

# Tracking Rotating Fluids in Realtime using Snapshots\*

Sai Ravela      John Marshall      Chris Hill      Andrew Wong  
Scott Stransky  
Earth, Atmospheric and Planetary Sciences,  
Massachusetts Institute of Technology  
ravela@mit.edu

## Abstract

*We present a model-based system for tracking rotating fluids, and apply it to a laboratory study of atmospheric circulation. Tracking is accomplished by filtering uncertain and high-dimensional states of a nonlinear general circulation model with optical measurements of the physical fluid's velocity. Realtime performance is achieved by using a nonuniform discretization of the model's spatial resolution, and by using time-snapshots of model-state to construct spatially-localized reduced-rank square-root representations of forecast uncertainty.*

*Realtime performance, economical and repeatable experimentation, and a direct connection to planetary flows implies that the proposed physical-numerical coupling can be useful for addressing many perceptual geophysical fluid dynamics problems. To the best of our knowledge, such a system has not hitherto been reported.*

## 1. Introduction

We present a coupled physical-numerical system for understanding rotating fluids<sup>1</sup> in the laboratory. This system contains sensors to take measurements from the physical fluid, a numerical 3-D fluid-model for forecasting, and inference algorithms that constrain the model with observations in *realtime*. Such a system opens up many exciting possibilities for addressing problems in perceptual geophysical fluid dynamics, including:

**Model Parametrization:** Using the coupled system, reality can be augmented with numerical quantities that cannot be observed. Thus it becomes an interactive tool for developing parametrizations (e.g. effective diffusivity) in weather and climate modeling.

---

\*This paper is supported in part by grants NSF CNS 0540248 and NSF CNS 0540259.

<sup>1</sup>A rotating fluid is subject to a global rotation, e.g. our oceans and atmosphere.

**Autonomous Observatories:** The coupled physical-numerical system can be used to devise robust planning and control algorithms for autonomous sampling of the oceans and atmosphere. Algorithms developed in the laboratory will have to account for uncertainties very similar to those in the real world (erroneous models, observations and tracking), thus providing a valuable prototyping tool.

**Numerical Weather Prediction:** Constraining models with observations is fundamental to ocean state estimation and numerical weather prediction. In these applications, predictions are made using general circulation models (GCMs) implementing the discretized governing equations of fluid motion. We know that when initial conditions are uncertain, even a perfect model will **not** predict the atmosphere perfectly [11]. Estimating states and their uncertainties is thus the primary means to constrain a GCM [26].

In this paper, we describe the architecture of the coupled-physical numerical system. Because it is out of the scope of this paper to explore all the aforementioned applications, we focus on the state estimation problem in a fluid dynamical analog of atmospheric circulation.

### 1.1. Planet in a Spinning Water Tank

Figure 1 depicts the 500mb heights in a weather map. The cyclonic and anti-cyclonic pressure features transport around the globe whilst winds (imagine a dye) flow along their contours. This *Hadley circulation* can be simulated in the laboratory using a well-known differentially-heated rotating annulus experiment [5]. When a spinning annulus' center is cooled relative to a warm periphery, the water near the center becomes dense and sinks. Warm waters from the periphery move in to replenish, thus setting up a radially overturning circulation. In the presence of Coriolis force and thermal-wind [16] flow becomes unsteady at strong temperature differences and high rotation rates, forming eddies and jets like the atmosphere (see Figure 1). Flow features form and evolve via an energy-entropy cascade [16], and differ from one physical experiment to the next.

Tracking this laboratory flow is convenient and use-




Figure 1. Image (a) shows the 500hPa 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 similar. We study the state-estimation problem for planetary flows using the laboratory analog.

ful because repeatable experiments with real data can be performed using far simpler logistics than the operational weather-prediction setting. Yet, key challenges associated with the large-scale problem must also be addressed in the laboratory: (a) Nonlinearity — the laboratory analog is nonlinear and the numerical model is the same used in planetary simulations. (b) Dimensionality — the size of the state of the numerical model is of the same order as planetary simulations. (c) Uncertainty — the initial conditions are unknown, and the model is imperfect relative to the physical system. (d) Realtime — forecasts must be produced in better than realtime. Solutions to these problems can accelerate operational acceptance of new methods as well as inform estimation and inference in many other coupled systems.

Our coupled-system operates by filtering states of the MIT General Circulation Model [13, 12] with optical measurements of the physical fluid’s velocity. Tracking is in realtime, achieved by using a nonuniform discretization of the model’s spatial resolution, and by using time-snapshots of model-state to construct spatially-localized reduced-rank square-root representations of forecast uncertainty. To the best of our knowledge, this is the first such system in existence. It is in routine use and data-sets are readily available to other researchers.

In the remainder of this paper, we relate to other work (Section 2), and present the coupled system components: the observatory (Sections 3 and 4), the numerical model (Section 5) and the estimation scheme(Section 6). A discussion of experimental results follows<sup>2</sup> (Section 7).

<sup>2</sup>Supplementary material includes high-resolution versions of figures and pointers to video.

## 2. Related Work

Thermally-driven rotating flows [10, 21, 5, 4, 14] are well-known analogs of large-scale circulation. It is a robust experiment, easily conducted in the laboratory, and has been used to study a variety of physical phenomena including geostrophic turbulence [14], convection [10], baroclinic instability [20], and chaos [21], and as a test-bed for evaluating the utility of numerical models [22, 8]. In more recent work [20], numerical studies are combined with laboratory experiments to study heat transport and effort has been afoot to study prediction and predictability problems [27]. To the best of our knowledge, however, this is the first coupled observation/numerical system of a laboratory experiment to operate in realtime.

For simulation, we chose the MIT GCM [13, 12] because it is readily available and belongs to the same class of models used in ocean and atmosphere simulations on a large-scale [6].

For estimation, several possibilities were explored. Variational [3, 26] inference was not amenable to realtime performance. Sequential estimation is appropriate, but smoothing [17, 9, 7, 18] is not. In filtering, nonlinearity precludes the use of the Kalman filter [9] and the state’s high-dimensionality precludes use of the extended Kalman filter [9] as well as the particle filter [2] as well as its localized variants [24]. The ensemble Kalman Filter [7] is thus used and comparisons to a variety of similar filters [15], and the particle filter [2] have been conducted. We adapt the ensemble filter by localizing it and using snapshots [23] to produce multiple samples per simulation.

## 3. Observatory for the Annulus Experiment

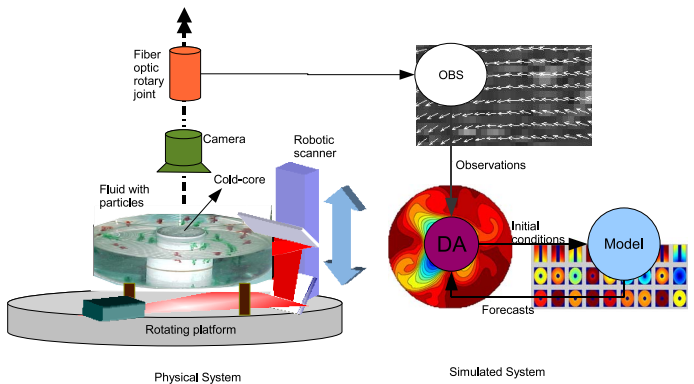


Figure 2. The observatory consists of a physical system including a rotating table on which a tank, camera and control system for illumination are mounted. The computational part consists of a measurement system for velocimetry, a numerical model, and an estimation system, as described more fully in the text.

The observatory, illustrated in Figure 2, has a physical

and computational component. For the annulus experiment, the physical component consists of a perspex annulus of inner radius  $8\text{cm}$  and outer radius of  $23\text{cm}$ , filled with  $15\text{cm}$  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. Fluorescent particles (Dantec Dynamics,  $\text{sg } 1.03\text{g/cc}$ ) are homogenized in saline water of equal density and respond to incident laser illumination. They appear as texture in the 12-bit,  $1\text{K} \times 1\text{K}$  images (Figure 4) taken from 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 Figure 3. The observation rig is carefully mounted and tested for vibrations (the camera must shake by less than  $0.1^\circ$  to have less than 10% motion noise).




Figure 3. The apparatus consists of (a) the rotating platform, (b) the motorized mirror, (c) the tank, (d) electronics, (e) a rig on which a camera is mounted, (g). Laser light comes from direction (f) and bounces off two mirrors before entering the tank. The fiber optic rotary joint (FORJ) (h) allows images to leave the rotating frame and is held stably by bungee chords (i). The square tank prevents the laser light from bending at the annulus interface.

The computational aspects of the observatory are also shown in Figure 2. A server acquires particle images and ships them to two processors that compute optic-flow in parallel (Figure 2, labeled OBS). Measured velocity vectors are passed to an estimation program (Figure 2, labeled DA) that combines them with model forecasts to estimate new states. These estimates become new initial conditions for the models. We now go on to discuss individual components of this system.

#### 4. Visual Observation

We spin up the rotating platform at the desired period (between  $3\text{s}$  and  $12\text{s}$ ). After twenty minutes or so the fluid comes into solid body rotation. The inner core is then cooled using a chiller (see Figure 4). Circulation forms within minutes as the flow becomes unsteady and baroclinic instability[16] sets in.




Figure 4. The camera's view of the rotating annulus in visible light (top-left), the laser illuminating the tank (top-right). The camera's view in laser light shows particles (bottom). Notice shadow due to the chiller in the middle. High-resolution version of this figure is in attachment.




Figure 5. Observed horizontal velocities (green) at  $100\text{ mm}$  above the bottom of the tank after circulation has formed. Maximum flow speeds are of the order of  $2\text{cm s}^{-1}$ .

Once cooling commences, we turn off the lights and turn on the continuous wave  $1\text{W } 532\text{nm}$  laser, which emits a horizontal sheet of light that doubles back through a periscope mounted on the robot to illuminate a horizontal sheet of the fluid volume (see Figure 4). The imaging subsystem in the rotating frame observes the developing flow using a camera looking down at the annulus. The ultra-



small particles move with the flow and thus, optic flow on successive image-pairs is used to measure the horizontal component fluid velocity. An example is shown in Figure 5.

Observations are gathered over several levels repeatedly. The mirror moves to a fluid level, the system captures images, flow is computed, and the mirror moves to the next level and so on. Note that a complete scan takes a few seconds and thus the model-state corresponding to a volumetric observation is somewhat uncertain. In section 6 we will incorporate this uncertainty in estimation.

## 5. Numerical Model

We use the MIT General Circulation Model [12, 13] to solve the equations of an incompressible Boussinesq fluid in hydrostatic balance. The governing equations are:

$$\frac{\partial \vec{v}_h}{\partial t} = G_{v_h} - \frac{1}{\rho_0} \nabla_h p \quad \text{horiz. momentum} \quad (1)$$

$$\nabla_h \vec{v}_h + \frac{\partial w}{\partial z} = 0 \quad \text{continuity} \quad (2)$$

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

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

Here, the three-dimensional velocity is denoted by  $\vec{v} = [\vec{v}_h; w]$  where  $\vec{v}_h$  is the horizontal velocity,  $w$  is the vertical velocity and  $\nabla_h$  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  $\rho_0$  a constant reference value and  $\theta$  is the temperature. The term  $G_{v_h}$  in the horizontal momentum equation includes inertial, Coriolis and frictional terms;  $G_\theta$  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 [13, 12].

No-slip boundary conditions are assumed on all solid boundaries and a linearized free surface is adopted. The temperature at the outer wall of the tank is held constant; at the inner core it is set to an observed vertical profile taken from a separate experiment (see Figure 6(b)). The bottom boundary is assumed to be thermally-insulating.

Finite difference forms of the above equations are solved in cylindrical coordinates, as shown in Figure 6(a), the natural geometry for representing flow in an annulus. In the experiments reported here the domain is divided into 23 bins in radius ( $0.65\text{cm}/\text{bin}$ ) and 120 bins in azimuth ( $3^\circ$  bins). The vertical coordinate is discretized using 15 levels non-uniformly distributed over the  $15\text{cm}$  depth of the fluid, as shown in Figure 6(b). The MIT-GCM discretizes variables on an Arakawa C-grid [1]. Momentum equations are time-stepped using a second-order Adams Bashforth technique and, in the calculations presented here,  $\theta$  is advected

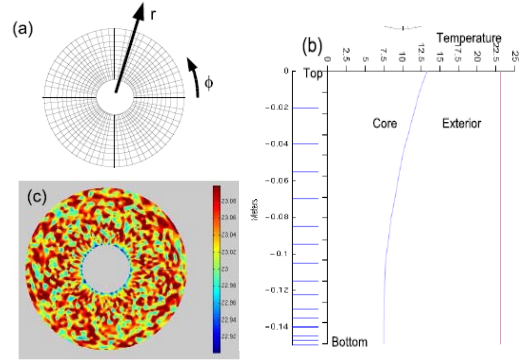


Figure 6. (a) The computational domain is represented in cylindrical coordinates. (b) Depth is discretized with variable resolution to enhance resolution near the bottom-boundary. The lateral boundary conditions on temperature are obtained by interpolating sparse temperature measurements on the boundary. The bottom boundary condition is one of zero heat flux. (c) Random initial conditions are used for the interior temperature field, shown here at a given level.

with an upwind-biased direct space-time technique using a Sweby flux-limiter [25]. The treatment of vertical transport is implicit. A 2-D equation for the surface pressure field is solved at each time-step using a conjugate gradient method ensuring that the flow remains non-divergent.




Figure 7. The planar velocity of a model forecast (in yellow) is shown with observed velocity (green) at a height of  $100\text{mm}$  from the bottom of the tank at the beginning of an estimation experiment. Maximum flow speeds are of order  $2 \text{ cm s}^{-1}$ .

The model is started with a random temperature field to initiate hydrodynamical instability. A 2-D horizontal slice is shown in Figure 6(c). In Figure 7 the model horizon-

tal currents are overlaid on the observed velocities after suitably registering the domain geometry to the physical tank. Despite an obvious uncertainty in initial conditions and other approximations, the model is capable of capturing the gross character of flow observed in the physical fluid, such as typical flow speeds and scales. However, as is to be expected, many flow details are different and observations must be used to constrain the evolving state of the model.

## 6. State Estimation

We develop a two-stage hybrid filter to sequentially reduce uncertainty in state estimates. Define  $\vec{X}_t = [\vec{v}_h(t); \vec{\theta}(t)]$  to be the state<sup>3</sup> at a discrete time  $t$ , and  $\vec{Y}_t$  as measurements, assumed to arise from a linear observation equation  $\vec{Y}_t = \mathbf{H}\vec{X}_t + \vec{v}_t$  and  $\vec{v}_t \sim \mathcal{N}(\vec{0}, \mathbf{R}_t = \sigma_o^2 \mathbf{I})$ . Further, we define  $\vec{X}_t^f$  to be the model forecast, with error covariance  $\mathbf{P}_t^f$ . Note that in the experiments here, state dimension is  $N = 124200$  (see Section 5).

In the first stage, *initialization*, the large bias between model and observation resulting from starting the model from a random initial condition (Figure 6) is reduced by using a Gaussian-spline to represent forecast error-covariance. This error-covariance is not propagated across time and domain-decomposition is employed to solve the filter equations.

In the second phase, *tracking*, the ensemble Kalman filter [7] is adapted to estimate both states and their uncertainties. In this formulation, multiple model simulations (Section 5) are used to forecast an ensemble of states. These  $S \ll N$  samples form a reduced-rank square-root approximation of the forecast error-covariance, from which a Kalman-update can be computed rapidly [7]. Further, the model is never linearized, covariance is never propagated explicitly, and the estimated state is a weakly nonlinear combination of the forecast ensemble. Unlike the particle filter, no re-sampling is necessary because the states are directly updated.

For high-dimensional and nonlinear systems, a large number of Monte-Carlo simulations may be necessary to rid the forecast uncertainty from spurious correlations. However, the complexity of filtering grows as the square of the ensemble size and it is also not economical to conduct many numerical simulations. We address these issues using snapshots to produce many samples per simulation and localizing filtering to eliminate long-range correlations.

### 6.1. Initialization

The initialization phase commences after a single simulation is spun-up from a random initial condition to remove

<sup>3</sup>The state for estimation consists of the horizontal velocities and temperature. Vertical velocity is implicit, pressure is diagnostic and salinity is unrepresented.

transients. Initialization consists of four steps, executed in sequence for a few estimation cycles:

**1. Interpolation in the vertical.** An interpolation function of horizontal velocities and temperature is estimated from the forecast. Let  $\vec{v}_h^f[i, j, k]$  be the forecast horizontal velocity at grid node  $i, j, k$  in the radial, azimuthal and vertical directions respectively. Let  $\vec{v}_h^f[i, j]$  be the column-vector of forecast velocities at all  $N_z = 15$  vertical levels at  $i, j$  and let  $\vec{v}_h^{fo}[i, j]$  be the corresponding vector of horizontal velocities at the  $N_o < N_z$  observed vertical levels. Similarly construct vectors  $\vec{\theta}^f[i, j]$  and  $\vec{\theta}^{fo}[i, j]$  from the forecast temperatures. Using samples in the forecast field, *learn* the matrices  $\Lambda_v$  and  $\Lambda_\theta$  by solving  $\vec{v}_h^f[i, j] = \Lambda_v \vec{v}_h^{fo}[i, j]$  and  $\vec{\theta}^f[i, j] = \Lambda_\theta \vec{\theta}^{fo}[i, j]$ .

**2. Estimating Horizontal Velocities at observation layers.** Horizontal velocities are estimated at each observed layer ( $k_o \in \{k_1 \dots k_{N_o}\}$ ) of the fluid separately. At every location  $i, j$  on the horizontal grid ( $N_r = 23 \times N_\phi = 120$ ) of an observed layer, we estimate the horizontal velocity from forecasts and observations using a spatial context of dimensions  $N_r^l$  radially and  $N_\phi^l$  azimuthally. The estimation is written as:

$$\begin{aligned} \mathbf{C}_{ij} &= \mathbf{P}_i^f \mathbf{H}_{ij}^T (\mathbf{H}_{ij} \mathbf{P}_i^f \mathbf{H}_{ij}^T + \mathbf{R}_{ij})^{-1} \quad (5) \\ \vec{v}_h^a[i, j, k_o] &= \vec{v}_h^f[i, j, k_o] + \mathbf{C}_{ij} [\vec{v}_h^{o, ij k_o} - \mathbf{H}_{ij} \vec{v}_h^{f, ij k_o}] \end{aligned}$$

The vector  $\vec{v}_h^f[i, j, k_o]$  is the forecast horizontal velocity at location  $i, j, k_o$ ,  $\vec{v}_h^a[i, j, k_o]$  is the corresponding estimated horizontal velocity,  $\vec{v}_h^{f, ij k_o}$  is the vector of forecast horizontal velocities in a  $N_r^l \times N_\phi^l$  area centered<sup>4</sup> at  $i, j, k_o$ , and  $\vec{v}_h^{o, ij k_o}$  are available observations in the same area. The local forecast covariance  $\mathbf{P}_i^f$  (size  $2N_r^l N_\phi^l \times 2N_r^l N_\phi^l$ ) is generated using a two-dimensional Gaussian. 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  $N_r^l \times N_\phi^l$  region. The matrix  $\mathbf{R}_{ij}$  is the corresponding observational uncertainty.

**3. Estimating Temperature at observation layers.** Once the horizontal velocities  $\vec{v}_h^a[i, j, k_o]$  are estimated at each grid node of observed layers, we compute temperature  $\theta^a[i, j, k_o]$  by solving an elliptic equation from the divergence of *thermal-wind* [16], which is:

$$\frac{\partial \vec{v}_h}{\partial z} = \frac{g\alpha}{2\Omega} \hat{k} \times \nabla \theta \quad (7)$$

The temperature boundary conditions are obtained from climatological measurements (see Section 5) and  $\alpha$  is the coefficient of thermal expansion and  $\hat{k}$  is the vertical unit-vector.

**4. Estimate Full State.** The precomputed vertical interpolation models are applied to the estimated horizontal velocity and temperature. Thus we estimate  $\vec{v}_h^a[i, j] =$

<sup>4</sup>Except near annulus boundaries, where the window is off-center.

$\Lambda_{\sqrt{v}} \vec{v}_h^{ao}[i, j]$  and  $\vec{\theta}^a[i, j] = \Lambda_{\theta} \vec{\theta}^{ao}[i, j]$ , where these vectors are defined analogously to step 1 (but using the estimated fields).

The estimated fields become the new state  $\vec{X}_t = [\vec{v}_h^a; \vec{\theta}^a]$  for the next forecast. These four steps are repeated for a few estimation cycles and then switch to a flow-dependent ensemble tracking method that can both estimate states and their uncertainties, discussed next.

## 6.2. Tracking

Throughout the tracking phase, 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 and the fact that tracking is designed to produce an evolving estimation uncertainty. For tracking, we use a variation of the ensemble Kalman filter in the following way:

**Creating the Ensemble:** The two prominent sources of uncertainty are the thermal boundary condition that drives the numerical system and flow uncertainty arising from time-staggered observations and numerical integration.

To model these, we use the output of the initialization step to drive  $N_b > 1$  simulations, each utilizing a thermal boundary condition perturbed from the climatological profile (see Section 5) to have a steeper or shallower *lapse rate* at the inner core.

Additionally, to capture flow uncertainty we use snapshots [23] and save the state every few time steps in during forward integration of a simulation. If we suppose that the state  $X[\vec{x} + \vec{v}\Delta t]$  at position  $\vec{x}$  is random due to timing errors  $\Delta t$  or flow uncertainty  $\vec{v}$  between simulation and physical flow, an ensemble of samples  $X[\vec{x} + \Delta\vec{x}]$  drawn from  $\Delta\vec{x} \sim \mathcal{N}(E[\vec{v}\Delta t], \sigma_t)$  captures this uncertainty. As shown in [19] the sample covariance captures spatial correlations projected onto the flow.

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, discussed next.




Figure 8. The estimation using the ensemble Kalman filter is localized within estimation windows  $E$ , influenced by observations from overlapping spatial-context windows  $C$ .

**Localized estimation.** Horizontal velocities are estimated in an *estimation window*  $E$  of size  $N_r^e \times N_\phi^e$  at

$i_e, j_e, k_o$ , using forecasts and observations in a spatial context window  $C$  that is indexed by location  $i_c, j_c, k_o$  and of size  $N_r^c \times N_\phi^c$  (Figure 8). Estimates over an entire horizontal layer are produced by tiling it with estimation windows (no overlap). Note, however, that adjacent estimation windows share substantial spatial context, as shown in Figure 8.

Let  $\mathbf{V}^{f, i_e j_e k_o}$  be the  $2N_r^e N_\phi^e \times S$  matrix representing forecast horizontal velocities 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  $2N_r^c N_\phi^c \times S$  matrix of forecast horizontal velocities of  $S$  ensemble members coincident with the context window  $C$  at  $i_c, j_c, k_o$ . Also, let  $\tilde{\mathbf{V}}^{(\dots)}$  represent the deviation of each column from the column-mean of  $\mathbf{V}^{(\dots)}$ . Using the observations  $\mathbf{V}^{o, i_c j_c k_o}$  and forecasts in the context window selected by  $\mathbf{H}_{i_c j_c}$  to construct the transformation  $\aleph_{i_c j_c k_o}$ , we may express the analysis ensemble  $\mathbf{V}^{a, i_e j_e k_o}$  as:

$$\mathbf{V}^{a, i_e j_e k_o} = \mathbf{V}^{f, i_e j_e k_o} \aleph_{i_c j_c k_o} \quad (8)$$

$$\aleph_{i_c j_c k_o} = \left[ \mathbf{H}_{i_c j_c} \tilde{\mathbf{V}}^{f, i_c j_c k_o} \right]^T \mathbf{C}_{i_c j_c k_o}^{-1} \Delta \mathbf{V}^{i_c j_c k_o} \quad (9)$$

$$\Delta \mathbf{V}^{i_c j_c k_o} = \mathbf{V}^{o, i_c j_c k_o} - \mathbf{H}_{i_c j_c} \mathbf{V}^{f, i_c j_c k_o}$$

$$\begin{aligned} \mathbf{C}_{i_c j_c k_o} &= \mathbf{H}_{i_c j_c} \tilde{\mathbf{V}}^{f, i_c j_c k_o} \left[ \mathbf{H}_{i_c j_c} \tilde{\mathbf{V}}^{f, i_c j_c k_o} \right]^T + \\ &\quad \tilde{\mathbf{V}}^{o, i_c j_c k_o} \left[ \tilde{\mathbf{V}}^{o, i_c j_c k_o} \right]^T \\ &= \mathbf{J}_{i_c j_c k_o} \mathbf{J}_{i_c j_c k_o}^T \quad (10) \end{aligned}$$

$$\mathbf{J}_{i_c j_c k_o} = \mathbf{H}_{i_c j_c} \tilde{\mathbf{V}}^{f, i_c j_c k_o} + \tilde{\mathbf{V}}^{o, i_c j_c k_o} \quad (11)$$

To compute  $\aleph_{i_c j_c k_o}$ , the matrix  $\mathbf{C}_{i_c j_c k_o}$  can be inverted implicitly using the square-root  $\mathbf{J}_{i_c j_c k_o}$ . In practice only the estimate corresponding to the last snapshot of the current forecast of each simulation is necessary to launch the next forecast, therefore  $\aleph_{i_c j_c k_o}$  need only be  $S \times N_b$  in size, with an appropriately ordered ensemble.

## 7. Estimation Experiment

For the experiments presented here, the reference density  $\rho_0 \approx 1037 \text{ kg m}^{-3}$ , the rotation rate is  $\Omega = 1.15 \text{ rad/s}$ , the annulus width  $L = 0.15 \text{ m}$ , the mean fluid depth  $D = 0.15 \text{ m}$ , and the mean temperature difference of fluid across annulus  $\Delta T = 6 \text{ K}$  (measured separately). The viscosity is  $\nu = 10^{-6} \text{ m}^2 \text{ s}^{-1}$ , the thermal diffusivity  $\kappa = 10^{-7} \text{ m}^2 \text{ s}^{-1}$ , and the thermal expansion coefficient  $\alpha = 3 \times 10^{-4} \text{ K}^{-1}$ . Thus, the Ekman number  $E = \frac{\nu}{2\Omega D^2} = 1.9 \times 10^{-5}$ , the thermal Rossby number  $R_\theta = \frac{g\alpha\Delta TD}{\Omega^2 L^2} = 0.09$ , the Prandtl number  $\mathcal{P}_\sigma = \frac{\nu}{\kappa} = 10$ . These numbers imply that expect flows of  $2 \text{ cm s}^{-1}$  can be expected and there exists a  $10 \text{ s}$  interval within which the forecast-observe-estimate cycle must be completed, and this is achieved handily.

Experiment starts when chiller is turned on. In about 300s circulation is established (Figure 5). The model is






Figure 9. The estimated velocity field at a time  $t = 100s$  for an ensemble member at 100mm above the bottom of the tank is shown (yellow). Observations at this layer are shown in (green).




Figure 10. The RMS-error between forecast and observed velocities at all observed locations as a function of time.

started (with chiller) from a random initial condition and climatological thermal-boundary condition shown in Figure 6. It is stepped forward for 300s to remove transients and form a circulation (Figure 7). The model performs in better than realtime. On one processor of an Altix-350 we can carry out a 10-second simulation in 8-seconds, due in large part to the non-uniform discretization of the domain using variable vertical levels.

Unconstrained, model waves have the wrong phase and incorrect amplitudes and velocities can be as much as twice the observed velocities. Velocity observations are acquired by sampling particle images 125 – 250ms apart using LaVision’s DaVis software. Flow is computed in  $32 \times 32$  windows with a 16 pixel uniform pitch across the image. It takes one second to acquire and compute the flow of a single  $1K \times 1K$  image pair by distributing the computation across two 2.8GHz processors.

Filtering, stage one, is then turned on. The robotic arm scans the fluid, observing it at  $N_o = 5$  different levels and observations of the whole fluid volume are available every 5 seconds to constrain the numerical model. We set  $\mathbf{R}_{ij} = \sigma_o^2 \mathbf{I}$  and  $\sigma_o = 1.2mm/s$ , by assuming a 0.5 pixel uncertainty in optic-flow calculations per image pair. The covariance  $\mathbf{P}_i^f$  is constructed as an un-normalized Gaussian with standard deviations  $1 \times 5$  (radially by azimuthally) with extent  $N_r^l = 5 \times N_\phi^l = 10$ , and scaled by an amplitude of  $\sigma_b = \sigma_o * 2$  to account for poor model skill. Note that each  $\mathbf{C}_{ij}$  is of size  $2 \times 100$  and is constructed *a priori*<sup>5</sup>.

The observation operator  $\mathbf{H}_{ij}$  rejects grid points in the the shadow region or where velocities are more than 3cm/s (e.g., see shadow-region boundaries in Figure 5). Stage one terminates when the average error between forecast and observed velocities is less than  $1.5 * \sigma_o$ , which corresponds to about 3 estimation cycles. In stage-2, multiple simulations ( $N_b = 3$ ) start from the stage-1 model-state, with steeper, shallower or equal inner boundary-temperature vertical *lapse rates* than the climatological profile (Figure 6). During forward simulation,  $N_s = 5$  snapshots of the model-state are extracted in the last  $N_s$  model seconds. The final forecast (at  $t = 10s$ ) is used to construct the vertical interpolation functions per simulation. Observations extracted over the immediately preceding 5 seconds are used to produce an updated ensemble. We choose  $N_r^c = 11$ ,  $N_\phi^c = 21$ ,  $N_r^e = 5$  and  $N_\phi^e = 11$  for the localized estimation and the observational uncertainty is identical to stage-1. A single estimation (all four steps) with  $S = 15$ , runs on a 2.8GHz processor in under 1.6s. The final (forecast) time estimated model-states are used to initialize the next 10s forecast for each simulation.

Figure 9 shows the estimated horizontal velocities and observations after 10 estimation cycles, corresponding to the last snapshot of a simulation with climatological thermal boundary condition (Figure 6). Note that the model is well adjusted and beautifully infills velocities in the shadow region.

Figure 10 shows the evolving root mean square (RMS) error between the forecast and observation over 30 estimation cycles in a 300 second experiment<sup>6</sup>. The model velocities come close to the observations within the inherent uncertainty ( $\sigma_o = 1.2mm/s$ ) to which observations are represented. Indeed both the amplitudes and phase are in good agreement as can be seen in Figure 9. Repeated experimentation with different regimes shows that it takes approximately 10 rotation periods or six estimation cycles before the model adjusts itself to be consistent with the observations. After 20 estimation cycles, we turn it off. As expected the error between forecasts and observations grows, and saturates in around 10 cycles.

<sup>5</sup>A large number of matrices  $\mathbf{C}_{ij}$  are identical, saving storage costs.

<sup>6</sup>Because the truth is unknown, we use the likelihood as the diagnostic.

## 8. Conclusions

The coupled realtime system for studying rotating fluid-flows is robust and scalable. Visual flow calculations, numerical simulations and localized filtering are highly parallelizable, and snapshots reduce the number of numerical simulations per sample.

The hybrid filtering algorithm uses a flow-free prior that allows subsequent use of reduced-rank flow-dependent covariances. Localization prevents spurious long-range correlations. Mixing snapshots and thermal boundary-condition perturbations is also very useful. In fast unsteady flows, snapshots capture the dominating flow-uncertainty. In nearly steady flows, the boundary-condition uncertainty dominates.

There is also room for improvement. Multi-resolution and adaptive approaches can improve performance. Whole-field laser-induced fluorescence can be used to constrain with temperature measurements. A new periscope with a rotating mirror and stationary paraboloid will improve the scan speed many fold and allow denser observation.

Even without these improvements, our observatory is remarkably versatile. The system is easy to replicate, data sets are readily available and it is used routinely in our laboratory. For rotating fluids, such a system has not been achieved before.

## References

- [1] A. Arakawa and V. Lamb. Computational design of the basic dynamical processes of the ucla general circulation model. *Methods in Computational Physics, Academic Press*, 17:174–267, 1977.
- [2] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, Feb. 2002.
- [3] A. E. Bryson and Y. C. Ho. *Optimal Control*. Blaisdell, 1969.
- [4] P. A. Davis. Laboratory experiments in geophysical fluid dynamics. In P. H. Roberts and A. M. Soward, editors, *Rotating Fluids in Geophysics*, pages 523–546. Academic Press, 1977.
- [5] P. A. Davis and G. Walin. Experiments with a rotating differentially-heated fluid annulus. *Tellus*, 27(6):574–596, 1975.
- [6] L. DeRose, K. Gallivan, and E. Gallopoulos. Experiments with an ocean circulation model on CEDAR. In *6th ACM International Conference on Supercomputing*, pages 397–408, Washington, D.C., 1992.
- [7] G. Evensen. The ensemble kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, 53:342–367, 2003.
- [8] J. E. Geisler, E. J. Pitcher, and R. C. Malone. Rotating-fluid experiments with an atmospheric general circulation model. *Journal of Geophysical Research*, 88(C14):9706–9716, 1983.
- [9] A. Gelb. *Applied Optimal Estimation*. MIT Press, 1974.
- [10] R. Hide. An experimental study of thermal convection in a rotating liquid. *Phil. Trans. Roy. Soc.*, A250:441–478, 1958.
- [11] E. N. Lorenz. Deterministic nonperiodic flow. *J. Atmos. Sci.*, 20:130–141, 1963.
- [12] J. Marshall, A. Adcroft, C. Hill, L. Perelman, and C. Heisey. A finite-volume, incompressible navier stokes model for studies of the ocean on parallel computers. *J. Geophysical Res*, 102(C3):5753–5766, 1997.
- [13] J. Marshall, C. Hill, L. Perelman, and A. Adcroft. Hydrostatic, quasi-hydrostatic and nonhydrostatic ocean modeling. *Journal of Geophysical Research*, 102(C3):5733–5752, 1997.
- [14] O. Morita and M. Uryu. Geostrophic turbulence in a rotating annulus of fluid. *Journal of Atmospheric Sciences*, 64(15), 1989.
- [15] L. Nerger, W. Hiller, and J. Schortner. A comparison of error subspace kalman filters. *Tellus-A*, 57(5):715–735, 2005.
- [16] J. Pedlosky. *Geophysical fluid dynamics*. Springer-Verlag, 1987.
- [17] H. E. Rauch. Solutions to the linear smoothing problem. *IEEE Transactions on Automatic Control*, 8:371–372, 1963.
- [18] S. Ravela and D. McLaughlin. Fast ensemble smoothing. *Ocean Dynamics*, 57(2):123–134, 2007.
- [19] S. Ravela, A. Torralba, and W. T. Freeman. An ensemble prior of image structure for crossmodal inference. In *International Conference on Computer Vision*, volume 1, pages 871–876, 2005.
- [20] P. L. Read. A combined laboratory and numerical study of heat transport by baroclinic eddies and axisymmetric flows. *J. Fluid Mech.*, 489, 2003.
- [21] P. L. Read, M. J. Bell, D. W. Johnson, and R. M. Small. Quasi-periodic and chaotic flow regimes in a thermally driven, rotating fluid annulus. *J. Fluid Mech.*, 238:599–632, 1992.
- [22] P. L. Read, N. P. J. Thomas, and S. H. Risch. An evaluation of eulerian and semi-lagrangian advection schemes in simulations of rotating, stratified flows in the laboratory. part i: Axisymmetric flow. *Monthly Weather Review*, 128:2835–2852, 2000.
- [23] L. Sirovich. Turbulence and the dynamics of coherent structures, part 1: Coherent structures. *Quarterly of Applied Mathematics*, 45(3):561–571, 1987.
- [24] E. Sudderth, M. Mandel, W. Freeman, and A. Willsky. Visual hand tracking using nonparametric belief propagation. In *Workshop on Generative Model-based Vision*, 2004.
- [25] P. K. Sweby. High resolution schemes using flux-limiters for hyperbolic conservation laws. *SIAM Journal of Numerical Analysis*, 21:995–1011, 1984.
- [26] C. Wunsch. *The Ocean Circulation Inverse Problem*. Cambridge University Press, 1996.
- [27] R. Young and P. Read. Breeding vectors in the rotating annulus as a measure of intrinsic predictability. In *Royal Met. Soc. Annual Student Conference*, 2006.