Exp Fluids (2010) 48:915–925
DOI 10.1007/s00348-009-0752-0

RESEARCH ARTICLE

A realtime observatory for laboratory simulation of planetary flows

Sai Ravela · John Marshall · Chris Hill · Andrew Wong · Scott Stransky

Received: 19 October 2007 / Revised: 5 April 2009 / Accepted: 6 April 2009 / Published online: 19 November 2009
© Springer-Verlag 2009

Abstract Motivated by the mid-latitude atmospheric circulation, we develop a system that uses observations from a differentially heated rotating annulus experiment to constrain a numerical simulation in real-time. The coupled physical-numerical system provides a tool to rapidly prototype new methods for state and parameter estimation, and facilitates the study of prediction, predictability, and transport of geophysical fluids where observations or numerical simulations would not independently suffice. A computer vision system is used to extract measurements from the physical simulation, which constrain the model-state of the MIT general circulation model in a hybrid data assimilation approach. Using a combination of parallelism, domain decomposition and an efficient scheme to select ensembles of model-states, we show that estimates that effectively track the fluid-state can be produced. To the best of our knowledge, this is the first realtime coupled system for this laboratory analog of planetary circulation.

1 Introduction

In the differentially heated rotating annulus experiment, a rotating annulus with a cold center (core) and warm periphery develops a circulation that is dynamically similar to the mid-latitude circulation of the atmosphere (see Fig. 1). It is a robust and easily conducted laboratory experiment, which has been used to study a variety of properties of geophysical fluids including geostrophic turbulence (Morita and Uryu 1989), convection (Hide 1958), baroclinic instability (von Larcher and Egbers 2005; Tajima and Nakamura 2003; Read 2003), and chaos (Read et al. 1992; Lee 1993). It has also been used as a test-bed for evaluating the utility of numerical models (Read et al. 2000; Geisler et al. 1983).

In this paper, we present a realtime observatory for the differentially heated rotating annulus experiment (see Fig. 2). The observatory is defined as a coupled physical-numerical system with sensors to take measurements of the evolving physical process, a numerical model for forecasting it, and algorithms that couple the model with observations. We envision the coupling to be two-way; observations constrain the model, and the model guides where and when to take measurements. In this way, the observatory produces an evolving state estimate in realtime that is closer to the laboratory flow than either observations and model alone.

A realtime observatory opens up a number of fundamentally exciting possibilities for experiments in geophysical fluids. The constrained numerical model can be used to study properties of fluids that are not directly measured (surface height, pressure fields, vertical velocities, radial heat transport, etc.), and can permit volumetric visualization of flows at a much higher resolution than observations. Further, because data gathering cannot be arbitrarily dense and realtime, the coupled system provides an alternative when few measurements gathered in realtime can effectively constrain the model. Optimally deciding when and where to observe requires, in general, guidance from the model, and thus the model must be integrated and constrained in realtime too. When realtime performance is achieved in observation, simulation and estimation, we may have a new way to experiment with fluids in many different dynamical regimes.

S. Ravela (✉) · J. Marshall · C. Hill · A. Wong · S. Stransky
Earth, Atmospheric and Planetary Sciences,
Massachusetts Institute of Technology,
54-1624, 77 Massachusetts Avenue,
Cambridge, MA 02139, USA
e-mail: ravela@mit.edu

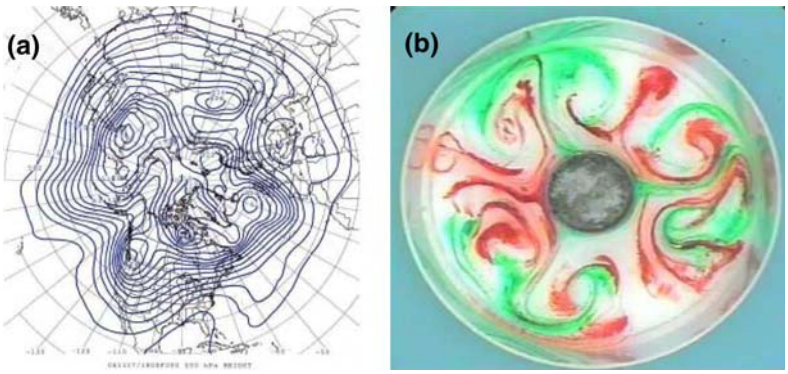

Fig. 1 Image **a** shows the 500 hPa heights for 11/27/06:1800Z over the northern hemisphere centered at the north pole. Winds flow along the pressure contours. Image **b** shows a tracer (dye) in a laboratory analog. The tank is spinning and the camera is in the rotating frame. Tracer droplets initially inserted at the periphery (*red dye, warm region*) and around the central chilled can (*green dye, cold region*) has evolved to form this pattern. The laboratory analog and the planetary system are dynamically akin to one-another. We study the state-estimation problem for planetary flows using the laboratory analog

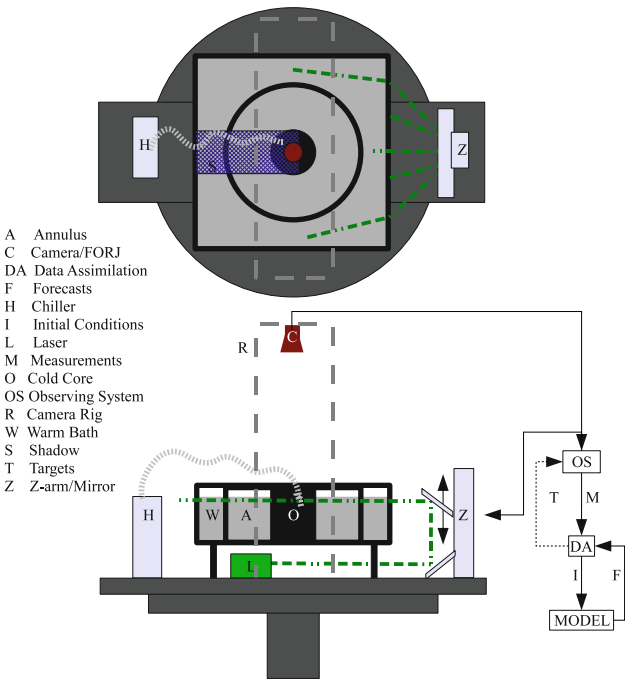

Fig. 2 The observatory and its components are schematically shown in plan (*top*) and elevation (*bottom*) views. The physical component consists of a rotating table on which a tank, camera and illumination control system are mounted. The computational part consists of a measurement system for velocimetry (OS), a numerical model, and an assimilation system (DA), as described more fully in the text

Of particular interest to us is the use of the observatory to accelerate research in prediction and predictability of the large-scale atmosphere. Topics such as state and parameter estimation, model error and adaptive sampling particularly benefit, but so do others. For example, we may quantify long term performance of models better by studying azimuthally integrated quantities over prolonged periods of time under a variety of temperature differentials.

It is not possible in this paper to explore each and every application. A large number of potential applications, however, will require the numerical system to track the fluid's state in real-time as the first step. Our focus in this paper, therefore, is on the design of the observatory for the differentially heated rotating annulus experiment, including a procedure to estimate model-states in real-time.

The tracking problem studied here is, to be sure, of direct importance to numerical weather prediction (NWP). In NWP, predictions are typically made using general circulation models (GCMs), which implement the discretized governing equations. GCMs typically have uncertain parameters and crude parametrizations, uncertain initial and boundary conditions, and their numerical schemes are approximate. Thus, not only will the error between physical truth and simulation evolve in a complex manner, but the probability density function (PDF) of the evolving model state's uncertainty is unlikely to retain the true state within it (Lorenz 1963). A way forward is to constrain the model with observations of the physical system (Wunsch 1996).

Studying the tracking problem in the laboratory is convenient because repeatable experiments with real data can be performed using far simpler logistics than the operational setting. It is also useful because key challenges in the large-scale tracking problem are also addressed in the laboratory. They include nonlinearity of the process, dimensionality of the numerical model, uncertainty of states and parameters, and realtime performance. Solutions found in a laboratory setting can not only accelerate operational acceptance of new methods, but also inform many other coupled numerical-physical experiments.

In a geophysical context, the rotating annulus experiment has already been used to explore the utility of numerical models. Read et al. (2000) use measurements from the annulus and numerical models based on Eulerian

schemes to report that such models can simulate baroclinic instability reasonably well, but this does not necessarily imply predictive skill. In more recent work (Read 2003), numerical studies are combined with laboratory experiments in the study of heat transport. Effort has also been afoot to study prediction and predictability problems using the laboratory setting (Young and Read 2006; Ravela et al. 2003, 2007). To the best of our knowledge, however, this is the first coupled system for the differentially heated rotating annulus to operate in realtime (Ravela et al. 2007).

Our coupled system continually observes the experiment and uses a hybrid estimation method to constrain the model-states of the numerical model. It is implemented using off-the-shelf hardware and commercially or publicly available software. Although it can operate in many dynamical regimes, in experiments presented here a realtime cycle of forecast-observe-estimate must be and is completed within roughly 10 s. The system is now in routine use.

2 The observatory

The observatory, illustrated in Fig. 2, has a physical and computational component. The physical component consists of a perspex annulus, of inner radius 8 cm and outer radius of 23 cm, filled with 15 cm of water and situated rigidly on a rotating table. A robotic arm by its side moves a mirror up and down to position a horizontal sheet of laser light at any depth of the fluid using a periscope arrangement. The light sheet is produced by a continuous wave 1 W 532 nm laser equipped with readily available line-generator optics. It is $\sim$0.5 mm thick between entry and exit in the annulus, and dead level. Fluorescent pliolite particles (Dantec Dynamics, sg 1.03 g/cc, 50 μm) are homogenized in saline water of equal density and scatter incident laser illumination. Particles appear as a plane of textured dots in the 12-bit quantized, 1 K $\times$ 1 K images (see Fig. 4) of an Imperx camera. These images are transferred out of the rotating frame using a Hitachi fiber-optic rotary joint (FORJ).

The actual configuration of these elements is shown in a photograph of our rig in Fig. 3. The observation rig is carefully mounted and tested for vibrations. To appreciate the importance, consider particles that can move at up to 2 cm/s. The camera with scale factor 0.5 mm/pixel (approx.) is positioned 50 cm away from the annulus. At a sampling rate of 1/4 s, the camera must shake by less than 0.1° from the horizontal to have less than 10% motion noise. We center the rig and hold the FORJ-assembly using four bungee chords, which have the appropriate stiffness (see Fig. 3) to damp vibrations and moments. The measured motion noise is $\sim$0.5 pixels, which implies that the camera shake can be no more than 0.03°.

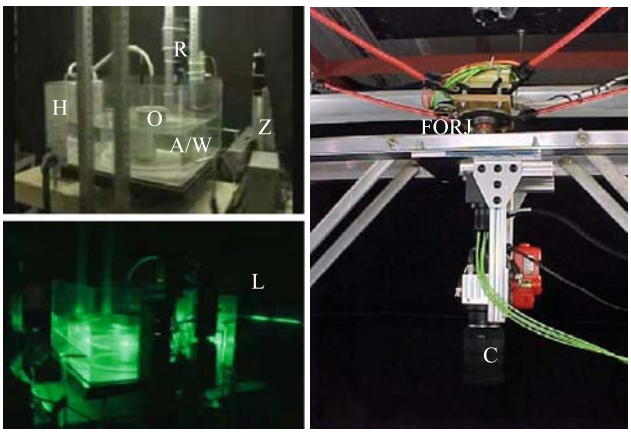

Fig. 3 The apparatus, depicted using symbols defined in Fig. 2. The fiber optic rotary joint (*FORJ*) allows image data to leave the rotating frame and is held stably by orange bungee chords. The square tank mitigates refractive effects at the annulus interface. It is insulated at the bottom by a thick black rubber pad

The computational aspects of the observatory are also shown in Fig. 2. A server acquires particle images and computes velocity using PIV on two processors (Fig. 2, labeled OS). Velocity vectors are passed to an assimilation program (Fig. 2, labeled DA) that combines them with model forecasts to estimate new states. These estimates become new initial conditions for the models. Estimates of states and their uncertainties will also be used in the future to target observations adaptively (T, dotted line). Here, we go on to discuss individual components of the current system.

2.1 Laboratory experiment and visual observation

We homogenize the fluid with neutrally buoyant particles and spin up the rotating platform at the desired period (between 3 and 12 s). After 20 min or so, the fluid comes into solid body rotation. The inner core is then cooled using a chiller (see Fig. 4). Within minutes, the water near the core cools and becomes dense. It sinks to the bottom to be replenished by warm waters from the periphery of the annulus, thus setting up a circulation. At high enough rotation rates eddies form (see Fig. 1) and baroclinic instability sets in.

Once cooling commences, we turn off the lights and turn on the continuous wave 1 W 532 nm laser, which emits a horizontal sheet of light that doubles back through a periscope to illuminate a sheet of the fluid volume (see Fig. 4). An imaging system in the rotating frame observes the developing flow using a camera looking down at the annulus. We measure the horizontal component of velocity from particle motion in image pairs, acquired 125–250 ms apart using LaVision's DaVis PIV software. Horizontal velocity is computed in 32 $\times$ 32 windows with a 50%

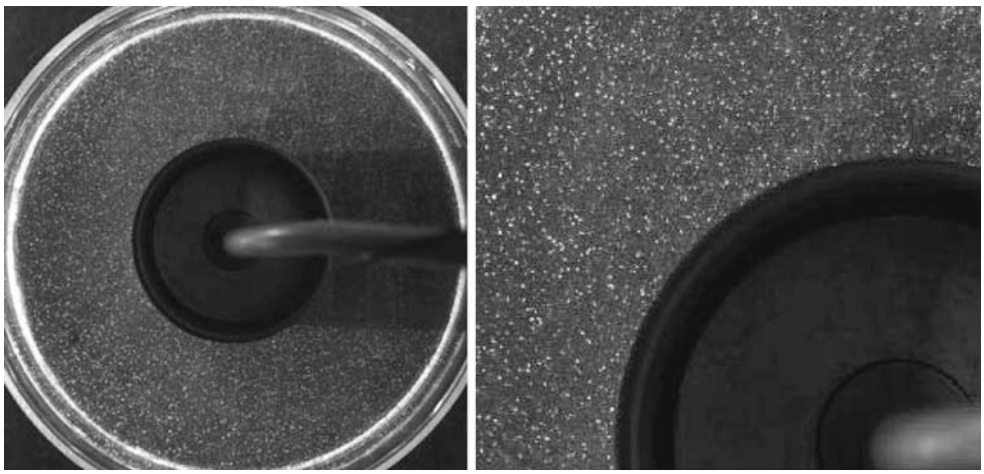

Fig. 4 The camera's view of a horizontal plane of the fluid in laser light (*left*). The chiller can be seen in the center. A magnified view of the *upper-left* quadrant shows embedded particles (*right*), which are used for PIV

overlap between windows. It takes one second to acquire and compute PIV of a single 1 K × 1 K image pair by distributing the computation across two 2.8 GHz processors. An example is shown in Fig. 6.

Observations are gathered at several levels on a repeating cycle. The mirror moves to a preset level, the system captures images, horizontal velocity is computed, and the mirror moves to the next programmed level and so on, scanning the fluid volume in levels. We typically measure the fluid velocity at five different levels. Thus, measurements of the whole fluid are available every 5 s to constrain the numerical model.

We also gather temperature measurements in a separate experiment to establish a climatological temperature boundary condition. Five temperature probes (RTDs) are distributed evenly along a vertical line spanning the fluid depth on the inner boundary. Temperature is recorded for a few minutes of circulation. This process is repeated multiple times at different randomly chosen azimuthal placements of the vertical line on the inner boundary. Climatology is then obtained by averaging all measurements at corresponding observed depths, and interpolating to all model levels (see Fig. 5). Measurements are similarly gathered on the outer boundary, but all outer-wall temperature measurements are averaged to represent the outer boundary condition.

3 Numerical model

We use the MIT GCM developed by Marshall et al (1997a, b) to numerically simulate the laboratory experiment. The MIT-GCM is freely available software and can be configured for a wide variety of simulations of atmosphere, ocean or laboratory flows. Here, the model is used to solve the equations that govern the evolution of an incompressible Boussinesq fluid in hydrostatic balance. The governing equations are:

$$\frac{\partial \vec{v}_\mathrm{h}}{\partial t} = G_{v_\mathrm{h}} - \frac{1}{\rho_0}\nabla_\mathrm{h} p \quad \text{horizontal momentum} \tag{1}$$

$$\nabla_\mathrm{h}\vec{v}_\mathrm{h} + \frac{\partial w}{\partial z} = 0 \quad \text{continuity} \tag{2}$$

$$\frac{\partial p}{\partial z} + g\rho = 0 \quad \text{hydrostatic balance} \tag{3}$$

$$\frac{\partial \theta}{\partial t} = G_\theta \quad \text{thermodynamic} \tag{4}$$

Here, the three-dimensional velocity is denoted by $\vec{v} = [\vec{v}_\mathrm{h}; w]$ where $\vec{v}_\mathrm{h}$ is the horizontal velocity, w is the vertical velocity and ∇_h is the horizontal gradient operator, p is the pressure, assumed to be in hydrostatic balance with the mass field, g is the acceleration due to gravity, $\rho = \rho(\theta)$ is the density with ρ_0 a constant reference value and θ is the temperature. The term G_{v_h} in the horizontal momentum equation includes inertial, Coriolis and frictional terms; G_θ is the corresponding term in the thermodynamic equation and includes advection and thermal diffusion. Explicit forms of the G's are discussed in detail in Marshall et al. (1997a, b).

No-slip boundary conditions are assumed on all solid boundaries and a linearized free surface is adopted (Marshall et al. 1997a, 1997b). The temperature at the outer wall of the tank is held constant; at the inner core, it is set to a vertical profile taken from a separate experiment (see Fig. 5b). The bottom boundary is assumed to be thermally-insulating.

Finite difference forms of the above equations are solved in cylindrical coordinates, the natural geometry for representing flow in an annulus. In the experiments reported here, the domain is divided into 23 bins in radius

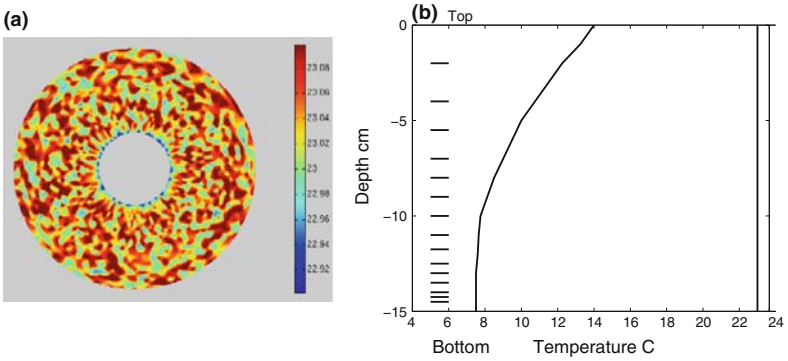

Fig. 5 **a** Random initial conditions are used for the interior temperature field, shown here at a given level. **b** Depth is discretized with variable resolution to enhance resolution near the bottom-boundary. Also shown are temperature *curves* estimated from sparse temperature measurements on the boundary and used as lateral boundary conditions. The *bottom* boundary condition is one of zero heat flux

(0.65 cm/bin) and 120 bins in azimuth (3° bins). The vertical coordinate is discretized using 15 levels non-uniformly distributed over the 15 cm depth of the fluid, as shown in Fig. 5b. The MIT-GCM discretizes variables on an Arakawa C-grid (Arakawa and Lamb 1977). Momentum equations are time-stepped using a second-order Adams Bashforth technique and, in the calculations presented here, θ is advected with an upwind-biased direct space-time technique using a Sweby flux-limiter (Sweby 1984). The treatment of vertical transport is implicit. A 2D equation for the surface pressure field is solved at each timestep using a conjugate gradient method ensuring that the flow remains non-divergent.

We initialize the model with a uniform temperature field to which a small random component is added to initiate hydrodynamical instability. A 2-d horizontal slice is shown in Fig. 5c. The model performs in better than realtime. On one processor of an Altix350, we can produce a single 10-second model simulation within 8 s. The use of non-uniform discretization of the domain using variable vertical levels enables economies to be made in model resolution without compromising resolution where it matters. In the coupled system, multiple simulations are performed (see Sect. 4). Four processors are used for implementing the "MODEL" and "DA" components in Fig. 2.

In Fig. 6, the model horizontal currents are overlaid on the measured velocities. This is done by converting cylindrical model velocities to cartesian velocities, projecting them into screen coordinates, and interpolating using a radial-basis function/splines. The projection matrices are obtained by manually registering the model's geometry with the screen coordinates of the annulus.

We note that despite an obvious uncertainty in initial conditions and other approximations, the model is capable of capturing the gross character of flow in the physical fluid. In addition to errors in initial conditions, and similar to large-scale scenario, model errors will exist. For example, the surface drag at the top is not modeled, heat may leak

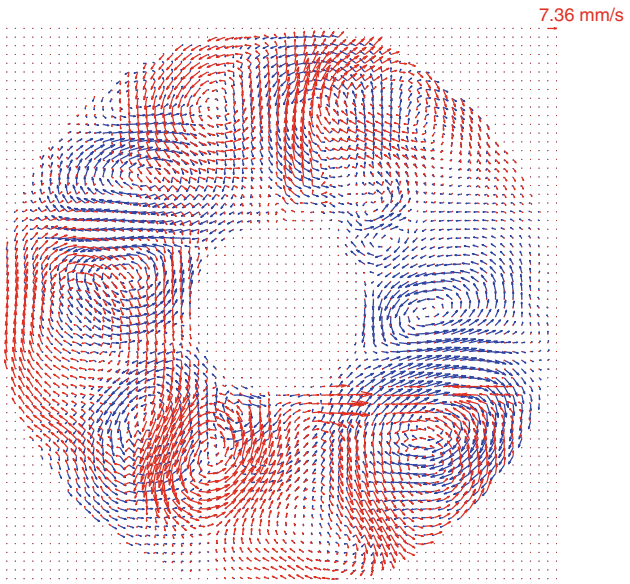

Fig. 6 The measured horizontal velocity (*red*) at a depth of 100 mm from the top of the tank 300 s after cooling commences and the horizontal model velocities at the corresponding time (*blue*). This marks the start of an assimilation experiment

from the bottom (despite an insulating rubber pad) and lateral boundary conditions are imperfect. Thus, it is expected that many flow details will be different. It will be necessary to use measurements to continually constrain the model.

4 State estimation

At a rotation period of 6 s, fluid parcels can traverse the annulus at up to 2 cm/s, with eddy length scales of 9–12 cm. The doubling time of the Eady growth rate (Pedlosky 1987) of instability is $\sim$10 s. It is used as the time period for an assimilation cycle. Since it typically takes 8 real-seconds for a 10 s numerical simulation, and 5 s for observation (in parallel), there are 2 s left for

communication and computational activities, before which the next forecast must be initiated. This is accomplished by sequential filtering using additional processors, and is described next.

Let $\vec{X}_t = [\vec{v}_h; \vec{\theta}]$ be the state[1] at a discrete time t, and measurements $\vec{Y}_t$ be assumed to arise from a linear observation equation $\vec{Y}_t = \mathbf{H}\vec{X}_t + \vec{v}_t$, where the observational noise is normally distributed with zero mean and diagonal covariance $\mathbf{R}_t$, that is $\vec{v}_t \sim \mathcal{N}(0, \mathbf{R}_t)$. Further, let $\vec{X}_t^{\mathrm{f}}$ be the model forecast, with error covariance $\mathbf{P}_t^{\mathrm{f}}$. Now the well-known update equation for analysis state $\vec{X}_t^{\mathrm{a}}$ can be written as:

$$\vec{X}_t^{\mathrm{a}} = \vec{X}_t^{\mathrm{f}} + \mathbf{P}_t^{\mathrm{f}}\mathbf{H}^T(\mathbf{H}\mathbf{P}_t^{\mathrm{f}}\mathbf{H}^T + \mathbf{R})^{-1}\left[\vec{Y}_t - \mathbf{H}\vec{X}_t^{\mathrm{f}}\right] \tag{5}$$

$$\vec{X}_t^{\mathrm{a}} = \vec{X}_t^{\mathrm{f}} + \mathbf{C}_t[\vec{Y}_t - \mathbf{H}\vec{X}_t^{\mathrm{f}}] \tag{6}$$

As shown in Gelb (1974) the Kalman and extended-Kalman filter are given by Eqs. 5 and 6. A dimensionality issue, however, often arises because computing and propagating the covariance explicitly may be numerically unfeasible even for modest sized domains. Therefore, we seek an approach that produces effective estimates while ameliorating the dimensionality problem. One way to address the problem is through domain decomposition (Demmel et al. 1997).

Another way is to use a reduced-rank spectral approximation of the forecast uncertainty. In the Ensemble Kalman Filter (Evensen 2003) method, for example, an ensemble of estimates at time $t - \Delta t$ are forecast to time t using the model. Since the filter operates at time t, we will drop the notation's explicit dependence of time. Let us call the forecast ensemble $\mathbf{V}^{\mathrm{f}} = [\vec{X}_1^{\mathrm{f}} \ldots \vec{X}_S^{\mathrm{f}}]$, where the columns of $\mathbf{V}^{\mathrm{f}}$ are the S samples of the ensemble of horizontal velocities at an observed layer. Thus, if we let $\mathbf{V}^{\mathrm{o}}$ represent a S-column matrix of perturbed observations[2], obtained by perturbing an observation $\vec{Y}$ with noise $\vec{v}$, and $\widetilde{\mathbf{V}}^{\mathrm{f}}$ be the deviation from mean[3] $\bar{\mathbf{V}}^{\mathrm{f}}$ of $\mathbf{V}^{\mathrm{f}}$, the update equation can be written as:

$$\mathbf{V}^{\mathrm{a}} = \mathbf{V}^{\mathrm{f}} + \mathbf{P}^{\mathrm{f}}\mathbf{H}^{\mathrm{T}}(\mathbf{H}\mathbf{P}^{\mathrm{f}}\mathbf{H}^{\mathrm{T}} + \mathbf{R})^{-1}[\mathbf{V}^{\mathrm{o}} - \mathbf{H}\mathbf{V}^{\mathrm{f}}] \tag{7}$$

$$= \mathbf{V}^{\mathrm{f}} + \widetilde{\mathbf{V}}^{\mathrm{f}}(\mathbf{H}\widetilde{\mathbf{V}}^{\mathrm{f}})^{\mathrm{T}}[\mathbf{H}\widetilde{\mathbf{V}}^{\mathrm{f}}(\mathbf{H}\widetilde{\mathbf{V}}^{\mathrm{f}})^{\mathrm{T}} + \widetilde{\mathbf{V}}^{\mathrm{o}}\widetilde{\mathbf{V}}^{\mathrm{oT}}]^{-1}(\mathbf{V}^{\mathrm{o}} - \mathbf{H}\mathbf{V}^{\mathrm{f}}) \tag{8}$$

$$= \mathbf{V}^{\mathrm{f}}\aleph \tag{9}$$

The posterior (or analysis) distribution is represented by mean $\bar{\mathbf{V}}^{\mathrm{a}}$ and covariance $\mathbf{P}^{\mathrm{a}} = \frac{1}{S-1}\widetilde{\mathbf{V}}^{\mathrm{a}}\widetilde{\mathbf{V}}^{\mathrm{aT}}$. This method is very useful because (a) the model is never linearized as in an extended Kalman Filter. (b) Covariance is never propagated explicitly. (c) The update equation is a weakly nonlinear combination of the forecasts. (d) The mixing matrix $\aleph$ can be computed very efficiently using square-root representations and will have very low-size (typically $S \times S$). For highly nonlinear systems, the large number of Monte-Carlo simulations necessary to capture the forecast uncertainty are often computationally not feasible. When only a few ensemble members are used, it is well-known that the forecast covariance can contain spurious long-range correlations. Thus, a localized version of the ensemble Kalman filter that filters out long-range correlations is often implemented, which in our paper is again based on domain decomposition.

Our estimation method consists of two phases. The first phase, initialization, seeks to reduce a large initial uncertainty in the model state to a level where model-states and observations can be thought of as arising from similar distributions. Initialization is based on an engineered forecast error-covariance, and it is not propagated across time (Eq. 6). Once initialized, we switch to the second phase, called tracking, where an ensemble method is used. In both phases, domain-decomposition is used; for addressing dimensionality in initialization and for removing long-range correlations in tracking. Thus, localized versions of Eqs. 6 and 9 will be implemented (see Fig. 7). We now go on to discuss these steps in detail.

4.1 Initialization

We spin up a single model simulation from a random initial temperature field (see Fig. 5). After a transient period has elapsed, the initialization phase commences, and is repeated for a few assimilation cycles. The initialization phase consists of four steps, executed in sequence:

4.1.1 Interpolation in the vertical

An interpolation function of horizontal velocities and temperature is estimated from the forecast. Let $\vec{v}_h^{\mathrm{f}}[i,j,k]$ be the forecast horizontal velocity at grid node i, j, k in the radial, azimuthal and vertical directions, respectively. Let $\vec{v}_h^{\mathrm{f}}[i,j]$ be the column-vector of forecast velocities at all $N_z = 15$ vertical levels corresponding to horizontal grid location i, j and let $\vec{v}_h^{\mathrm{fo}}[i,j]$ be the corresponding vector of horizontal velocities at the $N_o = 5$ observed vertical levels. Similarly construct vectors $\vec{\theta}^{\mathrm{f}}[i,j]$ and $\vec{\theta}^{\mathrm{fo}}[i,j]$ from the forecast temperatures. Using samples in the forecast, we

[1] The state for assimilation consists of the horizontal velocities and temperature. Vertical velocity is implicit, pressure is diagnostic and salinity is unrepresented.

[2] This formulation is discussed for its simplicity. Other variations, e.g. explicit inverse, will be useful for small state sizes.

[3] $\bar{\mathbf{V}}^{\mathrm{f}} = \frac{1}{S}\sum_{i=1}^{S}\mathbf{V}^{\mathrm{f}}[:,i]$

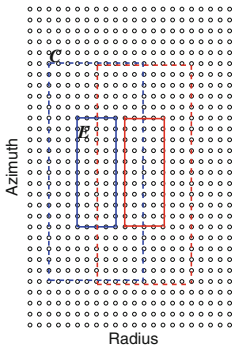

Fig. 7 The estimation using the ensemble Kalman filter is localized within estimation windows E, influenced by observations from overlapping spatial-context windows C

estimate the matrices $\lambda_{\mathbf{v}}$ and λ_θ by solving equations of the form $\vec{v}_{\mathrm{h}}^{\mathrm{f}}[i,j] = \lambda_{\mathbf{v}}\vec{v}_{\mathrm{h}}^{\mathrm{fo}}[i,j]$ and $\vec{\theta}_{\mathrm{h}}^{\mathrm{f}}[i,j] = \lambda_\theta\vec{\theta}_{\mathrm{h}}^{\mathrm{fo}}[i,j]$.

4.1.2 Estimating horizontal velocities at observation layers

At each observed layer ($k_{\mathrm{o}} \in \{k_1 \ldots k_5\}$) of the fluid, initialization occurs with a deterministic scheme. Since this step is repeated for each observation level, it is sufficient to consider the assimilation at any single observed layer k_{o}. At every location i, j on the horizontal grid ($N_{\mathrm{r}} = 23 \times N_\phi = 120$) of an observed layer, we estimate the horizontal velocity from forecasts and observations using a spatial context of dimensions $N_{\mathrm{r}}^{\mathrm{l}}$ radially and N_ϕ^{l} azimuthally. The estimation is written as:

$$\vec{v}_{\mathrm{h}}^{\mathrm{a}}[i,j,k_{\mathrm{o}}] = \vec{v}_{\mathrm{h}}^{\mathrm{f}}[i,j,k_{\mathrm{o}}] + \mathbf{P}_i^{*}\mathbf{H}_{ij}^{\mathrm{T}}(\mathbf{H}_{ij}\mathbf{P}_i^{\mathrm{f}}\mathbf{H}_{ij}^{\mathrm{T}} + \mathbf{R}_{ij})^{-1}\left[\vec{v}_{\mathrm{h}}^{\mathrm{o},ijk_{\mathrm{o}}} - \mathbf{H}_{ij}\vec{v}_{\mathrm{h}}^{\mathrm{f},ijk_{\mathrm{o}}}\right] \quad (10)$$

$$\vec{v}_{\mathrm{h}}^{\mathrm{a}}[i,j,k_{\mathrm{o}}] = \vec{v}_{\mathrm{h}}^{\mathrm{f}}[i,j,k_{\mathrm{o}}] + \mathbf{C}_{ij}[\vec{v}_{\mathrm{h}}^{\mathrm{o},ijk_{\mathrm{o}}} - \mathbf{H}_{ij}\vec{v}_{\mathrm{h}}^{\mathrm{f},ijk_{\mathrm{o}}}] \quad (11)$$

Here, $\vec{v}_{\mathrm{h}}^{\mathrm{f},ijk_{\mathrm{o}}}$ is the vector of forecast horizontal velocities in a $N_{\mathrm{r}}^{\mathrm{l}} \times N_\phi^{\mathrm{l}}$ area centered[4] at grid node i, j, k_{o}, and $\vec{v}^{\mathrm{o},ijk_{\mathrm{o}}}$ are available observations in the same area. The local forecast covariance $\mathbf{P}_i^{\mathrm{f}}$ is generated using a 2D Gaussian per velocity component. It only varies radially (so as to account for annulus borders) but not in depth k_{o} or azimuth j. Each local observation operator $\mathbf{H}_{ij}$ selects locations where observations are valid in the corresponding $N_{\mathrm{r}}^{\mathrm{l}} \times N_\phi^{\mathrm{l}}$ region. The matrix $\mathbf{R}_{ij}$ is the corresponding observational uncertainty. We typically choose up to $N_{\mathrm{r}}^{\mathrm{l}} = 5$ and $N_\phi^{\mathrm{l}} = 10$; motivated by the estimated auto-correlation length-scales in the corresponding directions. Each $\mathbf{C}_{ij}$ is at most 2×100 and is constructed a priori[5]. The vector $\vec{v}_{\mathrm{h}}^{\mathrm{f}}[i,j,k_{\mathrm{o}}]$ is the forecast horizontal velocity at location i, j, k_{o} and $\vec{v}_{\mathrm{h}}^{\mathrm{a}}[i,j,k_{\mathrm{o}}]$ is the corresponding estimated (sometimes called assimilated or analysis) horizontal velocity.

[4] Except near annulus boundaries, where the window is off-center.

[5] A large number of matrices $\mathbf{C}_{ij}$ are identical, thus saving storage costs.

4.1.3 Estimating temperature at observation layers

Once the horizontal velocities $\vec{v}_{\mathrm{h}}^{\mathrm{a}}[i,j,k_{\mathrm{o}}]$ are estimated at each grid node of observed layers, we compute temperature $\theta^{\mathrm{a}}[i, j, k_{\mathrm{o}}]$ by solving the elliptic equation obtained by taking the divergence of the *thermal-wind* equation at each observed layer Pedlosky (1987), given by:

$$\frac{\partial\vec{v}_{\mathrm{h}}}{\partial z} = \frac{g\alpha}{2\Omega}\hat{k} \times \nabla\theta \quad (12)$$

Here, α is the coefficient of thermal expansion and $\hat{k}$ is the vertical unit-vector.

4.1.4 Estimate full State

The precomputed vertical interpolation models are applied to the estimated horizontal velocity and temperature. Thus, we estimate $\vec{v}_{\mathrm{h}}^{\mathrm{a}}[i,j] = \lambda_{\mathbf{v}}\vec{v}_{\mathrm{h}}^{\mathrm{ao}}[i,j]$ and $\vec{\theta}^{\mathrm{a}}[i,j] = \lambda_\theta\vec{\theta}^{\mathrm{ao}}[i,j]$, where these vectors are defined analogously to step 1 (but using the analysis fields).

The estimated fields become the new state $\vec{X}_t = [\vec{v}_{\mathrm{h}}^{\mathrm{a}}; \vec{\theta}^{\mathrm{a}}]$ for the next forecast. We repeat this four step process for a few assimilation cycles and then switch to a flow-dependent ensemble tracking method that can both estimate states and their uncertainties, as discussed next.

4.2 Tracking

Throughout the tracking phase, the steps **1, 3**, and **4** remain the same and thus are not discussed again. The only difference between initialization and tracking is the process of constraining horizontal velocities at observed layers. For tracking, we use a variation of the ensemble Kalman filter in the following way:

4.2.1 Creating the ensemble

The two prominent sources of uncertainty are the thermal boundary condition that drives the numerical system and the flow uncertainty due to time-staggered observations and numerical integration. To model these, we use the output of the initialization step to drive several simulations, each utilizing a thermal boundary condition perturbed from the climatological profile (see Sect. 3). Additionally, motivated by the method of snapshots (Sirovich 1987), we also save the state every few time steps in the forward integration of a simulation. The forecast ensemble is

therefore constructed as a *mixture* of two distributions, one representing boundary condition uncertainty (multiple simulations) and the other due to uncertainty in flow (snapshots during the model integration). Assuming there are N_s snapshots and N_b simulations, we have an ensemble of $S = N_s N_b$ forecast samples. These samples are used for estimation, as we now describe.

4.2.2 Localized estimation

Akin to the localization during deterministic initialization, we also localize the ensemble Kalman filter during tracking. Estimation at each observed horizontal layer of the fluid k_o follows the illustration in Fig. 7. Estimates are produced separately for each component of horizontal velocity in an *estimation window* E of size $N_r^e \times N_\phi^e$ indexed by location i_e, j_e, k_o, using forecasts and observations in a spatial context window C that is indexed by location i_c, j_c, k_o and of size $N_r^c \times N_\phi^c$. Estimates over an entire layer are produced by tiling it with estimation windows (no overlap). Note, however, that adjacent estimation windows share substantial spatial context, as shown in Fig. 7.

Let $\mathbf{V}^{f,i_c j_c k_o}$ be the matrix representing forecast horizontal velocity component of S ensemble members coincident with the estimation window E at i_e, j_e, k_o, and $\mathbf{V}^{f,i_c j_c k_o}$ be the matrix of forecast horizontal velocity components of S ensemble members coincident with the context window C at i_c, j_c, k_o. Using the observations $\mathbf{V}^{o,i_c j_c k_o}$ and forecasts in the context window to construct $\aleph_{i_c j_c k_o}$, we may express the analysis ensemble $\mathbf{V}^{a,i_e j_e k_o}$ as:

$$\mathbf{V}^{a,i_e j_e k_o} = \mathbf{V}^{f,i_e j_e k_o} \aleph_{i_c j_c k_o} \tag{13}$$

In practice, because only the analysis corresponding to the last snapshot of the current forecast of each simulation is necessary to launch the next forecast, $\aleph_{i_c j_c k_o}$ need only be $S \times N_b$ in size, with an appropriately ordered ensemble. Note that our approach is related to LEKF (Ott et al. 2003), but with substantial differences in how estimation and context windows are designed and used. A single assimilation is completed within the realtime constraints.

5 Experiments

For the experiments presented here, the reference density $\rho_0 \approx 1,037$ kg m^{-3}, the rotation rate is $\Omega = 1.15$ rad/s, the annulus width $L = 0.15$ m, the mean fluid depth $D = 0.15$ m, and the mean temperature difference of fluid across the annulus $\Delta T = 6$ K (measured separately). The viscosity is $\nu = 10^{-6}$ m^2 s^{-1}, the thermal diffusivity $\kappa = 10^{-7}$ m^2 s^{-1}, and the thermal expansion coefficient $\alpha = 3 \times 10^{-4}$ K^{-1}. Thus, the Ekman number $E = \frac{\nu}{2\Omega D^2} = 1.9 \times 10^{-5}$, the thermal Rossby number $R_\theta = \frac{g\alpha\Delta T D}{\Omega^2 L^2} = 0.09$, the Prandtl number $\mathcal{P}_\sigma = \frac{\nu}{\kappa} = 10$. In comparison, assuming an average of seven planetary waves on the 45°N latitude circle, a mean tropopause height of 13.5 km, a temperature differential of 30 K around a mean of 288 K, and an eddy viscosity of 1 m^2/s, the Ekman number is 5×10^{-5} and the Thermal Rossby number is 0.08. Both are in good agreement with the experimental regime and appropriately small. The Prandtl number of the atmosphere (~ 1) is based on the existence of a turbulent boundary layer that is not present in the annulus experiment.

We cool the core after the fluid attains solid body rotation. A circulation is established in about 300 s, and an example of a well-formed circulation is shown in Fig. 6 at a layer 100 mm below the water surface.

The MIT-GCM is started from a random initial condition with a climatological thermal-boundary condition shown in Fig. 5. Using the parameters described in Sect. 3, the model is integrated forward to remove transients and establish a circulation, albeit unconstrained with measurements. The horizontal velocity field at 100 mm below the top of the tank is shown in Fig. 6 along with corresponding observations. It shares the gross characteristics of the circulation, but the waves have the wrong phase and incorrect amplitudes. Over several experiments, we note that model velocities can be as much as twice that of the measured velocities.

We then turn on the assimilation component. The local observational uncertainty $\mathbf{R}_{ij} = \sigma_o^2 \mathbf{I}$. This uncertainty can arise due to a number of factors. The dominant source is vibration, and we note a 0.5 pixel jitter when tracking a calibration grid. This implies a velocity uncertainty of $\sim \sigma_o = 1.2$ mm/s.

During initialization, the covariance $\mathbf{P}_i^f$ is constructed as an un-normalized 2D Gaussian per velocity component with standard deviation of 1 (radially) and 2 azimuthally, with extent of 5 grid nodes (radially) and 10 grid nodes (azimuthally). The Gaussian is initially scaled by an amplitude of $\sigma_b = \sigma_o * 2$, to account for the fact that unconstrained model velocities have less skill than the observations. The observation operator $\mathbf{H}_{ij}$ admits grid points in the domain outside the shadow region and where observations pass a simply quality control of being less than 3 cm/s. Doing so excludes impulse noise, seen for example at the edges of the shadow region in Fig. 6.

With these parameters, the deterministic assimilation scheme is run till the root mean square error between measured and forecast horizontal velocities over is at least within $1.5*\sigma_o$. This takes ~ 3 assimilation cycles.

After the initialization, the system alternates with an ensemble scheme. We run different simulations and each of them start from the model-state estimated during initialization, but with a perturbed temperature boundary profile and initial condition. Each simulation runs on a separate processor of the Altix350, and integrates the model 10 s

forward in ~ 8 s of clock-time. Snapshots of the model-state (horizontal velocity and temperature) are extracted from each simulation during the integration. Thus, at the end of the 10- second period, an ensemble becomes available. The final forecast (at $t = 10$ s) is used to estimate the interpolation functions in the vertical. Observations in the immediately preceding 5 s are used in the ensemble assimilation scheme discussed in Sect. 4. The observational uncertainty is identical to the deterministic case (forecast covariance is inflated). We choose at most $N_{\mathrm{r}}^{\mathrm{c}} = 11$, and nominally select $N_{\mathrm{r}}^{\mathrm{c}} = N_{\mathrm{r}}^{\mathrm{e}} = 6$ and $N_{\phi}^{\mathrm{c}} = N_{\phi}^{\mathrm{e}} = 5$ with $S = 30$. Figure 8 shows the estimated horizontal velocities and observations after 10 assimilation cycles at a depth of 100 mm from the top of the tank. The estimate depicted here corresponds to the last snapshot of the simulation with a climatological thermal boundary condition profile in Fig. 5. The final time estimated model-states are used to re-initialize it for the next 10 s forecast.

Figure 9 shows the evolving root mean square (RMS) error between the estimated and observed velocities over 30 assimilation cycles in a 300-second assimilation experiment. Please note that this graph depicts the likelihood and not the a posteriori error between the estimate and truth, because the truth is unknown. Nevertheless, it is a useful measure in that it shows model velocities come close to the measurements nearing the inherent uncertainty ($\sigma_{\mathrm{o}} = 1.2$ mm/s) with which they are observed. Indeed both the amplitudes and phase are in good agreement, as can be seen in Fig. 8. After 20 assimilation cycles, we turn the assimilation off and simply compute the error between forecast velocities and experimental measurements. As expected this error grows, and saturates in around 10 cycles. Figure 10 shows the model and measured velocities at $t = 300$ s for the ensemble member corresponding to Fig. 8, at 100 mm above the bottom of the tank. The model has once again departed from the system trajectory. Similarly configured experiments suggest that it takes ~ 10 rotation periods or six assimilation cycles before the model adjusts itself to be consistent with observations.

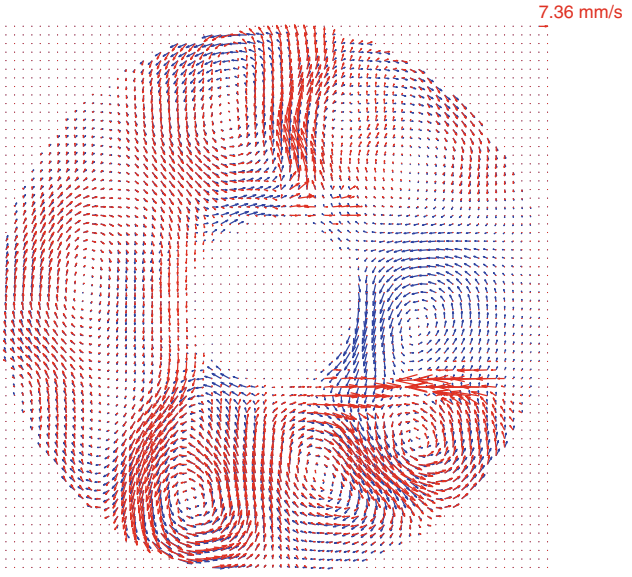

Fig. 8 The horizontal model velocity field of an ensemble member, $t = 200$ s into the assimilation experiment, ~ 100 mm below the surface (*blue*), and corresponding measurements (*red*)

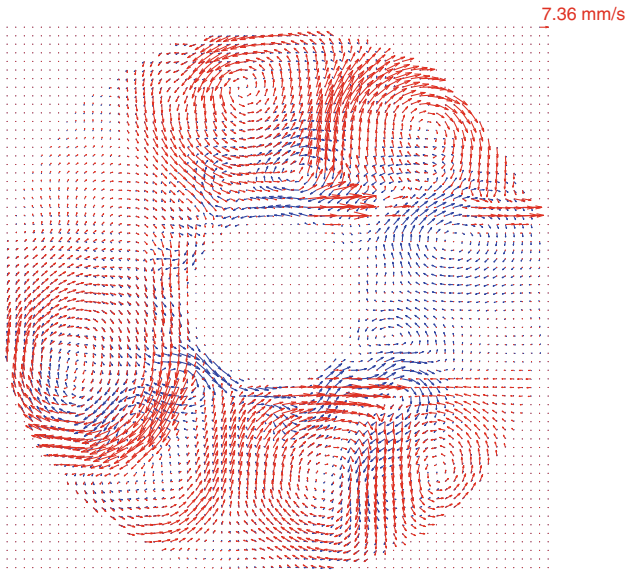

Fig. 9 Once the assimilation is terminated the model diverges from the observations. Shown here is the model velocity for an ensemble member at 100 mm below the top of the tank (*blue*) and corresponding measurements (*red*) at $t = 300$ s

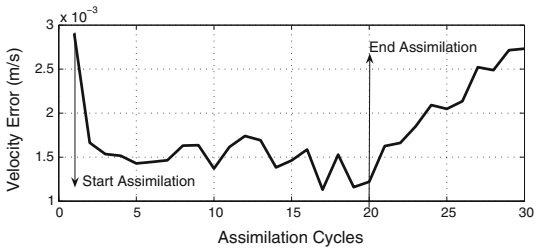

Fig. 10 The RMS-error between estimated and measured velocity at observed location as a function of time

6 Discussion and conclusions

The coupled physical-numerical system described here is an effective way to study a variety of rotating flows. In particular, it can accommodate flows with a wide range of thermal boundary conditions and rotation periods. The hybrid assimilation scheme is motivated by several considerations. Early analysis showed that a variational approach (Wunsch 1996) would not meet realtime needs and that an ensemble-filter provided the best prospect, if a large

number of numerical simulations is to be avoided. It is in this sense that initialization and tracking are synergistic. Initialization helps condition forecast uncertainty, after which snapshots capture the smaller of the uncertainties within the tracking loop and boundary-condition perturbations capture the larger uncertainty of the boundaries. In fast-evolving flows, the flow uncertainty starts to dominate, but in slowly evolving flows, the boundary-condition uncertainty dominates. In any flow situation, the use of the proposed scheme prevents an ensemble collapse by maintaining a justifiable representation of the uncertainty. Further, the proposed representation requires fewer numerical simulations than purely sampling initial conditions and produces well-ranked ensembles during assimilation.

Our system scales to a variety of experiments and flows. The PIV and MIT-GCM are parallelizable beyond that described here. In our assimilation approach, localization not only prevents spurious long-range correlations but also lends to an easily parallelizable algorithm. Updates in individual windows are performed in parallel. Realtime performance is achieved here through parallelism (observations), domain-decomposition (model, estimation), spectral-reduction (estimation) an efficient method to generate samples and compute updates (estimation).

There is, however, a computational trade-off in localization. If we choose a window with W grid nodes and have S ensemble members at hand, the complexity of initialization is $O(W^2)$ and of tracking is $O(WS^2)$. Thus, we may prefer a smaller window (initialization) or ensemble size (tracking). Care must be taken nonetheless because the window size must not be so small as to lose the advantages that correlations in the model's field offers estimation and the same argument holds for the ensemble size. The present system can thus scale with the addition of computational resources.

There are also several limitations of the existing system. The domain boundaries are not resolved at high resolution, which may be essential for certain flows. Adaptive resolution in PIV, the model and assimilation is a promising direction. Temperature measurements have not been used, except to provide climatological temperature boundary conditions. Newer methods for whole-field LIF measurements or sparse measurements for assimilation or verification would be useful. We presently observe 5 layers, in large-part due to the latency associated with physical motor movement. A newer periscope design with a rotating mirror and paraboloid will improve the scan speed many fold. The assimilation method uses a fixed local context. A multiscale extension and comparisons with contemporary methods is beyond the scope of this paper but will appear in a forthcoming article.

Even without these improvements, our observatory works remarkably well to produce state estimates consistent with the observations in the current application, and with largely off-the-shelf and relatively inexpensive components. Thus, the analog serves as a new, easy-to-use, testbed to explore annulus dynamics and analysis techniques. To the best of our knowledge, a realtime observatory of this kind has not been achieved before.

Acknowledgments This work is funded by CNS-0540259 and NSF Grant CNS-0540248. The authors thank Ryan Abernathy for helping with the hardware platform development.

References

Arakawa A, Lamb V (1977) Computational design of the basic dynamical processes of the ucla general circulation model. Methods Comput Phys 17:174–267 (Academic Press)

Demmel JW, Demmel JW, Heath MT, Heath MT, Vorst HAVD, Vorst HAVD (1997) Applied numerical linear algebra. SIAM

Evensen G (2003) The ensemble kalman filter: theoretical formulation and practical implementation. Ocean Dyn 53:342–367

Geisler JE, Pitcher EJ, Malone RC (1983) Rotating-fluid experiments with an atmospheric general circulation model. J Geophys Res 88(C14):9706–9716

Gelb A (1974) Applied optimal estimation. MIT Press, Cambridge, MA, USA

Hide R (1958) An experimental study of thermal convection in a rotating liquid. Phil Trans Roy Soc A250:441–478

Lee C (1993) Basic instability and transition to chaos in a rapidly rotating annulus on a beta-plane. PhD thesis, University of California, Berkeley

Lorenz EN (1963) Deterministic nonperiodic flow. J Atmos Sci 20:130–141

Marshall J, Adcroft A, Hill C, Perelman L, Heisey C (1997a) A finite-volume, incompressible navier stokes model for studies of the ocean on parallel computers. J Geophys Res 102(C3):5753–5766

Marshall J, Hill C, Perelman L, Adcroft A (1997b) Hydrostatic, quasi-hydrostatic and nonhydrostatic ocean modeling. J Geophys Res 102(C3):5733–5752

Morita O, Uryu M (1989) Geostrophic turbulence in a rotating annulus of fluid. J Atmos Sci 46(15):2349–2355

Ott E, Hunt BR, Szunyogh I, Zimin A, Kstelich E, Corazza M, Kalnay E, Patil DJ, Yorke JA (2003) A local ensemble kalman filter for atmospheric data assimilation. Technical report. arXiv:physics/0203058 v4, University of Maryland

Pedlosky J (1987) Geophysical fluid dynamics. Springer, New York

Ravela S, Hansen J, Hill C, Marshall J, Hill H (2003) On ensemble-based multiscale assimilation. In: European geophysical union annual congress

Ravela S, Marshall J, Hill C, Wong A, Stransky S (2007) A real-time observatory for laboratory simulation of planetary circulation. In: Lecture notes in computer ccience, vol 4487, pp 1155–1162

Read PL (2003) A combined laboratory and numerical study of heat transport by baroclinic eddies and axisymmetric flows. J Fluid Mech 489:301–323

Read PL, Bell MJ, Johnson DW, Small RM (1992) Quasi-periodic and chaotic flow regimes in a thermally driven, rotating fluid annulus. J Fluid Mech 238:599–632

Read PL, Thomas NPJ, Risch SH (2000) An evaluation of eulerian and semi-lagrangian advection schemes in simulations of rotating, stratified flows in the laboratory. part i: axisymmetric flow. Mon Weather Rev 128:2835–2852

Sirovich L (1987) Turbulence and the dynamics of coherent structures, part 1: cohrerent structures. Q Appl Math 45(3):561–571

Sweby PK (1984) High resolution schemes using flux-limiters for hyperbolic conservation laws. SIAM J Numer Anal 21: 995–1011

Tajima T, Nakamura T (2003) Experiments to study baroclinic waves penetrating into a stratified layer by a quasi-geostrophic potential vorticity equation. Experiments in Fluids 34(6):744–747

von Larcher T, Egbers C (2005) Experiments on transitions of baroclinic waves in a differentially heated rotating annulus. Nonlinear Process Geophy 12:1044–1041

Wunsch C (1996) The Ocean Circulation Inverse Problem. Cambridge University Press, Cambridge, UK

Young R, Read P (2006) Breeding vectors in the rotating annulus as a measure of intrinsic predictability. In: Royal Met. Soc. Annual Student Conference