Stochastic parametrization of multi-scale processes using a dual grid and 'real-time computer games physics'

Glenn Shutts[1], **Thomas Allen and Judith Berner**

[1]*Met Office*
FitzRoy Road, Exeter, EX1 3PB, UK
glenn.shutts@metoffice.gov.uk

ABSTRACT

Some speculative proposals are made for extending current stochastic sub-gridscale parametrization methods using techniques adopted from the field of computer graphics and flow visualization. The idea is to emulate sub-filter-scale physical process organization and time evolution on a fine grid and couple the implied coarse-grained tendencies to a forecast model. A two-way interaction is envisaged so that fine grid physics (e.g. deep convective clouds) responds to forecast model fields. The fine grid model may be as simple as a two-dimensional cellular automaton or as computationally-demanding as a cloud-resolving model similar to the coupling strategy envisaged in 'super-parametrization'. Computer codes used in computer games and visualization software illustrate the potential for cheap but realistic simulation where emphasis is placed on algorithmic stability and visual realism rather than pointwise accuracy in a predictive sense. In an ensemble prediction context a computationally-cheap technique would be essential and some possibilities are outlined. Idealized proof-of-concept simulations are described which highlight technical problems such as the nature of the coupling.

1 Introduction

The inclusion of stochastic physical parametrization terms into ensemble prediction systems was motivated by insufficient member spread (relative to forecast error) and justified by the inherent statistical nature of the physical processes being parametrized (Buizza et al, 1999; Palmer,2001). Buizza et al (1999) devised a simple technique for perturbing the physical parametrization tendencies using a latitude/longitude 'tile' pattern of random numbers. The random number associated with each tile was chosen to lie between 0.5 and 1.5 with uniform probability density function (hereafter *pdf*) and the parametrization tendency for any gridpoint lying within the tile is multiplied by it. Positive impact on both spread and skill was found to be optimized for 10 degree tiles and a six-hourly update frequency. The resulting improvements to spread and skill were modest and other approaches were investigated.

Shutts (2005) proposed a kinetic energy backscatter scheme that parallels the use of stochastic backscatter in Large Eddy Simulation of turbulence (Mason and Thomson, 1992) but computes a generalized dissipation rate function composed of contributions from numerical diffusion, mountain drag and deep convection. The underlying pattern of two-dimensional streamfunction forcing arising from this was determined from a cellular automaton (CA) that had spatial characteristics reminiscent of mesoscale cellular convection. A revised version of this scheme replaces the CA with a spectral pattern generator in which each spectral component is evolved in time as a first-order autoregressive process with prescribed power spectrum (Berner et al, 2007). These schemes should be contrasted with Mason and Thomson's backscatter forcing pattern which was determined by smoothing the *Curl* of a 3D vector field whose components were random numbers and updating the numbers every other time step. No attempt was made to mimic the statistical properties of the filter-scale eddies.

It would seem to be desirable for the forcing patterns used in stochastic parametrizations to have the same statistical properties and auto-correlation functions as the model error fields (e.g. the statistical properties of the difference field obtained by subtracting fields from simulations made with 100 and 10 km horizontal resolution). However patterns generated from an algorithm driven by random numbers, whilst being plausible as 'model error' fields may or may not be likely given the local state of the forecast model flow. Stochastic backscatter schemes are only flow dependent to the extent that the magnitude of the forcing pattern is scaled by the square-root of the dissipation rate of the resolved flow. The random character of the error pattern does not account for its likelihood given the form of the background flow. The impact of this assumption is hard to know since uncommon error patterns may grow rapidly and common error patterns may decay. In addition no attempt is currently made to advect the error pattern with the flow. This might be prove to be important in getting the appropriate energy transfer between the sub-filter-scale eddy forcing and the explicitly-resolved flow. To address these issues one could attempt to characterize the properties of the model error fields (obtained from the differences between forecasts at very different resolutions) by space-time correlation functions and devise a forcing function whose statistical behaviour matches that of the error fields (Hermanson, 2006). Of course there may still be substantial residual error in the high-resolution 'truth' simulation. For instance, deep convection is neither resolvable nor parametrizable on a 10 km horizontal grid and one must expect important errors in the treatment of mesoscale convective systems and their subsequent upscale impacts.

From the perspective of physical parametrization (and in particular the parametrization of deep convection and gravity wave drag) the phenomena being represented comprise scales that span the whole range between the sub-gridscale and resolved scales of the model. For instance, deep tropical convective systems are driven from the convective updraught scale (< 1 km) but have mesoscale outflows extending over hundreds of kilometres. The theory of convective parametrization is predicated on the assumption that an ensemble of clouds exist within each grid column of the model so that the predicted tendencies of temperature, moisture and momentum are ensemble averages. Shutts and Palmer (2007) quantified the statistical fluctuations of convective warming in a big-domain cloud-resolving model simulation by computing the probability distribution functions of coarse-grained apparent warming, conditioned on a measure of the parametrized convective warming. They showed that the variance of the convective warming is typically greater than the mean when coarse-grained to 120 x 120 km gridboxes. It was found also that the variance tends to increase in proportion to the mean - as implied by the Buizza et al stochastic physics ansatz. Randall et al (2003) proposed a 'super-parametrization' methodology (now named Multiscale Modelling Framework) which avoids assuming statistical equilibrium (and the use of the ensemble average) by a brute force approach. In its original inception (Grabowski and Smolarkiewicz, 1999) this involved embedding 2D cloud-resolving models within each grid-box of a coarse grid climate model and using the net vertical fluxes from the 2D models as a replacement for convective parametrization. The technique has been subsequently demonstrated to provide superior treatment of the Madden-Julian oscillation and other low frequency variability but this comes at a substantial computational cost.

In the very different world of computer games technology there is a inexorable drive towards greater graphical realism and interactivity with the natural environment. The Hollywood film industry is increasingly using computer-generated imagery for special effects involving fluid motion and cloud/smoke rendering. For instance, flight simulators need to render smoke and cloudscapes at useable interactive frame rates for an immersive user experience (e.g. 40 frames per second). Even more remarkable is the development of fluid and particle motion simulators that can simulate the growth of clouds and describe the interaction of a moving aircraft with the cloud itself. This is possible through a combination of powerful graphics processor technology, highly-simplified physics and efficient coding techniques (e.g. Harris, 2003). In the world of visualization, predictive accuracy (in the sense of forecasting where a cloud will be at a particular time and what form it might take) is not required, only visual plausibility. However code stability and speed are paramount to the games programmer. Interestingly, the semi-Lagrangian advection algorithm used in numerical weather prediction has been adopted in real-time fluid simulation for visualization because of its stability with respect to choice of time step (Stam, 2003). Much of the computational overhead in the dynamics comes from the pressure solver which enforces the continuity of mass. Considerable saving can be obtained by reducing the accuracy requirement e.g. by reducing the number of iterations in an iterative scheme.

It has been proposed in Shutts and Allen (2007) that the games programmer's approach to representing natural phenomena could be used to devise fast stochastic pattern generators that couple directly with NWP model flows (e.g. by replacing the streamfunction forcing scheme with a 2D barotropic fluid solver of much higher

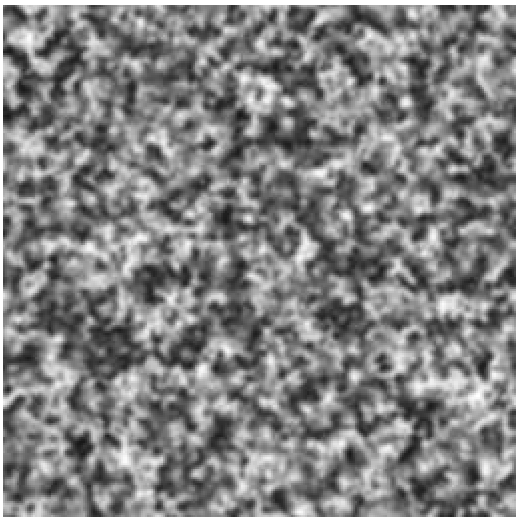

Figure 1: Pattern generated by a cellular automaton where each cell can have 50 'lives' with a life being lost at every step. The grey-scale intensity is a measure of the number of lives remaining with 50 being white and 0 being black. The resulting animated pattern is reminiscent of convective cloud organization in a cold air outbreak over the sea.

resolution than the forecast model). Also, they envisage an efficient - though highly-simplified - cloud-resolving model that could be coupled to the forecast model as a *dual grid* system. The simulated cloud systems on the finer grid would have realistic spatial structure and temporal evolution and this would be passed to the coarser (NWP) model through coarse-grained fluxes and relaxation terms. This dual grid approach applied to each member of an ensemble prediction system would represent a natural extension of current stochastic forcing techniques whilst at the same time opening the door to physical parametrization that breaks free from the vertical column paradigm.

2 Pattern generators for stochastic parametrization

2.1 Cellular automata

The use of cellular automata (CA) was suggested by Palmer (1997) as a means of describing near-gridscale variability associated with unresolved physical processes like deep convection (see Gardner (1970) for a discussion of John Conway's 'Game of life'). He envisaged a probabilistic cellular automaton in which the state of a cell (alive or dead) is governed by a probability distribution function that depends on the state of the cell's nearest neighbours and is also a function of an associated specification of topographic data i.e. land/sea, orographic height and land type. Organized convection could be modelled on this fine grid of cells by making convection (represented by a living cell) more likely if neighbouring cells were already convecting - an assumption that causes the CA convection to cluster. The topographic data could be used to make the state of a cell more likely to be alive (i.e. convecting) over land in the daytime, crudely representing the effect of surface solar heating. Some measure of gross moist atmospheric instability (e.g. Convective Available Potential Energy or CAPE) could be used to control the probability of convection so that the probability is zero (or very small) when the CAPE is zero or negative. The scale and degree of clustering of living cells might even be made a function of some forecast model field such as the CAPE. Considerable interest exists in methods for designing probabilistic CAs to fit complex natural processes but their application in Meteorology is still largely untested.

The CA used in the kinetic energy backscatter scheme by Shutts (1985) (i.e. Cellular Automaton Backscatter Scheme or CABS) was inspired by a CA of the *generations* class know as 'forest fire' or 'prairie on fire'. In it, each cell can have a certain number of lives (say NLIVES, which was equal to 32 in the backscatter scheme) and failure to meet a survival condition causes this to be decremented by 1. A new cell is born provided one out of a set of birth rules is satisfied. The survival condition for a cell is that either 3, 4 or 5 of its neighbouring 8 cells are alive : the birth condition is that 2 or 3 of a dead cell's neighbours are alive *and have not lost any lives*. In medical-speak, any cell that loses a life becomes infertile and reduces the space available for new cells until it loses all of its lives. The value of NLIVES and a notional timestep between successive CA states effectively puts a time scale into the CA's evolution. Since sets of living cells can generate new cells on their leading boundary at each step, a maximum propagation speed is determined by the cell's physical size. This combined with the value of NLIVES implies a length scale. Therefore one has a limited amount of control over the pattern's space and time autocorrelation scales (see Figure 1). In principle NLIVES could be made a function of the forecast model's state so that the autocorrelation scales vary. The CA grid was regular in latitude and longitude and so the spatial autocorrelation scales become smaller with increasing latitude.

For each member of an ensemble forecast the initial state of the CA was seeded differently but from then on the CAs are deterministic. Some provision was made in the code for random seeding of the CAs during the course of the forecast but this was never tried. This would have been desirable since the CA described above is overly periodic in its temporal behaviour. Further development at ECMWF and the Met Office shifted to the use of a spectral forcing approach which was easier to implement and whose behaviour was better understood (see below). The CA approach has yet to be thoroughly tested and it is now believed that Palmer's original Probabilistic CA (PCA) proposal offers the most promise.

2.2 Spectral method with time evolution by autoregressive process

The Spectral Stochastic Backscatter Scheme (SSBS) of Berner et al (2007) uses a pattern generator based on a truncated expansion in terms of spherical harmonics. If $\psi(\lambda,\phi,t)$ is a two-dimensional pattern field on the sphere (with ϕ and λ the latitude and longitude respectively), each spectral coefficient (ψ_n^m) in the expansion is evolved according to the relation:

$$\psi_n^m(t+\Delta t) = (1-\alpha)\psi_n^m + g_n\sqrt{\alpha}\varepsilon(t) \tag{1}$$

where m and n refer to the zonal wavenumber and spherical harmonic function degree respectively; $(1-\alpha)$ is the damped, linear autoregressive parameter with $\alpha \in [0,1]$, g_n is a wavenumber-dependent noise amplitude and $\varepsilon(t)$ is a random number drawn at time t with time step Δt. The total non-dimensional wavenumber squared is $n(n+1)$ and the function g_n can be chosen to provide the desired power spectrum. The autocorrelation time scale τ is given by:

$$\tau = -\frac{\Delta t}{\ln(1-\alpha)} \tag{2}$$

and if large relative to Δt, $\alpha \sim \Delta t/\tau$. The spectral function g_n is assumed to be given by a power law $g_n = n^p$. In the experiments described in Berner et al (2007), $\alpha = 0.83$ (corresponding to an autocorrelation time of about 6 hours) and $p = -5/3$. A backscatter streamfunction forcing is computed by multiplying this pattern by the square root of a local model dissipation rate (due to numerical diffusion, mountain wave drag and convection).

The SSBS has undergone extensive testing and tuning in the ECMWF ensemble prediction system at T255 resolution and improvements in skill and member spread have been clearly demonstrated. Fig. 2 (from Berner et al, 2007) shows the positive impact of SSBS and reduced initial perturbations on two probability skill scores - the ranked probability skill score and the ignorance skill score. A small positive impact is seen at every day in these 10-day forecasts and is particularly noticeable in the ignorance skill score (for details see Roulston and Smith, 2002). Improvements are more striking in the tropics where the ensemble system has insufficient spread and the model dynamics are not active enough.

Fig. 3 shows the impact of SSBS on the systematic error in the northern hemisphere winter 500 hPa geopotential height field in a multi-year integration run at T95 resolution. It can be seen that the tendency of the forecasts

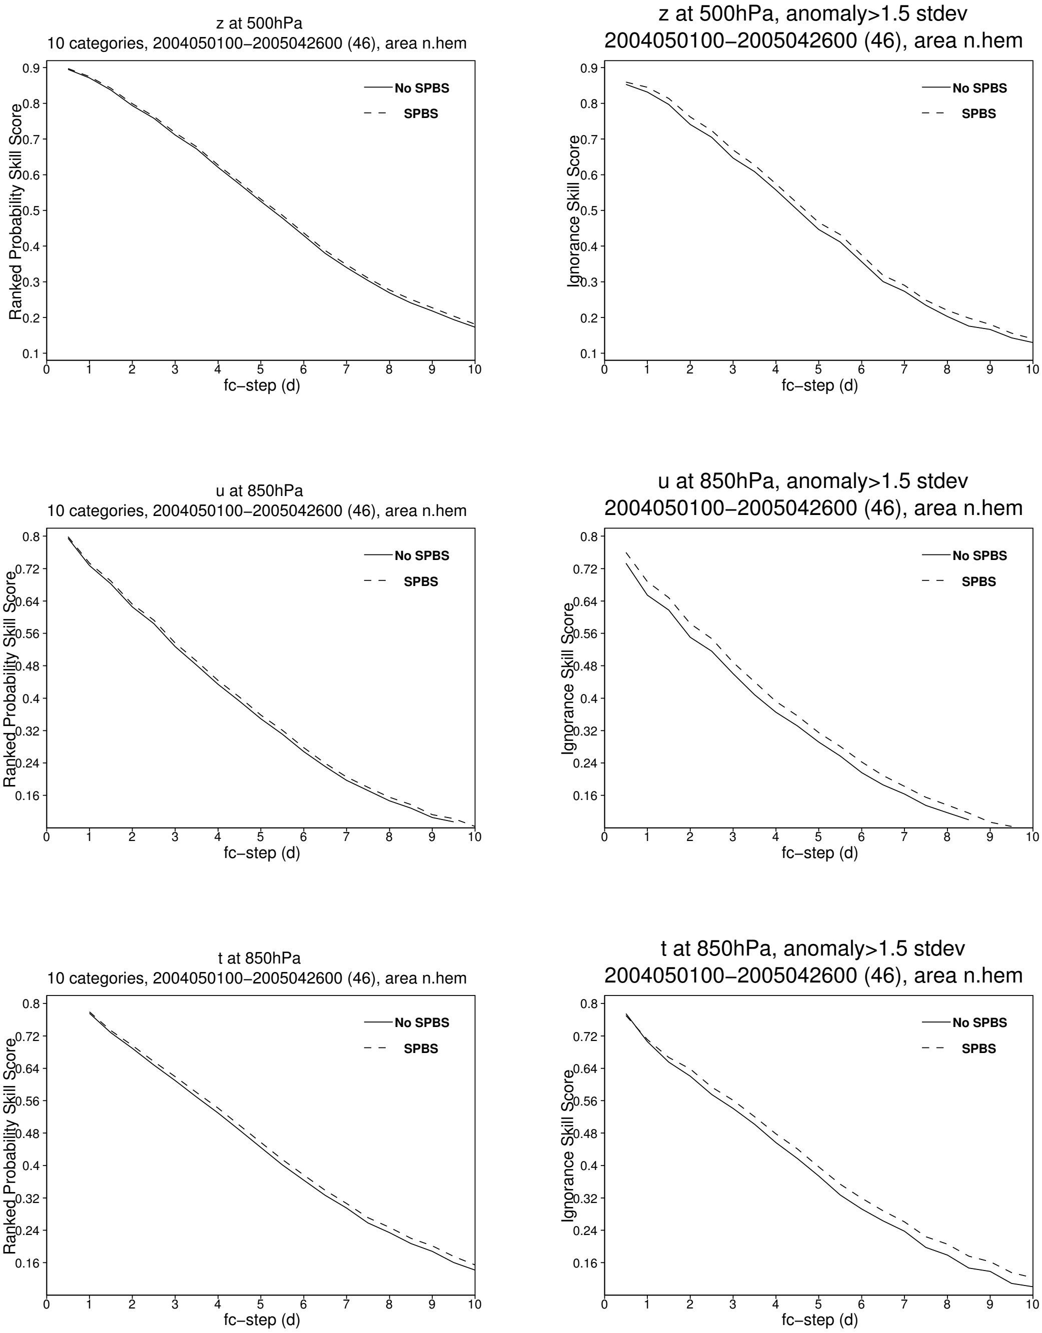

Figure 2: Rank probability skill score (left column) and ignorance skill score (right column) for T255 forecast events in the Northern Hemisphere 'extratropics' which exceed $+1.5$ standard deviations ; '500 hPa geopotential height (top pair); zonal wind at 850 hPa (middle pair), and temperature at 850 hPa (bottom pair).

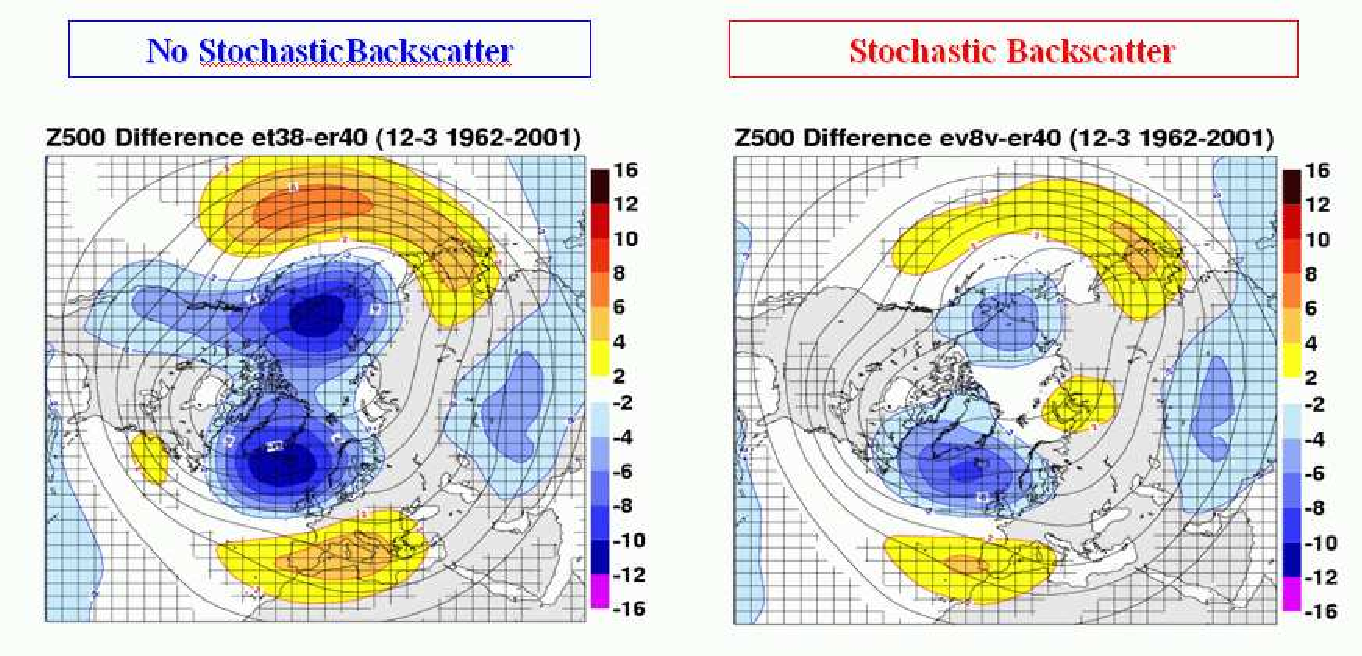

Figure 3: Reduction of the systematic error in the mean 500 *hPa geopotential height for December-March from 1962 to 2001. 6-month long forecasts using the ECMWF model (at T95 resolution and 60 levels) are started on November 1st of each year and the mean is obtained after discarding the first month. Systematic error is defined as the difference between the resulting mean fields and those derived from the ERA40 re-analysis (Kållberg et al (2005))*

without backscatter (left figure) to have excessively low heights in the Icelandic and Aleutian regions is greatly reduced by stochastic backscatter along with a corresponding reduction in the positive height errors over the subtropics. This is matched by an increased frequency of synoptic blocking events near these regions and this increase helps to improve the tendency of the model to underpredict blocking.

2.3 Other approaches

There are many other types of 'microscopic model' that could be used to generate patterns that mimic sub-filter scale atmospheric flow phenomena.

For instance, Khoulder et al (2003) developed a stochastic 'spin-flip' model of tropical convection where Convective INhibition (CIN) is modelled on a fine grid by a set of interaction rules governing the way sites of CIN interact with each other and with a larger-scale model. Coupling between coarse-grained variables and the microscopic states is achieved through an interaction potential. In an idealized model of tropical circulation they were able to demonstrate plausible effects of the stochastic convection on climatology.

Another technique that could be adapted to suit the needs of stochastic parametrization is the popular Lattice-Boltzmann method where collections of particles on a lattice have a probability distribution function governing their velocities and collisions take place between particles at neighbouring grid points (see Chen and Doolen, 1998). In classical applications, the Navier-Stokes equations are obtained for the time-averaged behaviour but the method is amenable to more complex descriptions of the interaction between microscopic and macroscopic behaviour and is well-suited to multi-phase problems and those with complex geometry (e.g. flow through porous media). Given that the evolution of model flow only involves nearest neighbour interactions the method is well suited to parallelization on supercomputers.

The Smoothed Particle Hydrodynamics (SPH) method for representing fluid motion (see the review by Monaghan, 1992) is another potential candidate for fluid pattern generation in stochastic backscatter schemes or direct coupling to a forecast model . In this approach the fluid is represented by a set of discrete lumps of fluid characterized by their mass, position, velocity and a thermodynamic energy or entropy variable. Continuous field descriptions of the flow are obtained using a weighting kernel through which the Newtonian laws of motion for each lump can be determined. The method has the advantage of not needing a mesh for computations and continuity of mass (for the fluid and any other chemical species or conserved property) is implicit in the formulation.

Lastly, the pattern generator could consist of a cloud of two-dimensional monopolar or dipolar point vortices whose individual motions are not only controlled by self-interaction (i.e. the fluid velocity due to each vortex advecting every other vortex) but also by the flow at some level in a forecast model (e.g. Bühler, 2002). The point vortex field could be regarded as a type of model error for the vorticity field at say jetstream level and a coarse-graining procedure could be defined for adding velocity increments to the forecast model. The introduction and removal of point vortices together with their strength could be governed by a probability distribution function that depends on the local model dissipation rate.

3 Computer animation and visualization techniques for fluid flow and cloud dynamics

Since the earliest personal computer systems became available in the early 1980s there has been a strong drive for physical realism in games and simulation software. With the vast increases in microprocessor power, this realism has come a long way from simply describing the parabolic path of a ball under the action of gravity. Flight simulators and Hollywood special effects need realistic and interactive descriptions of the movement of fluids and particulates together with convincing imagery of cloudscapes and smoke. Gone are the days when cartoon, shape-preserving images of clouds drifted by in the field of view. Computer animation of fluid motion and cloud is now based on the same physical laws and equation sets that meteorologists and oceanographers use and in many ways the constraints under which the games programmers operate are similar to those involved in NWP; i.e. the need for speed and stability. Whereas the forecast model has to at least keep ahead of the weather (and in practice is about 100 times faster than the real phenomenon) a cloud modeller might be happy just to keep pace with real time. In a flight simulator game one might hope to have realistic-looking 3D cloud motion, evolution and illumination at 40 frames per second - something that current cloud-resolving models would not be capable of as a direct output. So why is this ?

Part of the reason lies in a crucially different motivation : the games programmer only has to create the illusion of reality whereas the cloud modeller (and particularly the weather forecaster) has to focus on accuracy of representation. For instance Stam (e.g. Stam, 1999, 2003) uses the semi-Lagrangian advection algorithm with very large timesteps to simulate neutrally-buoyant fluid flow carrying smoke or cloud droplets. Since low-order interpolation of flow variables to the departure point acts to unnaturally smooth the fields, Stam has adopted the 'vorticity confinement' technique of Steinhoff and Underhill (1994) to counteract spurious dissipation and maintain vorticity features close the the model's gridscale. Steinhoff's technique has even been used by the Hollywood film industry in the creation of realistic computer-generated smoke scenes.

Inspired by the work of Stam, a flow-past-buildings code was written to confirm the extra speed attainable when the requirement for accuracy is relaxed. The code executes sufficiently quickly for the user to be able to interact with the display and place smoke tracers into the computational domain via mouse pointer actions. The authors were impressed by the efficiency and realism of the simulation and realized the potential for this approach to be used as a more sophisticated replacement for the cellular automaton pattern generator in stochastic parametrization. The simplest way of viewing vorticity confinement in two-dimensions is the addition of an apparent force to the momentum equation that acts tangentially to the vorticity surface with a strength proportional to the magnitude of the vorticity itself. The momentum equation can be written in the

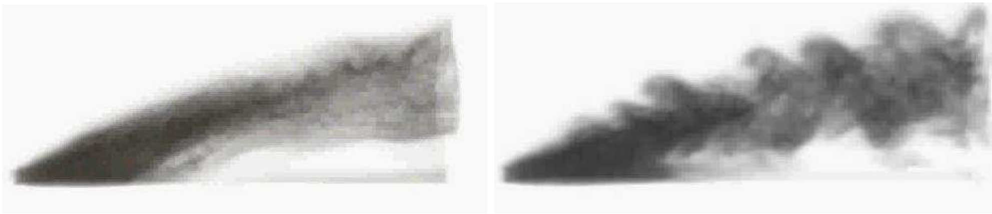

Figure 4: Tracer concentration in a simulation of a jet emerging at 30 degrees to a flat plate; without (left) and with (right) vorticity confinement (from Steinhoff et al, 2006). The computational grid uses $129 \times 65 \times 65$ *points and there is no cross-wind velocity.*

form :

$$\frac{\partial \mathbf{V}}{\partial t} + (\mathbf{V} \cdot \nabla)\mathbf{V} + \frac{1}{\rho}\nabla p = \nu \nabla^2 \mathbf{V} - \varepsilon \mathbf{S} \quad (3)$$

where $\mathbf{V}$ is the velocity, p is the pressure, ρ is the density and ν is the viscosity. The vorticity confinement term $\varepsilon \mathbf{S}$ is defined such that ε is a constant and

$$\mathbf{S} = \mathbf{n} \times (\nabla \times \mathbf{V}) \quad (4)$$

where the unit normal to the vorticity surface $\mathbf{n}$ is given by:

$$\mathbf{n} = \frac{\nabla \eta}{|\nabla \eta|} \quad (5)$$

and

$$\eta = |\nabla \times \mathbf{V}|. \quad (6)$$

An example of its effect on the simulation of fluid flow is given in Steinhoff et al (2006) where a three-dimensional jet emerging from a flat plate is simulated with and without the vorticity confinement term in eq. (3). The simulation without confinement is rather bland and lacking small-scale turbulent eddies but with vorticity confinement looks more realistic and shows rolled-up Kelvin-Helmholtz eddies at the edge of the plume.

A variant of this general approach for the semi-Lagrangian algorithm is to analyse the diffusive nature of the numerical error term for the interpolation of fields to the departure point and create an approximate term that cancels the error. This step is not equivalent to merely using a higher-order interpolant but has a similar effect at reduced computational cost.

The basic idea behind this method is to perform a "modified equation" analysis (see Morton and Mayer, 2005) for the first order semi-Lagrangian advection scheme. When applied to the one-dimensional advection equation for ϕ:

$$\phi_t + u\phi_x = 0 \quad (7)$$

one obtains

$$\hat{\phi}_t + u\hat{\phi}_x = \kappa \hat{\phi}_{xx}, \quad (8)$$

where $\hat{\phi}$ is the equivalent numerical solution and the numerical diffusivity κ is given by:

$$\kappa = \frac{r(1-r).(\delta x)^2}{2\delta t} \quad (9)$$

(see also McCalpin, 1988). Here δx and δt are the horizontal grid spacing and time step respectively while the parameter $r \in [0,1]$ is the (relative) distance of the departure point from the nearest grid point and is related to the fractional part of the Courant number. The fact that $\kappa \geq 0$ is independent of the Courant number is why the method is stable. Note that for non-constant advection there is also a dispersive correction to (8) which is ignored since we are only interested in the dissipation. The fact that r is calculated before the semi-Lagrangian advection is performed suggests that the numerical dissipation terms could be used for **S** in eq. (3) in place of the vorticity confinement term (eq. (4)) but now with ε close to but less than unity. The "ill-posed" nature of this anti-diffusion means that some care must be taken with the implementation of such a term in order to maintain numerical stability. However the potential advantage of the method is that it can be applied to all fields and the parameter ε is less arbitrary. The beneficial effects of this method can be seen in Figure 5 which shows the vorticity field for flow over a series of boxes/buildings. The white area indicates the dispersion of a passive tracer. The four simulations in this figure comprise:

- the standard first order scheme
- the first order scheme applied to momentum but with monotone cubic interpolation of the tracer
- cubic interpolation of all variables
- the first order scheme with the proposed anti-diffusion ($\varepsilon = 0.8$).

The simulation that used the anti-diffusion scheme was computationally as cheap as the standard first-order scheme and typically about 4 or 5 times faster than those using cubic interpolation.

As with vorticity confinement the emphasis is on creating the effect of higher resolution without the attendant pointwise accuracy. The additional terms are like deterministic backscatter and probably do much to improve the statistical properties of the flow without attempting to have extra accuracy in the Lagrangian advection step.

Another way the games programmer has achieved speed is to simplify the equations of moist thermodynamics, cloud microphysics and radiative transfer. It is not unreasonable to question the extent to which the complexity of cloud microphysics is necessary for the modelling of the overall statistical character of deep convection. Indeed it is an interesting challenge to perform a kind of 'algorithmic compression' and see if much simpler parametrization could suffice. The reason that this is not quite the outrageous assertion that it seems to be is that cloud-resolving models (and the latest generation of ultra-high limited area forecast models) are frequently run with an order of magnitude less resolution than a true cloud-resolving model needs (e.g. 2 km horizontal gridlength when 200 m is required). Much of the complexity of the description of cloud microphysics is compromised by excessive numerical diffusivity in the cloud scales. Shutts and Allen (2007) describe a simple 'fake' cloud-resolving model that attempts to create a highly-simplified moist physics while permitting believable deep convective clouds.

It is convenient for our purposes to use the term *emulator* to describe a class of computer model that mimics a conventional numerical simulation (i.e. the *simulator* such as a cloud-resolving model) but employs a much higher level of simplification than one would regard as acceptable for a research study. The *emulator* therefore imitates the *simulator* and the *simulator* models reality. For the purposes of parametrization the emulator must at the very least be numerically stable and sufficiently accurate to produce plausible patterns for use in stochastic parametrization. At the other end of the spectrum the emulator may use the same equations as the simulator and ultimately converge to it if the accuracy requirements are successively increased.

4 Proposal for a 'dual grid' class of stochastic parametrization

The combined stochastic pattern generator (or emulator) and the forecast model constitute a dual grid system in which the former acts to force statistical fluctuations due to sub-filter scale processes. The potential level of complexity of the emulator could range from a simple 2D cellular automaton to a 3D cloud-resolving model with simplified physics (similar to the super-parametrization concept of Randall et al, 2003). In the computationally-cheapest scenario all the conventional column-based parametrization schemes are retained and the emulator modulates their output. The computational demands of this kind of emulator are very small

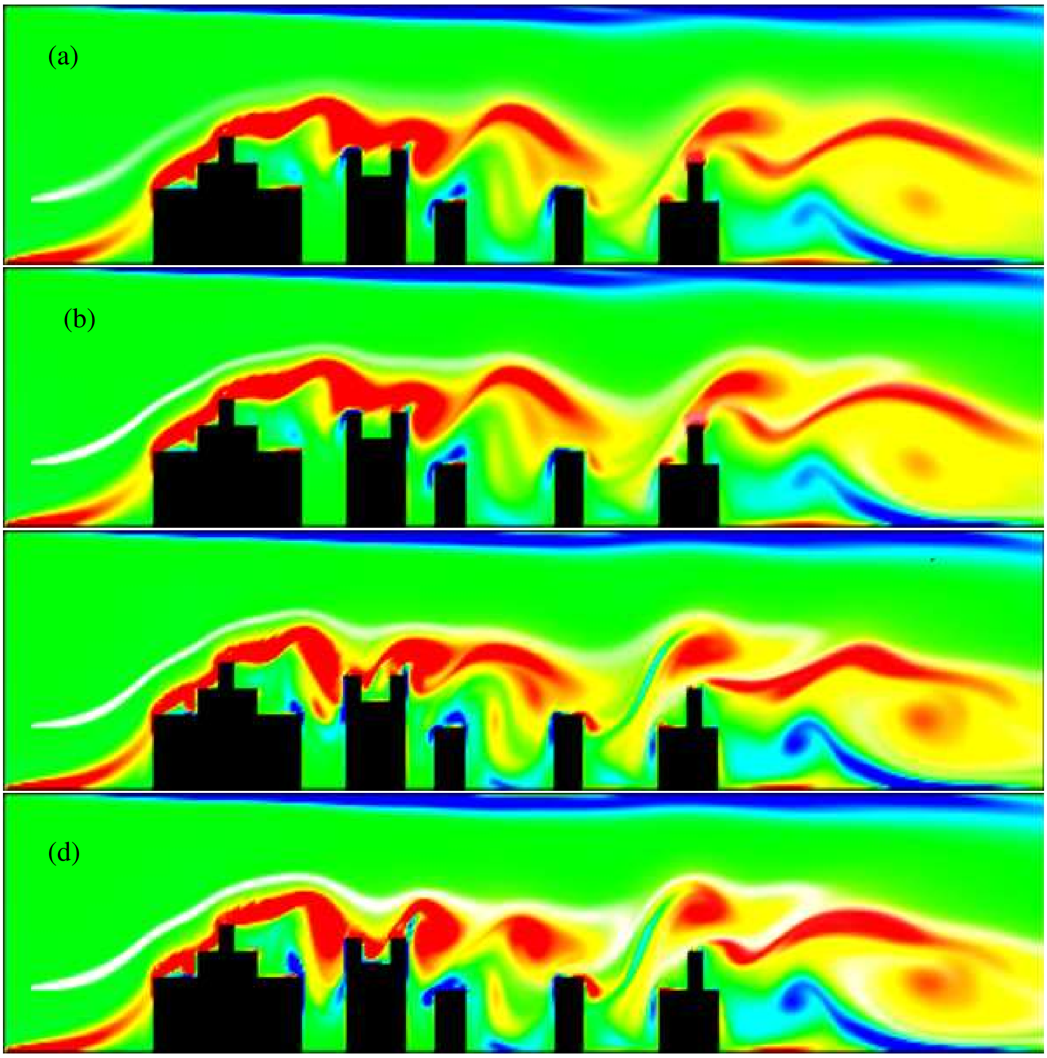

Figure 5: A snapshot of the vorticity field from an interactive movie of flow past a series of buildings. 500×125 *points were used and the location of buildings is enforced by masking the flow to zero. The white streak is a tracer field introduced at its leading edge on the left of the domain. On a 2.4 GHz Pentium 4 personal computer using OpenGL graphics the animation was able to run at 15 frames a second. (a) Standard linear scheme; (b) using monotone cubic interpolation on tracer; (c) monotone cubic on all variables; (d) linear interpolation with anti-diffusion terms*

relative to the conventional forecast model. In contrast, the computational burden of the cloud-resolving model limit would exceed that of the forecast or climate model. The role of the forecast model would then be to merely orchestrate sub-mesoscale processes in the fine-scale model and remove the need for many of the sub-gridscale parametrizations (e.g. deep convection and mountain drag). Wherever the computational demand lies it is important to remember that the fine-scale emulator has substantial compromises in its pointwise accuracy relative to the coarser NWP or climate model. In effect, the combined system does its dynamics and physics at different resolutions and on different grids with the physics being computed at higher resolution but degraded accuracy for speed.

On a related side note ECMWF have recently compared forecast accuracy in the IFS when the dynamics and physical parametrizations are computed at different resolutions. Somewhat surprisingly it was found that forecasts made with a higher resolution for the physics grid than the dynamics gave more forecast skill than the combination of higher resolution dynamics and coarser resolution physics (Martin Miller, personal communication). In the former case the physical parametrizations derive their input from interpolated dynamical fields whereas in the latter case it was obtained from coarse-grained fields. The idea of using different grids to compute dynamics and physics is similar to our proposed dual grid system in which 'the physics' is played out on a fine grid using an emulator.

As a proof-of-concept exercise a simple fluid dynamical problem has been modelled on a fine grid and coupled to a coarse grid i.e. a dual grid simulation. Coupling was attempted in several different ways and the vorticity confinement technique tested with different values of ε. The fine grid has ten times more resolution in both directions than the coarse model but uses the same time step, therefore implying a Courant number ten times larger (see Shutts and Allen (2007) for details). The simulated flow consists of three convective plumes ascending through an isentropic atmosphere at rest. The plumes are forced by three heat sources situated at the surface, each of which spans 3 grid points on the fine mesh. The heat sources are entirely sub-gridscale for the coarse model and therefore have no direct influence in the absence of diffusion and a boundary layer scheme. Coupling is achieved by forcing the coarse model with coarse-grained eddy fluxes computed from the fine model. Although unimportant for the present problem a relaxation term was added to the fine-scale model which nudged the fine variables towards the linearly-interpolated coarse ones. Some type of two-way relaxation/control is required in order that the two models do not drift too far apart over the course of a long time integration. Excessive relaxation from the coarse to fine model suppresses some of the flow structures that one would hope would be passed from fine to coarse model.

Figure 6 shows the buoyancy on both coarse- and fine-mesh model at some time after the plume has ascended up to the model's lid. From this figure it can be seen that, even though the fine mesh solution would be considered highly inaccurate, many of the actual features of an accurate converged solution are present. It can also be seen that the coarse model, which only sees the plumes through the coupling to the fine grid, reproduces most of the features of the rising plume. It is important to note that, apart from the relaxation term (which could be set to zero in this instance), there are no tunable parameters in this form of parametrization. It is unlikely that a column-based parametrization could generate the upscale influence found in the coupled grid simulation and the additional computational expense, provided not too prohibitive, would be worthwhile.

The feasibility of the dual grid approach will depend primarily on the speed achievable for the fine grid model and on the nature of the coupling between the two grids. In the context of ensemble prediction systems the fine grid model would necessarily have to be simple because the computational burden would be too high for each forecast member to have a 3D cloud-resolving model associated with it. More realistically one might consider using a fast 2D fluid emulator (with stochastic forcing) as a fine grid model associated with each member and relax the large scales of the fine model to the forecast model whilst relaxing the small scales of the forecast model towards the emulator. This could be treated as an enhancement of the SSBS to include vorticity advection.

The coupling between these fine-scale and coarse-scale models may also be achieved by coarse-graining eddy fluxes derived from the fine model and using their flux divergences to drive the coarse model. In practice one could imagine having a mix of linear relaxation and eddy forcing for the coupling of the two models.

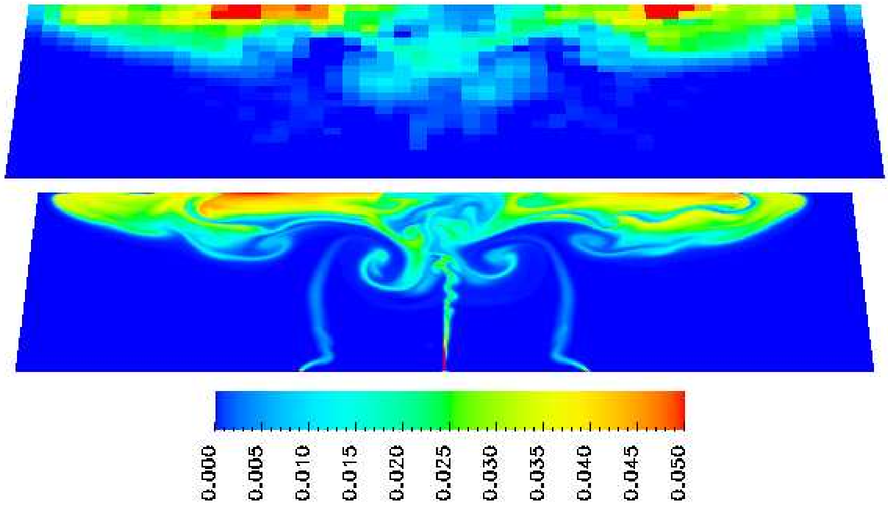

Figure 6: The buoyancy field (ms^{-2}) on the coarse (upper) and fine (lower) model grids after the plume has ascended to the top of the model domain (upper edge) from three surface heat sources. The model has 25×50 *points and uses a 25 second time-step.*

5 Discussion

It is proposed to use techniques similar to those employed by computer games and visualization programmers for emulating the statistical behaviour of some sub-gridscale physical processes. These might be as simple as using a probabilistic CA to describe patterns of convective cloud organization to use as driver in stochastic parametrization scheme or as ambitious as a 3D cloud system emulator that directly couples to a forecast model (in an analogous way to the super-parametrization concept except with much lower computational cost than a conventional cloud-resolving model). In this latter limit conventional column-based parametrization would be made redundant, though realistically one should expect that some blend of emulator and conventional parametrization would be necessary in practice - depending on the computer resources available. It would also be possible to use 2D convection emulators along each latitude circle. This would resemble the use of a horizontally-anisotropic grid (e.g. Shutts, 2006) but would have less computational cost.

In the context of ensemble prediction, each forecast member would see a different evolving pattern of stochastic forcing resulting from the underlying stochastic nature of the pattern maker or emulator. The current ECMWF spectral backscatter scheme (SSBS) uses a first-order autoregressive model to describe the time evolution of streamfunction forcing spectral amplitudes. An interesting next step would be to replace the autoregressive model with a high-resolution barotropic fluid emulator that is itself being stochastically forced. The largest scales of the fluid emulator would be relaxed towards the jetstream level flow in the forecast member. In this way the fluid emulator acts as an elaborate flow-dependent dynamical filter for the random numbers used to make it stochastic. The nature of the coupling between the emulator and the forecast model would probably be of critical to the success of this approach.

As computer resources become available the fluid emulator would become three-dimensional and include a more realistic description of the physics (though still an emulator in the sense of computer games and visualization). Ultimately one envisages the continual refinement of the emulator until it takes over the role of the NWP model - thereby defining a convergent process. At this stage one must regard the dual grid strategy defined here as speculative. Apart from the barotropic fluid emulator, future work will be directed at the design of a fast and stable deep convection emulator that can be coupled to a forecast model and hopefully speed up the replacement of column-based convection parametrization.

Acknowledgements

We thank John Steinhoff for Fig. 4 and Clive Temperton who alerted the first author to the paper by McCalpin.

References

Berner, J., G. J. Shutts, M. Leutbecher and T. N. Palmer (2007). A spectral stochastic backscatter scheme and its impact on flow-dependent predictability in the ECMWF ensemble prediction system. (*in prep*)

Bühler, O. (2002). Statistical mechanics of strong and weak point vortices in a cylinder, *Phys. Fluids*, **14**, 2139-2149.

Buizza, R., M. Miller and T. N. Palmer (1999). Stochastic representation of model uncertainty in the ECMWF Ensemble Prediction System, *Q.J.R.Meteorol. Soc*, **125**, 2887-2908.

Chen, S. and Doolen, G. D. (1998). Lattice Boltzmann method for fluid flows. *Annu. Rev. Fluid Mech.* **30**, 329-364.

Gardner, M. (1970). Mathematical Games. The fantastic combinations of John Conway's new solitaire game "life", *Scientific American*, **223**, 120-123.

Grabowski, W. W. and P. K. Smolarkiewicz (1999). CRCP: A cloud resolving convection parametrization for modelling the tropical convecting atmosphere. *Physica D*, **133**, 171-178.

Harris, M. J. (2003). Real-time cloud simulation and rendering. *University of North Carolina, USA*, 173 pp.

Hermanson, L. (2006). Stochastic physics: a comparative study of parametrized temperature tendencies in a global atmospheric model. *PhD thesis, Dept. of Meteorology, Reading University*, 145 pp.

Kållberg, P., P. Berrisford, B.J. Hoskins, A. Simmons, S. Uppala, S. Lamy-Thépaut and R. Hine (2005), ERA-40 Atlas. ERA-40 Project Report Series No. 19, ECMWF, Reading, UK, 191 pp.

McCalpin, J. D. (1988). A quantitative analysis of the dissipation inherent in semi-Lagrangian advection. *Mon. Wea. Rev.*, **116**, 2330-2336.

Monaghan, J J. (1992). Smoothed particle hydrodynamics. *Ann. Rev. Astron. and Astrophysics*, **30**, 543-574.

Morton, K. W. and D. F. Mayers (2005). Numerical solution of partial differential equations: An Introduction. *Cambridge University Press*, 227, pp.

Palmer, T. N. (2001). On parametrizing scales that are only somewhat smaller than the smallest resolved scales, with application to convection and orography. *ECMWF Workshop on New Insights and Approaches to Convective Parametrization, 4-7 November 1996. Shinfield, Reading, UK*

Randall, D. and M. Khairoutdinov and A. Arakawa and W. W. Grabowski (2003). Breaking the cloud parametrization deadlock. *Bull. Amer. Meteor. Soc.*, **84**, 1547-1564.

Roulston, M. S. and L. A. Smith (2002). Evaluating probabilistic forecasts using information theory. *Mon. Wea. Rev.*,**130**, 1653-1660.

Shutts, G. J. (2005). A kinetic energy backscatter algorithm for use in ensemble prediction systems, *Q.J.R.Meteorol. Soc*,**131**, 3079-3102.

Shutts, G. J. (2006). Upscale effects in simulations of tropical convection on an equatorial beta-plane. *Dyn. Atmos. Oceans*, **42**, 30-58.

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**, 187-202.

Shutts, G. J. and T. Allen (2007). Sub-gridscale parametrization from the perspective of a computer games animator. *Atmos. Sci Letts., accepted*

Stam, J. (1999). Stable fluids. *Proceedings of SIGGRAPH 1999*, 121-128.

Stam, J. (2003). Real-time fluid dynamics for games. *Proceedings of the Game Developer Conference, March 2003*.

Steinhoff, J. and D. Underhill (1994). Modification of the Euler equations for ”vorticity confinement”: Application to the computation of interacting vortex rings, *Phys. Fluids*, **6**, 2738-2744.

Steinhoff, J., N. Lynn, W. Yonghu, M. Fan, L. Wang and W. Dietz (2006). Turbulent flow simulations using vorticity confinement. from *Chapter 12: Implicit Large Eddy Simulation: Computing turbulent flow dynamics, Cambridge University Press*.