A Realtime Observatory for Laboratory Simulation of Planetary Flows*

Sai Ravela John Marshall Chris Hill Andrew Wong
Scott Stransky

Massachusetts Institute of Technology
54-1624, MIT, Cambridge, MA 02139
ravela@mit.edu

Abstract

Motivated by the large-scale circulation of the atmosphere and ocean, we develop a system that uses observations from a laboratory analog to constrain, in real time, a numerical simulation of the laboratory flow. This 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 of the physical simulation. Observations are used to constrain the model-state of the MIT General Circulation Model in a probabilistic, ensemble-based 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 such observatory for laboratory analogs of planetary circulation that functions in real time.

1. Introduction

Laboratory experiments have been extensively used to understand the properties of fluids [2, 23, 9, 14, 8, 19]. Of particular focus here is the large-scale circulation of the atmosphere, for which a well-known laboratory analog is a thermally-driven rotating flow [8, 19, 10, 3, 2, 14]. In this experiment, a rotating annulus with a cold center (core) and warm periphery (exterior) develops a circulation that has dynamical similarity to the mid-latitude circulation in the atmosphere, as shown in Figure 1. It is

*This research in this paper is funded by NSF Grant CNS 0540248.

a robust experiment, easily conducted in the laboratory. It has been used to study a variety of physical phenomena including geostrophic turbulence [14], convection [8], baroclinic instability [9, 23, 18], and chaos [19, 10], and as a test-bed for evaluating the utility of numerical models [20, 6].

We present an automated real-time observatory for this laboratory experiment. By observatory, we mean a coupled physical-numerical system with the following components: sensors to take measurements of the evolving physical system, a numerical model trying to forecast the system, and inference algorithms that constrain the model with observations to produce an evolving state that is closer to the laboratory flow than either observations and model alone.

We contend that a number of exciting possibilities open up once the observatory operates in real-time. Properties of the fluid that cannot easily be observed (surface height, pressure fields, vertical velocities, radial heat transport etc.) can be studied using the model. Studies in tracer transport can be conducted using the observatory in real-time. New algorithms that address prediction and predictability issues of state and parameter estimation, model error and targeting can be rapidly validated. Across disciplines, the platform provides application to new distributed computing, visualization and augmented reality applications. Whilst it is not possible, in one paper, to explore each and every application, a large number of potential applications will require the coupled numerical-physical system to track the fluid's state in real-time. Therefore, we focus here on the design of the observatory, including a procedure to estimate model-states in real-time.

Tracking, when formulated as a state and parameter estimation problem, is known to be fundamentally challenging in weather forecasting [11]. In the context of ocean state estimation or weather forecasting, predictions are typically made using general circulation models (GCMs), which implement the discretized governing equations. GCMs typically have uncertain parameters and crude parameterizations, 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 PDF of the

evolving model state's uncertainty is unlikely to retain the true state within it. A way forward is to constrain the model with observations of the physical system [24].

We posit that studying the estimation problem in the laboratory is convenient and useful. Firstly, repeatable experiments with real data can be performed using far simpler logistics than the operational setting. Secondly, the following key challenges demanded by the large-scale problem must also be addressed in the laboratory setting:

1. Nonlinearity — the laboratory analog is nonlinear and the numerical model is the same used in planetary simulations

2. Dimensionality — the size of the state of the numerical model is of the same order as planetary simulations

3. Uncertainty — the initial conditions are unknown, and the model is imperfect relative to the physical system

4. Realtime — forecasts must be produced in better than realtime.

These are all interesting problems in their own right and solutions found in a laboratory setting can accelerate acceptance of new methods in operational settings and could be useful in many other coupled numerical-physical systems.

Before going on it should be noted that the rotating annulus experiment has already been used to explore the utility of numerical models. Read et al. [20] use the annulus to study how well numerical transport schemes compare to real observations and report that Eulerian schemes, such as used in this paper, have skill. In more recent work Read [18] combines numerical studies with laboratory experiments in the study of heat transport and effort has been afoot to study prediction and predictability problems using the laboratory setting [25, 26, 16, 17]. To the best of our knowledge, however, this is the first coupled observation/numerical system of a laboratory experiment to operate in realtime [17].

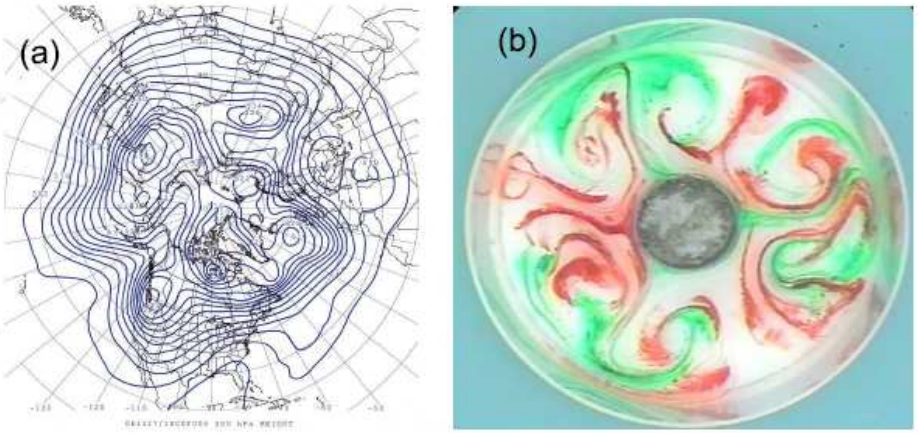

Figure 1. Image (a) shows the 500hPa heights for 11/27/06:1800Z over the northern hemisphere centered at the north pole. Winds flow along the pressure contours. Image (b) shows a tracer (dye) in a laboratory analog. The tank is spinning and the camera is in the rotating frame. Tracer droplets initially 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 akin to one-another. We study the state-estimation problem for planetary flows using the laboratory analog.

The observatory operates by continually taking measurements of the physical system and uses a probabilistic, ensemble-based estimation method to constrain the model-states of a numerical model. Our system does this using off-the-shelf components for the laboratory experiments, commercially available software to extract observations, the publicly available MIT-GCM, and a new hybrid filter that combines deterministic and probabilistic filtering. The probabilistic component is related to the Local Ensemble Kalman Filter [15], which is derived from the ensemble Kalman filter [4, 5].

Our coupled system can operate in a variety of dynamical regimes. For the experiments presented here, an assimilation cycle (forecast-observe-assimilate) must be completed within roughly 10 seconds, a typical rotation period. Our system accomplishes this using domain decomposition, spectral-reduction, distributed computation, and a new way to generate and select ensembles. It is now in routine use and data-sets are readily available to other researchers[1]

[1] Videos are enclosed as supplementary information with this paper. They depict the operational use of our observatory.

2. The Observatory

The observatory, illustrated in Figure 2, has a physical and computational component. The physical component consists of a perspex annulus of inner radius $8cm$ and outer radius of $23cm$, filled with $15cm$ 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' pliolite particles sg $1.03g/cc$) are homogenized in saline water of equal density and respond to incident laser illumination. They appear as a plane of textured dots in the $12-$bit quantized, $1K \times 1K$ images (see Figure 4) of an Imperx camera. These images are transferred out of the rotating frame using a Hitachi fiber-optic rotary joint (FORJ or slip-ring).

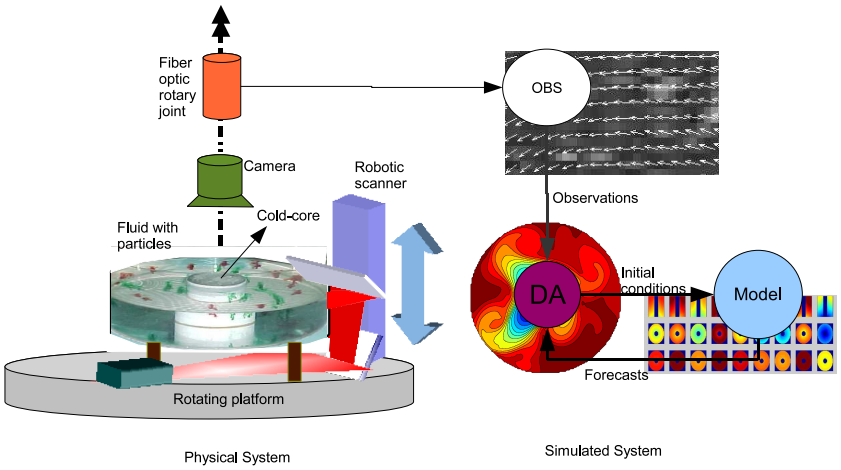

Figure 2. The laboratory observatory consists of a physical system: a rotating table on which a tank, camera and control system for illumination are mounted. The computational part consists of a measurement system for velocimetry, a numerical model, and an assimilation system, as described more fully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

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. To see this, consider that particles can move at up to $2cm/s$. The camera scale factor is approximately $0.5mm/pixel$ and it is positioned $50cm$ away from the annulus. At a sampling rate of $1/4s$, the camera must shake by less than 0.1^o to have less than 10% motion noise. Therefore, one must be very careful to eliminate vibrations. We center the rig and hold the FORJ-assembly using four bungee chords, which have the appropriate stiffness (see

Figure 3) to damp vibrations and moments.

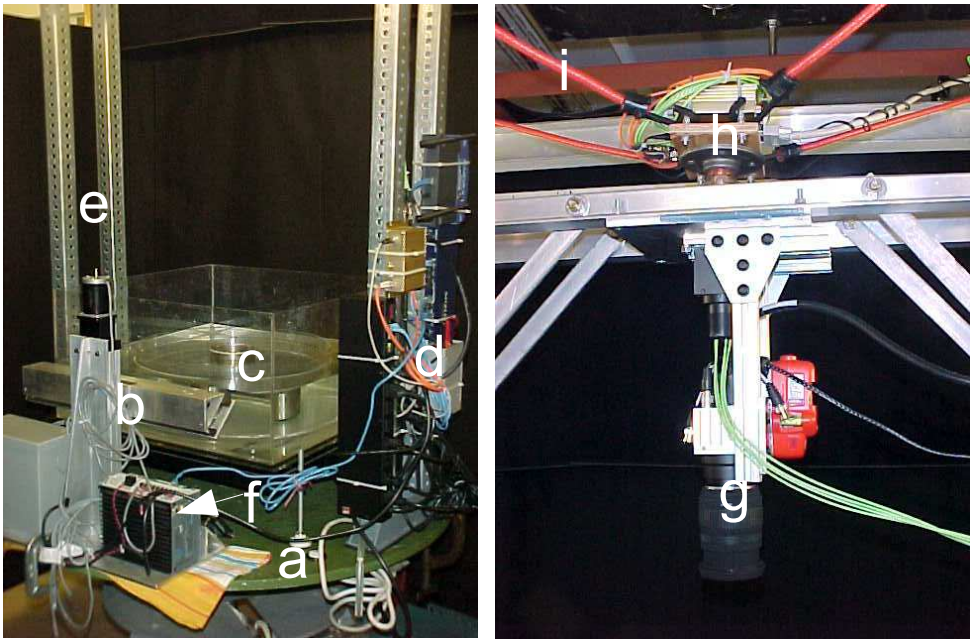

Figure 3. The apparatus consists of (a) the rotating platform, (b) the motorized mirror, (c) the tank, (d) electronics, (e) a rig on which a camera is mounted, (g). Laser light comes from direction (f) and bounces off two mirrors before entering the tank. The fiber optic rotary joint (FORJ) (h) allows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

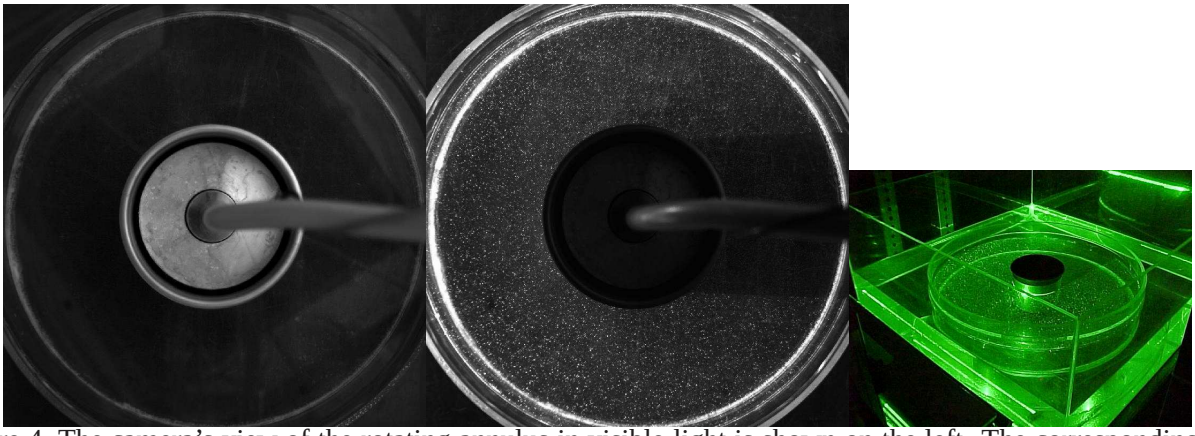

Figure 4. The camera's view of the rotating annulus in visible light is shown on the left. The corresponding view in laser light depicts the embedded particles and is shown in the middle image. A picture of the laser illuminating the tank is shown on the right. Notice the shadow due to the chiller in the middle. The square tank is used to prevent 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). Flow vectors are passed to an assimilation 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.

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 $3s$ and $12s$). After twenty minutes or so the fluid comes into solid body rotation. The inner core is then cooled using a chiller (see Figure 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 Figure 1) and baroclinic instability sets in.

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

The ultra-small pliolite particles move with the flow. We observe the horizontal component and compute optical flow from image pairs acquired $125 - 250ms$ apart using LaVision's DaVis software. Flow is computed in 32×32 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. An example is shown in Figure 5.

Observations are gathered over several levels on a repeating cycle. The mirror moves to a preprogrammed level, the system captures images, flow is computed, and the mirror moves to the next preprogrammed level and so on, scanning the fluid volume in levels. We typically observe the fluid at five different levels and so observations of the whole fluid are available every 5 seconds and used to constrain the numerical model of the laboratory experiment.

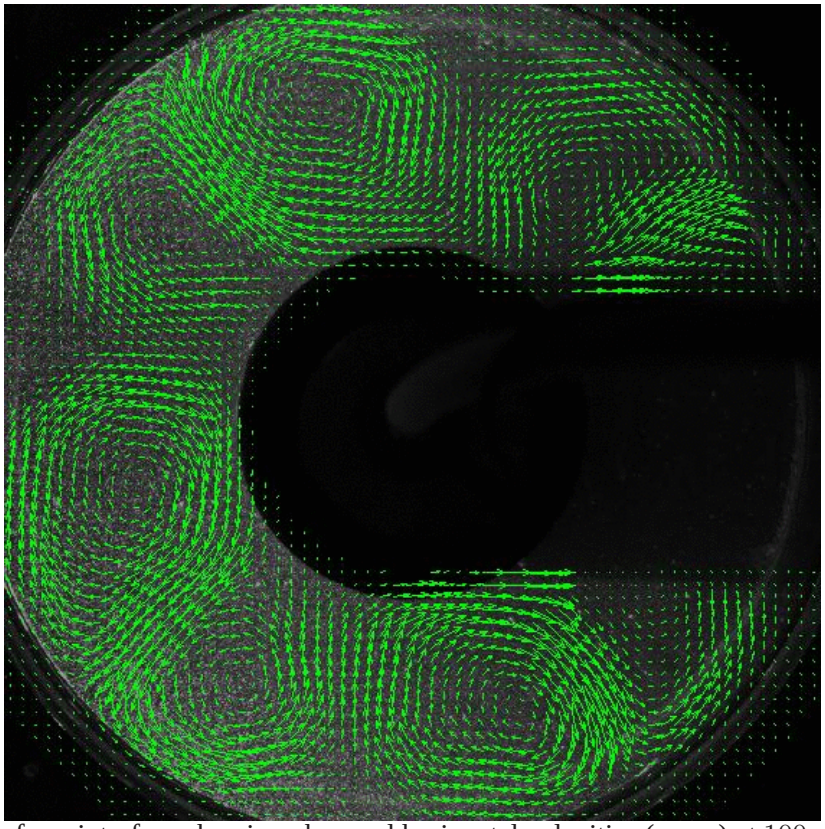

Figure 5. A snapshot of our interface showing observed horizontal velocities (green) at $100\ mm$ above the bottom of the tank after circulation has formed. Maximum flow speeds are of the order of $2cms^{-1}$.

3. Numerical Model

We use the MIT General Circulation Model developed by Marshall et al. [13, 12] to numerically simulate the laboratory experiment. The MIT-GCM is freely available software and can be configured for a 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:

$$\frac{\partial \vec{v}_h}{\partial t} = G_{v_h} - \frac{1}{\rho_0}\nabla_h p \quad \text{horizontal momentum} \tag{1}$$

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

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

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

Here, the three-dimensional velocity is denoted by $\vec{v} = [\vec{v}_h; w]$ where $\vec{v}_h$ 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. [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 ($0.65cm/$bin) and 120 bins in azimuth (3^o bins). The vertical coordinate is discretized using 15 levels non-uniformly distributed over the $15cm$ depth of the fluid, as shown in Figure 6(b). The MIT-GCM discretizes variables on an Arakawa C-grid [1]. Momentum

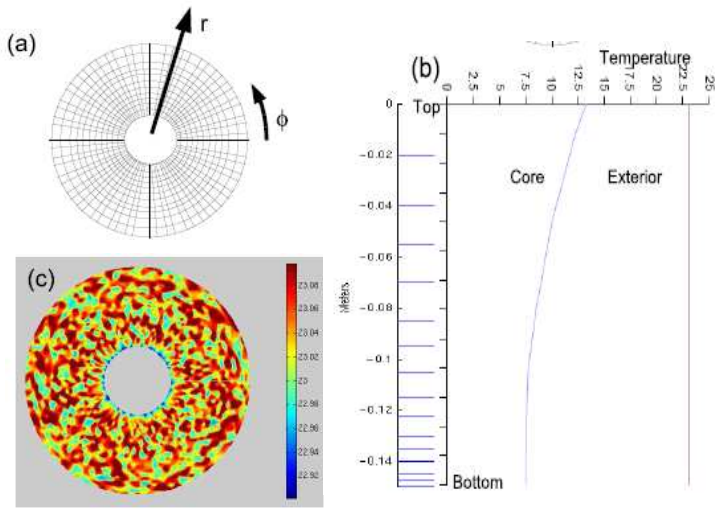

Figure 6. (a) The computational domain is represented in cylindrical coordinates. (b) Depth is discretized with variable resolution to enhance resolution near the bottom-boundary. The lateral boundary conditions on temperature are obtained by interpolating sparse temperature measurements on the boundary. 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.

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 [22]. The treatment of vertical transport is implicit. A 2-d 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 Figure 6(c). The model performs in better than realtime. On one processor of an Altix350 we can carry out a 10-second simulation in 8-seconds. 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 Figure 7 the model horizontal currents are overlaid on the observed velocities after suitably registering the domain geometry to the physical tank. We observe that 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

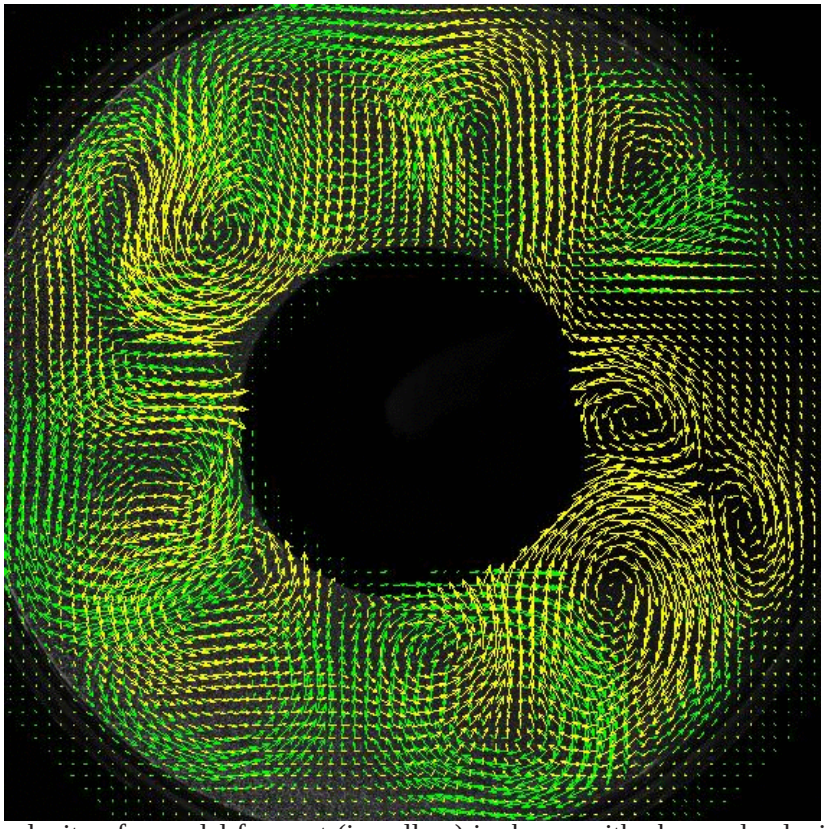

Figure 7. The planar velocity of a model forecast (in yellow) is shown with observed velocity (green) at a height of $100mm$ from the bottom of the tank at the beginning of an assimilation experiment. Maximum flow speeds are of order $2\ cms^{-1}$.

flow details are different. We therefore now go on to describe how we use observations to constrain the evolving state of the model.

4. State Estimation

At a rotation period of six seconds, fluid parcels can traverse the annulus at up to $2cm/s$, leaving approximately 10 seconds to complete an assimilation cycle. Since it typically takes 8 real-seconds for a $10s$ numerical simulation, and 5 seconds to gather observations (in parallel), there are 2 seconds left for communication and computational activities, before which the next forecast must be initiated. Using sequential filtering, we now describe an estimation method that produces a state estimate within these two seconds.

Let $\vec{X}_t = [\vec{v}_h; \vec{\theta}]$ be the state[2] at a discrete time t, and measurements $\vec{Y}_t$ be assumed to arise from from a linear observation equation $\vec{Y}_t = \mathbf{H}\vec{X}_t + \vec{\nu}_t$, where the observational noise is normally distributed with zero mean and diagonal covariance $\mathbf{R}_t$, that is $\vec{\nu}_t \sim \mathcal{N}(0, \mathbf{R}_t)$. Further, let $\vec{X}_t^f$ be the model forecast, with error covariance $\mathbf{P}_t^f$. Now the well-known update equation for analysis state $\vec{X}_t^a$ can be written as:

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

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

As shown in [7] the Kalman and extended-Kalman filter are given by Equations 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.

Another way is to use a reduced-rank spectral approximation of the forecast uncertainty. In the Ensemble Kalman Filter [4] formulation, 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 $\mathbf{V}^f = [\vec{X}_1^f \ldots \vec{X}_S^f]$, where the columns of $\mathbf{V}^f$ are the S samples of the ensemble of horizontal velocities at an observed layer. Thus, if we let $\mathbf{V}^o$ represent a S-column matrix of perturbed observations, obtained by perturbing an observation $\vec{Y}$ with noise $\vec{\nu}$, $\tilde{\mathbf{V}}^f$ be the deviation from mean[3] $\bar{\mathbf{V}}^f$ of $\mathbf{V}^f$ the update equation can be written as:

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

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

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

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

$$= \mathbf{V}^f\aleph \quad (9)$$

The posterior (or analysis) distribution is represented by mean $\bar{\mathbf{V}}^a$ and covariance $\mathbf{P}^a = \frac{1}{S-1}\tilde{\mathbf{V}}^{\mathbf{a}}\tilde{\mathbf{V}}^{\mathbf{a}^T}$. 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, 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. Once initialized, we switch to the second phase, called tracking. An ensemble method is used for tracking, during which both states and their uncertainties are estimated. In both phases, domain-decomposition is used. In initialization for addressing dimensionality and in tracking for removing long-range correlations. Thus, localized versions of Equations 6 and 9 will be implemented. 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 Figure 6). 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:

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 corresponding to horizontal grid location i,j and let $\vec{v}_h^{fo}[i,j]$ be the corresponding vector of horizontal velocities at the $N_o = 5$ 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, estimate the matrices $\mathbf{\Lambda_v}$ and $\mathbf{\Lambda}_\theta$ by solving equations of the form $\vec{v}_h^f[i,j] = \mathbf{\Lambda_v}\vec{v}_h^{fo}[i,j]$ and $\vec{\theta}_h^f[i,j] = \mathbf{\Lambda}_\theta\vec{\theta}_h^{fo}[i,j]$.

2. Estimating Horizontal Velocities at observation layers. At each observed layer ($k_o \in \{k_1 \ldots k_5\}$ 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 ($N_r = 23 \times N_\phi = 120$) 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_ϕ^l azimuthally. The estimation is written as:

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

$$\vec{v}_h^a[i,j,k_o] = \vec{v}_h^f[i,j,k_o] + \mathbf{C}_{ij}[\vec{v}_h^{o,ijk_o} - \mathbf{H}_{ij}\vec{v}_h^{f,ijk_o}] \quad (11)$$

Here, $\vec{v}_h^{f,ijk_o}$ is the vector of forecast horizontal velocities in a $N_r^l \times N_\phi^l$ area centered[4] at grid node

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

i, j, k_o, and $\vec{v}^{o,ijk_o}$ 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. We typically choose $N_r^l = 5$ and $N_\phi^l = 10$, therefore each $\mathbf{C}_{ij}$ is of size 2×100 and is constructed *a priori*[5]. The vector $\vec{v}_h^f[i, j, k_o]$ is the forecast horizontal velocity at location i, j, k_o and $\vec{v}_h^a[i, j, k_o]$ is the corresponding estimated (sometimes called assimilated or analysis) horizontal velocity.

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 *thermal-wind* equation at each observed layer. The temperature boundary conditions are obtained from climatological measurements, as discussed in Section 3.

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

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

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

following way:

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 Section 3). Additionally, motivated by the method of snapshots [21], 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, discussed next.

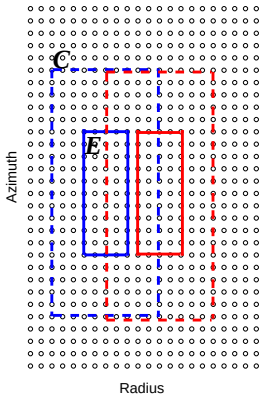

Figure 8. The estimation using the ensemble Kalman filter is localized within estimation windows E, influenced by observations from overlapping spatial-context windows C.

Localized estimation. 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 Figure 8. Estimates of horizontal velocities are produced for nodes 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 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. 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:

$$\mathbf{V}^{a,i_e j_e k_o} \quad = \quad \mathbf{V}^{f,i_e j_e k_o} \aleph_{i_c j_c k_o} \tag{12}$$

In practice 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.

A single assimilation (all four steps) with $S = 15$, runs on a $2.8GHz$ processor in under $1.6s$. Note that our approach is related to LEKF [15], with substantial differences in how estimation and context windows are designed and used.

5. Experiments

For the experiments presented here, the reference density $\rho_0 \approx 1037kgm^{-3}$, the rotation rate is $\Omega = 1.15rad/s$, the annulus width $L = 0.15m$, the mean fluid depth $D = 0.15m$, and the mean temperature difference of fluid across annulus $\Delta T = 6K$ (measured separately). The viscosity is $\nu = 10^-6m^2s^{-1}$, the thermal diffusivity $\kappa = 10^{-7}m^2s^{-1}$, 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 $\mathcal{P}_\sigma = \frac{\nu}{\kappa} = 10$.

We cool the core after the fluid attains solid body rotation. A circulation is established in about $300s$, and an example of a well-formed circulation is shown in Figure 5 at a layer $100mm$ high from the bottom

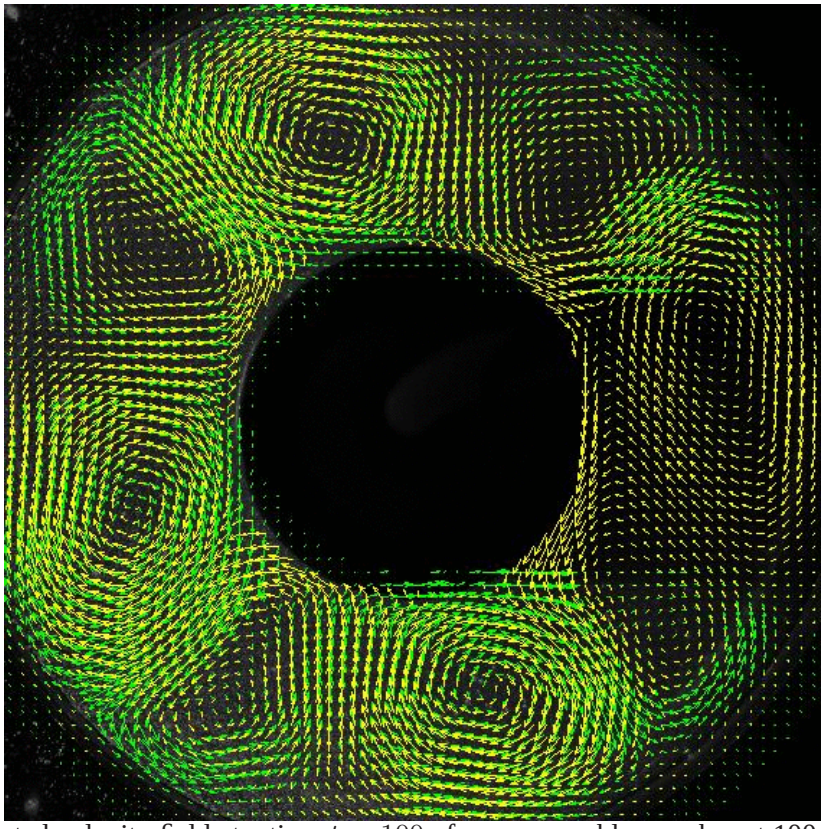

Figure 9. The assimilated velocity field at a time $t = 100s$ for an ensemble member at 100mm above the bottom of the tank is shown (yellow). Observations at this layer are shown in (green).

of the tank.

The MIT-GCM is started from a random initial condition with a climatological thermal-boundary condition shown in Figure 6. Using the parameters described in Section 3, the model is integrated forward for $300s$ to remove transients and establish a circulation, albeit unconstrained with observations. The horizontal velocity field at $100mm$ from the bottom of the tank is shown in Figure 7 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 observe that model velocities can be as much as twice that of the observed velocities.

We then turn on the assimilation component. The local observational uncertainty $\mathbf{R}_{ij} = \sigma_o^2\mathbf{I}$. A 0.5 pixel uncertainty in PIV calculations per image pair is a reasonable assumption and translates to velocity uncertainty of approximately $\sigma_o = 1.2mm/s$. This uncertainty can arise from representativeness error

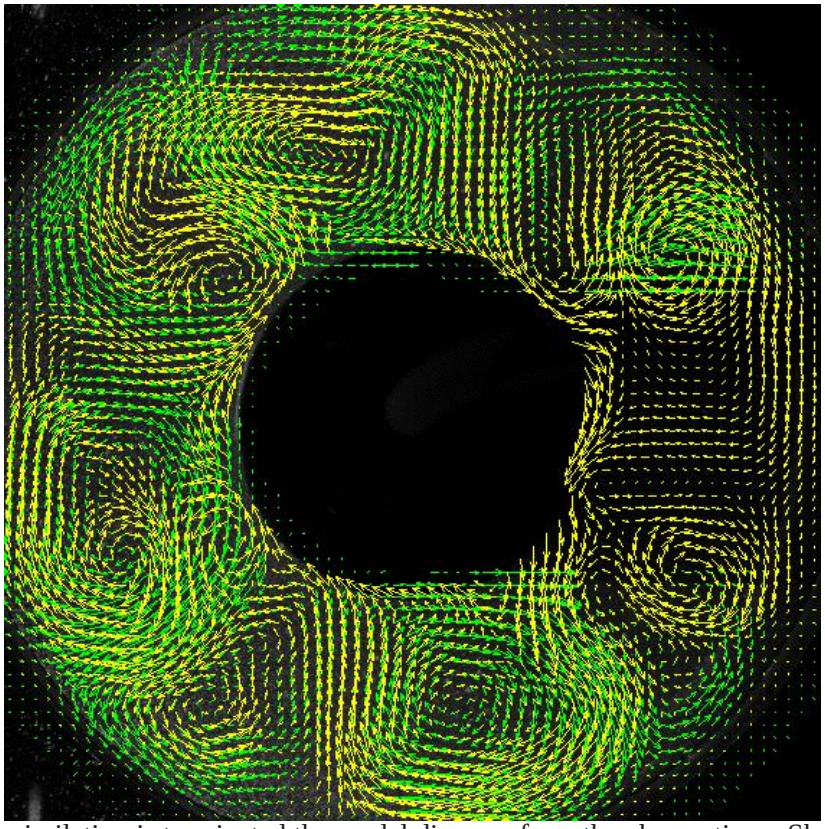

Figure 10. Once the assimilation is terminated the model diverges from the observations. Shown here is the model velocity for an ensemble member at 100mm above the bottom of the tank (yellow) and corresponding observations (green) at $t = 300s$.

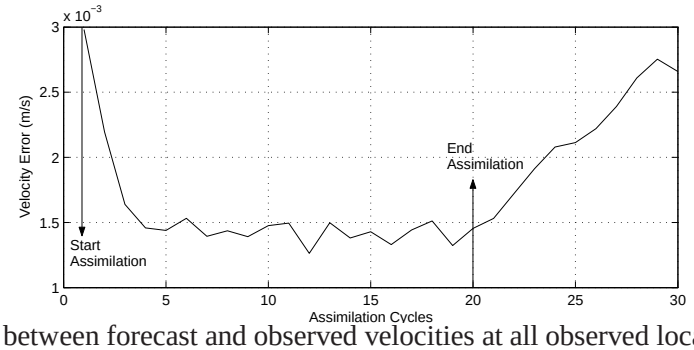

Figure 11. The RMS-error between forecast and observed velocities at all observed locations as a function of time.

due to size of the PIV window, change of focus with depth, noise and other factors.

During initialization, the covariance $\mathbf{P}_i^f$ is constructed as an un-normalized two-dimensional Gaussian with standard deviation of 1 (radially) and 2 azimuthally, with extent of 5 grid nodes (radially) and 10 grid nodes (azimuthally). The Gaussian is scaled by an amplitude of $\sigma_b = \sigma_o * 2$, to account for the observation 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 $3cm/s$. Doing so excludes impulse noise, seen for example at the edges of the shadow region in Figure 5.

With these parameters the deterministic assimilation scheme is run till the root mean square error between observed and forecast horizontal velocities over all observed locations is less than $1.5 * \sigma_o$. This corresponds to approximately 3 assimilation cycles.

After the initialization, the system switches to an ensemble scheme. We run three different simulations ($N_b = 3$), each of which start from the model-state estimated during initialization, but with temperature boundary conditions perturbed to have steeper or shallower *lapse rates* than the climatological profile. Each simulation runs on a separate processor of the Altix350, and integrates the model $10s$ forward in approximately $8s$ of clock-time. At every second of the last five seconds of this integration a snapshot of the model-state (horizontal velocity and temperature) is extracted from each simulation($N_s = 5$). Thus, at the end of the 10 second period, an ensemble with 15 members becomes available. The final forecast (at $t = 10s$) is used to estimate the interpolation functions in the vertical separately for each simulation. The observations extracted in the immediately preceding 5 seconds are used in the ensemble assimilation scheme discussed in Section 4. The observational uncertainty is identical to the deterministic case and we choose $N_r^c = 11$, $N_\phi^c = 21$, $N_r^e = 5$ and $N_\phi^e = 11$. Figure 9 shows the estimated horizontal velocities and observations after 10 assimilation cycles at a height of $100mm$ from the bottom of the tank. The estimate depicted here corresponds to the last snapshot of the simulation with a climatological thermal boundary condition profile in Figure 6. The final time estimated model-states are used to re-initialize it for the next $10s$ forecast.

Figure 11 shows the evolving root mean square (RMS) error between the forecast and observation 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 observations nearing 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. After 20 assimilation cycles we turn the assimilation off and simply compute the error between forecast velocities and observations. As expected this error grows, and saturates in around 10 cycles. Figure 10 shows the model velocities and observations at $t = 300s$ for the ensemble member corresponding to Figure 9, at 100mm above the bottom of the tank. The model has once again departed from the system trajectory. Similarly configured experiments suggest that it takes approximately 10 rotation periods or six assimilation cycles before the model adjusts itself to be consistent with the observations.

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 [24] 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 require 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 can be 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). Thus the present system can 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. Observations are presently gathered in 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 in its current application. Moreover the components used are largely off-the-shelf and relatively inexpensive. 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. Its utility extends to many problems in prediction and predictability, oceanography and meteorology. Our hope is that other researchers will be able to make use of this system, the methods used or the data-sets generated.

Acknowledgment

The authors thank Ryan Abernathy for his help in the development of the hardware platform.

References

[1] Arakawa, A., Lamb, V.: Computational design of the basic dynamical processes of the ucla general circulation model. Methods in Computational Physics, Academic Press **17**, 174–267 (1977)

[2] Davis, P.A.: Laboratory experiments in geophysical fluid dynamics. In: P.H. Roberts, A.M. Soward (eds.) Rotating Fluids in Geophysics, pp. 523–546. Academic Press (1977)

[3] Davis, P.A., Walin, G.: Experiments with a rotating differentially-heated fluid annulus. Tellus **27**(6), 574–596 (1975)

[4] Evensen, G.: The ensemble kalman filter: Theoretical formulation and practical implementation. Ocean Dynamics **53**, 342–367 (2003)

[5] Evensen, G.: Samplingstrategies and square root analysis schemes for the enkf. Ocean Dynamics **54**, 593–560 (2004)

[6] Geisler, J.E., Pitcher, E.J., Malone, R.C.: Rotating-fluid experiments with an atmospheric general circulation model. Journal of Geophysical Research **88**(C14), 9706–9716 (1983)

[7] Gelb, A.: Applied Optimal Estimation. MIT Press (1974)

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

[9] von Larcher, T., Egbers, C.: Experiments on transitions of baroclinic waves in a differentially heated rotating annulus. Nonlinear Processes in Geophysics **12**, 1044–1041 (2005)

[10] Lee, C.: Basic instability and transition to chaos in a rapidly rotating annulus on a beta-plane. Ph.D. thesis, University of California, Berkeley (1993)

[11] Lorenz, E.N.: Deterministic nonperiodic flow. J. Atmos. Sci. **20**, 130–141 (1963)

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

[13] Marshall, J., Hill, C., Perelman, L., Adcroft, A.: Hydrostatic, quasi-hydrostatic and nonhydrostatic ocean modeling. Journal of Geophysical Research **102**(C3), 5733–5752 (1997)

[14] Morita, O., Uryu, M.: Geostrophic turbulence in a rotating annulus of fluid. Journal of Atmospheric Sciences **64**(15) (1989)

[15] Ott, E., Hunt, B.R., Szunyogh, I., Zimin, A., Kstelich, E., Corazza, M., Kalnay, E., Patil, D.J., Yorke, J.A.: A local ensemble kalman filter for atmospheric data assimilation. Tech. Rep. arXiv:physics/0203058 v4, University of Maryland (2003)

[16] Ravela, S., Hansen, J., Hill, C., Marshall, J., Hill, H.: On ensemble-based multiscale assimilation. In: European Geophysical Union Annual Congress (2003)

[17] Ravela, S., Marshall, J., Hill, C., Wong, A., Stransky, S.: A real-time observatory for laboratory simulation of planetary circulation. In: Lecture Notes in Computer Science, vol. 4487, pp. 1155–1162 (2007)

[18] Read, P.L.: A combined laboratory and numerical study of heat transport by baroclinic eddies and axisymmetric flows. J. Fluid Mech. **489** (2003)

[19] Read, P.L., Bell, M.J., Johnson, D.W., Small, R.M.: Quasi-periodic and chaotic flow regimes in a thermally driven, rotating fluid annulus. J. Fluid Mech. **238**, 599–632 (1992)

[20] Read, P.L., Thomas, N.P.J., Risch, S.H.: 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)

[21] Sirovich, L.: Turbulence and the dynamics of coherent structures, part 1: Cohrerent structures. Quarterly of Applied Mathematics **45**(3), 561–571 (1987)

[22] Sweby, P.K.: High resolution schemes using flux-limiters for hyperbolic conservation laws. SIAM Journal of Numerical Analysis **21**, 995–1011 (1984)

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

[24] Wunsch, C.: The Ocean Circulation Inverse Problem. Cambridge University Press (1996)

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

[26] Young, R., Read, P.L.: Intrinsic predictability measures of baroclinic chaos and quasi-periodic flow in the rotating annulus. In: Geophysical Research Abstracts. EGU General Assembly (2007)