planetary Physics, Clarendon Laboratory, Parks Road, Oxford, OX1 3PU.
E-mail:	t.n.palmer@atm.ox.ac.uk
Other researchers:	Hannah Christensen (nee Arnold) Antje Weisheimer

Project Title:	Constraining stochastic parametrisation schemes through coarse graining

If this is a continuation of an existing project, please state the computer project account assigned previously.	**SP** ________________	
Starting year: (Each project will have a well defined duration, up to a maximum of 3 years, agreed at the beginning of the project.)	2015	
Would you accept support for 1 year only, if necessary?	YES ☒	NO ☐

Computer resources required for 2015-2017: (The maximum project duration is 3 years, therefore a continuation project cannot request resources for 2017.)	**2015**	**2016**	**2017**
High Performance Computing Facility (units)	2,000,000	3,000,000	4,000,000
Data storage capacity (total archive volume) (gigabytes)	2,400	3,600	4,800

*An electronic copy of this form **must be sent** via e-mail to:* *special_projects@ecmwf.int*

Electronic copy of the form sent on (please specify date):

23 June 2014……..........……………

Continue overleaf

[1] The Principal Investigator will act as contact person for this Special Project and, in particular, will be asked to register the project, provide an annual progress report of the project's activities, etc.

Principal Investigator: Tim Palmer

Project Title: Constraining stochastic parametrisation schemes through coarse graining

Extended abstract

It is expected that Special Projects requesting large amounts of computing resources (500,000 SBU or more) should provide a more detailed abstract/project description (3-5 pages) including a scientific plan, a justification of the computer resources requested and the technical characteristics of the code to be used. The Scientific Advisory Committee and the Technical Advisory Committee review the scientific and technical aspects of each Special Project application. The review process takes into account the resources available, the quality of the scientific and technical proposals, the use of ECMWF software and data infrastructure, and their relevance to ECMWF's objectives. - Descriptions of all accepted projects will be published on the ECMWF website.

SPPT and Coarse Graining

Stochastically Perturbed Parametrisation Tendencies (SPPT) is an attractive stochastic parametrisation scheme due to its ease of use and beneficial impact on ensemble forecast reliability. It runs in conjunction with operational physics parametrisation schemes, so can be quickly adapted for use in different models. It was originally developed for, and is in is use operationally in, the IFS (ECMWF), and testing is underway for implementation in COSMO (MeteoSwiss, USAM). SPPT is also available for use in climate models, including CESM (NCAR).

The SPPT scheme addresses model uncertainty due to the parametrisation process by perturbing the parametrised physics tendencies using multiplicative noise. SPPT perturbs the sum of the parametrised tendencies:

$$T = \frac{\partial X}{\partial t} = D + K + (1+e)\sum_{i=1}^{5} P_i$$

where T is the total tendency in X. D is the tendency from the dynamics, K is horizontal diffusion, P_i is the tendency from the ith physics scheme, and e is the zero mean random perturbation. The scheme perturbs the tendency for four variables: T, U, V and q. Each variable tendency is perturbed using the same random number field. The perturbation field is generated using a spectral pattern generator. The pattern at each time step is the sum of three independent random fields with horizontal correlation scales of 500, 1000 and 2000 km. These fields are evolved in time using an AR(1) process on time scales of 6 hours, 3 days and 30 days respectively, and the fields have standard deviations of 0.52, 0.18 and 0.06 respectively. It is expected that the smallest scale (500 km and 6 hours) will dominate at a 10 day lead time – the larger scale perturbations are important for monthly and seasonal forecasts.

Despite its popularity, the SPPT remains ad hoc in its assumptions. The imposed spatial and temporal correlations have not been derived from theory or observation, and have simply been tuned to give the best results. However, the magnitude and use of multiplicative noise were retrospectively justified using coarse graining studies. Shutts and Palmer (2007) defined an idealised cloud resolving model (CRM) simulation as truth. The resultant fields and their tendencies were then coarse grained to the resolution of a NWP model to study the sub-grid scale variability which a stochastic parametrisation seeks to represent. The effective heating function for the nth coarse grid box, Q_n, was calculated by averaging over nine fine grid boxes. This was compared to the heating calculated from a convective parametrisation scheme, $\overline{Q1} = Q1\left(\overline{X}\right)$, where $\overline{X}$ represents the coarse grained CRM fields. The validity of the multiplicative noise in the SPPT scheme was analysed by studying histograms of Q_n conditioned on different ranges of $\overline{Q1}$. The mean and standard deviation of Q_n is observed to increase as a function of $\overline{Q1}$ providing some support for the SPPT scheme.

Secondly, SPPT does not distinguish between the different parametrisation schemes. It is likely that the parametrisation schemes could have very different error characteristics so this assumption may not be valid. A recent coarse graining study by Shutts and Pallares (2014) has measured the standard deviation of the error for each physics tendency as a function of the parametrised tendency. The tendencies from the IFS at T1279 (16km) were defined to be "truth", and were compared to forecast tendencies from the IFS at T159 (130km). Both resolutions were coarse grained in time (using a 12 hour triangular filter) and in space (to a 500km grid), before the mean and standard deviation of the T1279 tendencies conditioned on the the T159 tendencies were calculated. The study revealed that the different schemes do, indeed, have very different error characteristics, with the uncertainty in the cloud and convection tendencies being much larger than the radiation tendency. Additionally, the standard deviation of the cloud and convection tendencies were shown to be proportional to the square root of the parametrised tendency, unlike the representation in SPPT.

Thirdly, SPPT assumes that the errors from each physics parametrisation scheme are perfectly correlated – one random number field is used to perturb all schemes. It is very unlikely that uncertainties in the different processes are precisely correlated, as modelled by SPPT. An alternative is to use independent random fields for each physics tendency. This "independent SPPT" (SPPTi) was considered by Arnold (2013) using computer resources at ECMWF as part of special project spgbtpuc (ending in December 2014) – a significant improvement in the reliability of ensemble forecasts was observed due to a correction of the under-dispersive ensemble spread. However, the scheme also increased the error in the ensemble mean for certain fields (T850) and resulted in over-dispersive ensembles for other fields (U850). It is also unlikely that the assumption of SPPTi, namely that the errors are precisely uncorrelated, is true. The different physics schemes in the IFS have been developed in tandem, and are called sequentially in the IFS to maintain a balance. If the IFS physics schemes have been closely tuned to each other, potentially with compensating errors, decoupling the two schemes by using independent perturbations could reduce the forecast skill and introduce errors into the forecasts.

The previous coarse graining experiments described above show the potential power of the technique for constraining stochastic parametrisation schemes. However, both studies have used fairly coarse atmospheric models as "truth", so can only provide a lower bound on the error statistics. Shutts and Pallares (2014) have assumed that the imposed SPPT spatial and temporal scales are correct, and coarse grained both resolutions to these scales before comparison. Additionally, the correlation of errors between the parametrisation schemes was not measured.

Questions of interest for this project are as follows:

1. Can we use high resolution model data sets as a more accurate "truth" for coarse graining experiments?
2. Is it possible to use these data to measure more characteristics of the stochastic noise term than have previously been measured?
    - spatial and temporal correlations?
    - correlation of errors between tendencies?
    - standard deviation?
3. How is this affected by ... ?
    - variable, location ...

Experimental Details

The project will consist of two phases.

Systematic Study of SPPT Correlations

Firstly, a systematic study will be performed testing the performance of SPPT with different imposed spatial and temporal correlations. Is it possible to find an optimal spatial and temporal correlation scale by scanning over parameter space and performing ensemble forecast experiments using each set of parameters in turn? This preliminary experiment will be an interesting test of the sensitivity of SPPT to the imposed correlations (to our knowledge, there has been little systematic testing of the sensitivity of SPPT in this way). It will show the degree to which the SPPT scheme can be improved using such "brute force" techniques. In addition, it will indicate the level of sensitivity which can be expected when we move on to the second phase of the project detailed below – deriving physically motivated spatial and temporal correlations, which can then be introduced into the SPPT scheme for evaluation.

Coarse Graining Experiments

Secondly, a similar technique will be followed as used in Shutts and Palmer (2007). Since the aim is to constrain SPPT for use operationally in the IFS, the high resolution data will be coarse grained to the current resolution of the IFS ensemble prediction system (EPS) – T639 or approximately 30km (alternatively, we could consider grid boxes at T1279 = 12km, which will be the operational resolution of the IFS EPS by Autumn 2014).

The high resolution data to be used as "truth" was produced as part of the CASCADE project. A limited area version of the Met Office's Unified Model was run for the Indian Ocean and West Pacific over a domain 4450km by 14950km for ten days (Holloway and Woolnough, 2013). The model was integrated at a variety of resolutions, including 1.5km and 4km. At 1.5km and 4km there is explicit convection – the convection parametrisation scheme was tuned down to only operate at low CAPE. It is these runs which we will use as our target or "truth" data set. Within the CASCADE project a similar set of simulations was also carried out for the West African monsoon region, detailed in Pearson et al (2013), which could provide an interesting second test case.

The following questions and experimental methods will be considered.

Aim: to measure the spatial and temporal correlations of the noise term in SPPT

As a proof of concept, these correlations will initially be measured for a single process. The first process considered will be the deterministic convection scheme. The methodology in Shutts and Palmer (2007) will be followed, using the coarse grained high-resolution dataset as input for the Bechtold convection parametrisation scheme. The spatial correlations of the error between the parametrised and "true" tendencies will be measured. Measuring the temporal correlation is also possible in this set-up, as feedbacks from the convection scheme onto the large scale state can be removed by considering the parametrised tendencies after just one time step, and always initialising the scheme from the 'true' initial conditions.

Aim: to measure the correlations of error between the different parametrisation schemes

Having performed the proof of concept experiments above, the focus of the project will shift to considering correlations of errors between the different parametrisation schemes. This is a more difficult aim. We measure the 'true' tendency, T, as the change in variables over a time step from the CRM data. The IFS (or our model of choice) is integrated forward for one time step, with the tendencies from each scheme, $T^*{}_i$ archived. The difference between the 'true' tendency and the sum of the parametrised tendencies, T^*, is measured as the total error.

$$T = D + \sum_i P_i \qquad \text{True tendency}$$

$$T^* = D^* + \sum_i P_i^* \qquad \text{Parametrised tendency}$$

$$T = T^* + \left(D - D^*\right) + \left(\sum_i P_i - \sum_i P_i^*\right)$$

$$T = T^* + e$$

It is assumed that the model dynamics is a good representation of the true dynamics (i.e. $D = D^*$), and that the only source of error is due to the parametrisation process. The error, e, must now be split up into its constituent errors from each parametrisation scheme in order to consider the correlations between errors. If it is assumed that multiplicative noise is the correct form of the noise, this provides additional information. Since the vertical tendency profiles for each scheme are known, the task is then to use the vertical parametrised tendency profiles as a basis set for expressing the total error.

$$T = D^* + \sum_i \left(T_i^* + e_i\right)$$

$$= D^* + \sum_i \left(1 + z_i\right) T_i^* \qquad \text{assuming multiplicative noise}$$

This fitting process will be carried out, and the value of z_i recorded for each scheme together with the residual. These will be analysed for:

> correlations between parametrisation schemes
> magnitude
> spatial and temporal correlation

Having measured the characteristics of the error term, this information will be incorporated into the EPS system, and the resultant forecasts tested against the operational system.

Technical requirements

The first phase of the project (systematically studying the effect of the imposed correlations in SPPT) will be computationally expensive, but can be carried out at low resolution in the first instance. At T159, 60,000 billing units is sufficient to perform a single experiment at T159 of 30, ten-day forecasts each with 50 ensemble members.

The second aim of this project is to measure the characteristics of the model error term in the IFS at operational EPS resolution of T639, necessitating integrations of the model at this high resolution. Prior to running the IFS at T639, it is expected that preliminary experiments will be carried out at lower resolution to test the methods used in the project.

For the (low) horizontal resolution of T159, 65 billing units are sufficient to perform a single fifteen-day forecast at T159, while the same experiment at T639 takes 945 billing units. It is anticipated that the data analysis aspect of the project will take up a large proportion of the requested resources, though it is difficult to estimate how this translates into billing units.

Having derived an error structure to be used in SPPT, we will implement this in the IFS. At T639, an experiment consisting of 15, ten-day forecasts each with 20 ensemble members costs 190,000 billing units.

References

Arnold, H. M., (2013). Stochastic parametrisation and model uncertainty. Ph.D. thesis, University of Oxford.

Holloway, C. E. and S. J. Woolnough, (2013). The effects of explicit versus parameterized convection on the MJO in a large-domain high-resolution tropical case study. Part I: Characterization of large-scale organization and propagation. J. Atmos. Sci., 70 (5), 1342–1369.

Pearson, K. J., G. M. S. Lister, C. E. Birch, R. P. Allan, R. J. Hogan, and S. J. Woolnough, (2013). Modelling the diurnal cycle of tropical convection across the ‘grey zone’. Q. J. Roy. Meteor. Soc., 140 (679), 491–499.

Shutts, G. J. and A. C. Pallares, (2014). Assessing parametrization uncertainty associated with horizontal resolution in numerical weather prediction models. Phil. Trans. R. Soc. A, 372 (2018).

Shutts, G. J. and T. N. Palmer, (2007). Convective forcing fluctuations in a cloud resolving model: Relevance to the stochastic parameterization problem. J. Climate, 20 (2), 187–202.