

PREDICTING CLIMATE REGIME CHANGE IN CHAOTIC CONVECTION

An Undergraduate Honors Thesis Presented

by

Kameron Decker Harris

Submitted to

Department of Mathematics and Statistics

Department of Physics

Honors College

Advised by

Christopher M. Danforth, PhD

In Partial Fulfillment of the Requirements

for College Honors

University of Vermont

April, 2009

Abstract

A toy climate analogous to the famous Lorenz system [1] is derived and compared to computational fluid dynamics (CFD) simulations in order to test new modeling techniques. In particular, methods of data assimilation and ensemble forecasting are used to predict regime changes and residencies in this toy climate. A climate “truth” is created using a finite-element simulation of a thermal convection loop, a physical apparatus designed to be the simplest model of convection in the Earth’s atmosphere. Forecasts of the climate are made using low-dimensional Lorenz-like models and synchronized to noisy observations of the truth using various Kalman filters. Forecasting of regime changes has been successfully demonstrated when the same model is used to create both the observations and the forecast [2], but never for realistic chaotic convection.

Acknowledgments

I must give credit to the multitude of excellent teachers that came before UVM, including but not limited to: Ms. Lewis, who taught us extra math in grade school; Marge Drexler and Frankie Dunleavy, for making Roman history interesting while beating into our heads conjugations, declensions, and nasty things like the gerundive of obligation; math teacher Lisa O'Connor; "Super Fizix" leader Carl Cierniewski; and Judith Vaill, for not letting me pass off second-rate papers and for sometimes being just a little crazy.

At UVM, I have had the pleasure of taking classes with a number of good teachers, including but not limited to: Ken Golden, Adam Wolfe, Peter Dodds, and Dennis Clougherty. The Honors College, especially Patty Redmond, Abu Rizvi, and Lisa Schnell, was a source of guidance and helpful in resolving a number of small issues.

Integral to this project were Floyd Vilmont, who is building the experiment, Co-PI Darren Hitt, and El Hassan Ridouane, who configured the FLUENT model, made the pretty pictures, and was a partner in many interesting conversations. Chris Danforth, my advisor for this beast, deserves special recognition for his accessibility, advocacy, and enthusiasm.

Mad River Glen, and our fantastic team of kids and coaches, also played into the development of this work. We shredded. Indeed, pondering the mysteries of nature is best done deep within a wooded glade: cold smoke underfoot and drifting through the sky, an untracked line ahead.

Lastly, hats off to my wonderful friends and family. Thank you, especially, to mom and dad for the genes, support, and for all those books.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Description of Models | 2 |
| 2.1 | Lorenz System | 2 |
| 2.1.1 | Stability Analysis | 3 |
| 2.1.2 | Bred Vector Growth | 4 |
| 2.2 | Ehrhard-Müller (EM) System | 5 |
| 2.2.1 | Derivation | 5 |
| 2.2.2 | Equilibrium Solutions | 7 |
| 2.3 | FLUENT CFD | 7 |
| 3 | Data Assimilation | 8 |
| 3.1 | Derivation of the Analysis Equations | 8 |
| 3.2 | Nonlinear Data Assimilation | 10 |
| 3.2.1 | 3D-Var | 10 |
| 3.2.2 | EKF | 11 |
| 3.2.3 | Ensemble Filtering: EnKF, EnSRF | 11 |
| 4 | Methods and Results | 13 |
| 4.1 | LL | 13 |
| 4.2 | EMF | 13 |
| 4.2.1 | Parameter Estimation | 13 |
| 4.2.2 | Figures | 15 |
| 5 | Discussion | 16 |
| | Bibliography | 20 |

1 Introduction

Many of the problems associated with weather and climate forecasting do not result from our lack of knowledge of these systems' governing equations. Even though such equations are well known, nonlinearities tend to amplify uncertainty in our estimate of the current atmospheric state, known as the initial condition (IC). In addition, uncertainty in the parameters used to tune the model so that it represents physical behavior lead to forecast divergence even for the best IC. Weather models define temperature, pressure, etc. at regularly-spaced gridpoints throughout the atmosphere, but meteorologists do not have observations at every one of these locations; our knowledge of the IC is inherently incomplete. To estimate the initial state of the atmosphere, they must make educated guesses for unavailable information. Over time, these guesses are validated by new measurements, in a guess-forecast-observe cycle called data assimilation (DA). Improving upon filters from control theory, modern DA techniques are used to estimate the best ICs for numerical weather prediction (NWP). However, DA is applicable to any modeling endeavor where partial knowledge of the "truth" is available.

Ensemble forecasting is one way of representing IC uncertainty. Instead of choosing a single state for the IC, an *ensemble* of states is used. The ensemble can be thought of as an approximation of the IC's probability distribution. The ensemble states are then integrated forward in time to create an ensemble of forecasts. The ensemble spread helps measure forecast uncertainty, and the ensemble mean often gives a better guess for the true state than any one member. The ensemble members must be chosen in a clever manner, because if the deviations of multiple ensemble members from the mean trajectory are linearly dependent, then the information they carry is redundant.

In this thesis, we examine a physical system that is a toy model for atmospheric convection. The toy climate investigated in this thesis is a thermosyphon, or natural convection loop. A thermosyphon is a non-mechanical heat pump, a device in which buoyant forces move fluid through a closed-loop system. They are used in solar water heaters, cooling systems for computers, the Trans-Alaska pipeline, and nuclear power plants. They also represent the simplest model of convection, an important process in the Earth's atmosphere. The toy climate undergoes bifurcations from stable conduction to stable, then chaotic, convection, and it is illustrative of the unpredictable behavior governing weather and climate. Because the model is approximate and observations contain noise, DA is needed to estimate the IC that leads to the best forecast.

Following previous experiments, we use a toroidal geometry, analogous to a vertically-oriented hula hoop. Hot water, pumped into channels surrounding the lower half ($-\frac{\pi}{2} < \phi < \frac{\pi}{2}$), heats the fluid in the bottom of the loop by maintaining a wall temperature T_h . Similarly, cold water cools the top section of the loop ($\frac{\pi}{2} < \phi < \frac{3\pi}{2}$) to temperature T_c (see Fig. 1). The forcing is constant, i.e. we consider the case of developed flow a long time after the wall temperatures have been established. As liquid in the bottom of the loop heats past the point of simple conduction, buoyant forces overcome those of friction and viscosity, and the hot fluid rises while the cool fluid falls. The cyclical flow of hot and cold fluid forms circular structures called convection cells, an example of Rayleigh-Bénard convection. Fluid in the thermosyphon is restricted by the geometry and must initially "choose" to rotate clockwise or counter-clockwise. Each direction of rotation is said to constitute a *regime* of behavior. The initial regime is determined randomly, since the symmetry of the tube's geometry and uniform heating do not favor any particular direction. The fluid will accelerate until the buoyant force is balanced by friction and gravity, and the flow stabilizes. With sufficient heating, i.e. large enough $\Delta T_W = T_h - T_c$, the flow aperiodically reverses direction in a chaotic manner. This means that small perturbations grow exponentially in time, or that there is at least one positive Lyapunov exponent. We call each flow reversal a *regime change*. Among the goals of this research is to develop an accurate method for predicting the duration of the current and next regime, given a series of observations.

The current work has been accomplished by simulating the toy climate using a CFD model, built with the commercial software package FLUENT. A numerical CFD solution is used as a synthetic

truth from which to make observations. A laboratory thermosyphon device is in construction; in the next stage of this research, we will attempt to forecast the regime changes of the physical experiment.

This thesis is structured in the following way: In Section 2, we describe the models used for forecasting and the finite-element model meant to represent the true state of the thermosyphon, including a brief overview of the Lorenz system and a derivation of the Ehrhard-Müller system. In Section 3, we give a derivation of the Kalman filter and describe the nonlinear filters used in this experiment. In Section 4, we present results, and a full discussion with implications for NWP follows in Section 5.

2 Description of Models

Several studies have examined the periodic [3] and unstable, chaotic [4–8] behavior of toroidal thermosyphons, a phenomenon closely related to the behavior of the Lorenz system [1]. Gorman et al. [7] developed a dynamical model for constant heat flux through the bottom half, appropriate when the source of forcing is heating tape. Ehrhard and Müller [8] derived similar equations of motion with fixed wall temperature for the top and bottom halves of the loop. The experiment described here replicates these boundary conditions.

Because chaotic flow can result, thermosyphons have been found particularly interesting by modelers and engineers. One recent study [9] investigated control of chaotic convection with numerics and physical experiment. Jiang and Shoji [10] used a multiscale analysis of the field equations to derive low-dimensional models for arbitrary boundary conditions. To our knowledge, Desrayaud et al. [11] implemented the first two-dimensional (2D) finite element simulation of a thermosyphon, which captures spatiotemporal details of the flow. Ridouane et al. [12] found similar results with another 2D finite element simulation. We use this same simulation to represent the true state of the toy climate.




Fig. 1: The thermosyphon has a simple toroidal geometry. The bottom wall is heated to a constant hot temperature T_h while the top wall is maintained at the temperature T_c . This creates a temperature inversion of hot fluid below cold, an unstable situation. Energy in the bottom half must find a way to the top half. If conduction alone cannot stabilize this temperature inversion, then the fluid will start to move and convection becomes the dominant method of heat transfer.

2.1 Lorenz System

MIT meteorologist Edward N. Lorenz, one of the fathers of chaos theory, wanted to show the National Weather Service that state-of-the-art linear techniques of NWP were inadequate. Lorenz settled upon the simplest realistic model of Rayleigh-Bénard convection, where fluid held between two plates is heated from below and cooled on top. The following equations were derived using a Fourier truncation of the Navier-Stokes field equations between two infinite plates [1]. Written in dimensionless form, the Lorenz equations read

$$\frac{dx}{dt} = s(y - x) \tag{1a}$$

$$\frac{dy}{dt} = rx - y - xz \tag{1b}$$

$$\frac{dz}{dt} = xy - bz \tag{1c}$$


where $s = 10$, $b = 8/3$, and $r = 28$ are the parameters chosen by Lorenz for chaotic behavior (these parameters were used unless otherwise indicated). Physically, x is proportional to the mean fluid velocity, y to the temperature difference across the convection cell or ΔT_{3-9} (between 3 o'clock and 9 o'clock), and z is proportional to the deviation of the vertical temperature profile from linear. Parameter b relates to the geometry of the flow, s is the Prandtl number, and r is the Rayleigh number.

Lorenz's model describes the convection cells that develop in fluid placed between two plates, or classic Rayleigh-Bénard convection. It was later discovered that the equations are also applicable to toroidal thermosyphons. They have been widely used in nonlinear dynamics to study chaos and in NWP as a toy model for testing DA techniques.

The stability of the Lorenz system was investigated using numerical simulations in MATLAB. The equations can be written as a vector function $\dot{\mathbf{x}} = F(\mathbf{x})$, where \mathbf{x} is the state vector $(x, y, z)^T$ and the dot is Newtonian notation for the time derivative. The numerical integration of this autonomous ODE was performed, as is customary, with a fourth-order Runge-Kutta method and timestep 0.01.

2.1.1 Stability Analysis

The Lorenz system exhibits three equilibrium points \mathbf{x}_f , i.e. solutions corresponding to $\dot{\mathbf{x}}_f = F(\mathbf{x}_f) = 0$. The equilibrium solution at the origin corresponds to steady conduction, i.e. a motionless fluid. The other two equilibrium solutions $(\pm\sqrt{b(r-1)}, \pm\sqrt{b(r-1)}, r-1)^T$ correspond to convective states where the fluid rotates at constant velocity. For the parameter values chosen by Lorenz, all equilibria are repelling; the Jacobian matrix evaluated at each fixed point has eigenvalues exclusively outside the unit circle in the complex plane. For Lorenz's chosen s and b , and for values of the Rayleigh number $r < 1$, conduction is attracting. For $1 \leq r < 470/19 \approx 24.74$, steady convection is the attracting behavior. For $r \geq 470/19$, the system undergoes time-dependent oscillations. (Consult Lorenz's original paper for details [1].)



(a) Siblings method, units: dimensionless time (b) Jacobian method, units: dimensionless distance

Fig. 2: Stability of the Lorenz system, attractor shown in pink. In both cases, blue points are more stable than red. (a) shows an “S”-shaped region of instability punctured by stable lines extending from the convecting fixed points located in the center of each lobe of the attractor. (b) shows more stable behavior extending in the directions of the lobes. The conducting fixed point at the origin sits on a manifold of instability that extends orthogonal to the attractor. These pictures show markedly different stabilities, but this is because the eigenvalue method (b) is local compared to the divergence method (a).

For the parameters chosen by Lorenz, the system's stability was examined for a 25×25 grid of ICs $(x, y)^T$ on a $z = 25$ slice through the state space by two methods:

- Fig. 2(a): Integrate each point and its sibling (a trajectory differing by a random perturbation

of magnitude < 0.001 to each coordinate in IC). The rate of sibling divergence was estimated by measuring the 2-norm of their vector difference as a function of time. Divergence is achieved once the distance between siblings reached 5% of a reference distance $d_{\text{ref}} = 17.72^1$. The integration time until divergence is plotted.

- Fig. 2(b): The eigenvalues of the Jacobian matrix are evaluated at each point and the most positive real part is plotted. This gives an estimate of the local Lyapunov number, a measure of stretching around a trajectory. An eigenvalue λ with $\text{Re}(\lambda) > 0$ implies that nearby trajectories move away exponentially quickly in the direction of the corresponding eigenvector.

2.1.2 Bred Vector Growth

To address IC uncertainty growth, an ensemble of perturbed states which forms a small ball around the best guess is used to represent the probability distribution of that guess. A nonlinear system will dynamically stretch and shrink such a ball around its trajectory as it moves through the attractor [13]. Small perturbations to points on a trajectory, $k = 3$ orthogonal vectors of length $r = 0.001$, were integrated forward $m = 8$ time steps. The differences between perturbed and unperturbed solutions, called *bred vectors* (BVs), are used to populate the columns of a matrix \mathbf{A} . The spread is then analyzed with the singular value decomposition (SVD), $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^T$ [14]. The unitary matrix $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k]$ forms an orthonormal basis for a subspace of the state space which preserves the direction of maximum growth. The singular values s_i , entries in the diagonal matrix \mathbf{S} , measure the exponential growth rate in the direction \mathbf{u}_i according to

$$g_i \equiv \frac{1}{m} \ln \left(\frac{s_i}{r} \right), \quad (2)$$

while the average growth is

$$\bar{g} = \frac{1}{k} \sum_{i=1}^k g_i. \quad (3)$$

This quantity is shown in Fig. 3 for a typical trajectory. The average growth rate is useful because it measures the local sensitivity to small perturbations. A fast BV growth rate implies a higher forecast error, which can be used to improve forecasts in DA by adaptively inflating the forecast error estimate. BVs and SVD can also be used in ensemble initialization to ensure that ensemble members are linearly independent and span the subspace of highest uncertainty [13].




Fig. 3: Trajectory points for a typical Lorenz solution are colored by average BV growth rate \bar{g} for each growth cycle. The most dramatic growth occurs at the extrema of large oscillations. At some points near $y = 0$ the growth rate is large, due to the instability of the conducting state, but the growth rate is generally slower for smaller $|y|$. Spikes in forecast error also tend to occur at such points. Evans et al. [15] used BV growth rates to predict the occurrence of regime changes with 91.4% success. Yang et al. [2] conclude that BV growth rates can be used to predict such “catastrophes”, thus improving forecasts in adaptive DA schemes.

¹This was found by averaging the distance between uncorrelated trajectories over 2500 time units; it is approximately half the size of the attractor.

2.2 Ehrhard-Müller (EM) System

The dynamics of the thermosyphon are dominated by convection, resulting in Lorenz-like flow. Lorenz derived his equations while thinking about convection cells in the atmosphere — unconstrained by the toroidal geometry of the thermosyphon. By deriving the dynamical equations governing the thermosyphon from basic principles based on its geometry and the imposed boundary conditions, one is left with parameters in terms of physically meaningful constants and the scalings needed to make the state variables dimensionless. In principle, the physical properties of the fluid and loop define all parameters of the model. In practice, this is more complicated, because some parameters must be measured in the laboratory.

2.2.1 Derivation

Following the derivations of Ehrhard and Müller [8] and Gorman and Widmann [7], we consider the forces acting upon a control volume of fluid in the loop. All fluid properties are cross-sectionally averaged, and the radial components of velocity and heat conduction within the fluid are neglected. Incompressible flow is assumed so that the velocity $u = u(t)$ is only a function of time, not space. Applying Newton's second law, the sum of all forces on the control volume must equal its change in momentum:

$$F_p + F_f + F_g = \rho\pi r^2 R d\phi \frac{du}{dt} \quad (4a)$$

where

$$F_p = -\pi r^2 R d\phi \nabla p = -\pi r^2 d\phi \frac{\partial p}{\partial \phi} \quad (4b)$$

$$F_f = -\rho\pi r^2 R d\phi f_w \quad (4c)$$

$$F_g = -\rho\pi r^2 R d\phi g \sin \phi. \quad (4d)$$

The total force in Eqn. (4a) is comprised of the net pressure (F_p), friction from shear within the fluid (F_f), and the force of gravity (F_g). The pressure term, Eqn. (4b), is the volume times the pressure derivative. Note that the negative sign implies that if pressure increases with angle ($\partial p/\partial \phi > 0$), this force acts against the motion. The friction term, Eqn. (4c), can be considered proportional to the volume. Here, f_w is a correlation that will depend on fluid velocity, to be discussed later.

Before we write the momentum equation, it is convenient to apply the Boussinesq approximation, which assumes that variations in fluid density are linear with temperature. In other words, $\rho = \rho(T) \approx \rho_0(1 - \gamma(T - T_0))$ where ρ_0 is the reference density, γ is the coefficient of volumetric thermal expansion, and $T_0 = \frac{1}{2}(T_h + T_c)$ is the reference temperature. The Boussinesq approximation also states that the density variation is insignificant except in terms multiplied by g . Thus, the density ρ is replaced by ρ_0 in all terms of Eqn. (4) except gravity, Eqn. (4d). Using the Boussinesq approximation, gathering terms, and dividing out common factors gives the momentum equation

$$\rho_0 \frac{du}{dt} d\phi = -d\phi \left(\frac{1}{R} \frac{\partial p}{\partial \phi} + \rho_0 (1 - \gamma(T - T_0)) g \sin \phi + f_w \right). \quad (5)$$

Integrating about the loop, the momentum equation is simplified because u and f_w are independent of ϕ and other terms drop out due to periodicity in ϕ .

$$\rho_0 \frac{du}{dt} = \frac{\rho_0 \gamma g}{2\pi} \int_0^{2\pi} d\phi T \sin \phi - f_w \quad (6)$$

We now must account for the transfer of energy within the fluid, and between the fluid and the wall. The energy rate of change in the control volume is

$$\rho_0 \pi r^2 R d\phi c_p \frac{dT}{dt} = \rho_0 \pi r^2 R d\phi c_p \left(\frac{\partial T}{\partial t} + \frac{u}{R} \frac{\partial T}{\partial \phi} \right) \quad (7)$$

which must be equal to the heat transfer through the wall

$$\Delta Q = -\pi r^2 R d\phi h_w (T - T_w), \quad (8)$$

where h_w is a correlation for the heat transfer coefficient, which depends on velocity, and T_w is the temperature at the wall. Combining Eqns. (7) and (8) gives the energy equation

$$\left(\frac{\partial T}{\partial t} + \frac{u}{R} \frac{\partial T}{\partial \phi} \right) = -\frac{h_w}{\rho_0 c_p} (T - T_w). \quad (9)$$

Together, Eqns. (6) and (9) represent a simple model of the flow in the loop.

The transport coefficients f_w and h_w characterize the interaction between the fluid and the wall. They are defined by the constitutive relations

$$h_w = h_{w_0} \left(1 + K|x_1|^{1/3} \right) \quad (10)$$

$$f_w = \frac{1}{2} \rho_0 f_{w_0} u, \quad (11)$$

where u is the average velocity of fluid in the loop and $x_1 \propto u$ is the dimensionless velocity. Thus, if $|u| > 0$, frictional forces will decelerate the fluid by an amount f_w . Dimensionally, f_w is an acceleration (m/s^2) and h_w is power per unit area per unit temperature ($\text{W/m}^2\text{K}$). The parameters h_{w_0} , f_{w_0} , and K are identified by measuring heat transfer and pressure loss as fluid passes through the loop in the laboratory; the friction correlation is proportional to the pressure loss. Ehrhard and Müller actually opened the laboratory loop midway to measure these parameters. We were able to estimate the parameters empirically; the methodology is explained in Section 4.2.1.

Ehrhard and Müller solved this system of two coupled, partial differential equations by introducing an infinite Fourier series for T . The essential dynamics can be captured by the lowest modes, i.e.

$$T(\phi, t) = T_0(t) + S(t) \sin \phi + C(t) \cos \phi. \quad (12)$$

Because this ansatz separates the variables ϕ and t , the problem is transformed into a set of ordinary differential equations. Substituting Eqn. (12) into Eqn. (6) and integrating gives the equation of motion for u . Similarly, Eqn. (9) is integrated by $\int d\phi \sin \phi$ and $\int d\phi \cos \phi$ to separate the two temperature modes S and C . The system is written in dimensionless form

$$\frac{dx_1}{dt'} = \alpha (x_2 - x_1) \quad (13a)$$

$$\frac{dx_2}{dt'} = \beta x_1 - x_2 \left(1 + K|x_1|^{1/3} \right) - x_1 x_3 \quad (13b)$$

$$\frac{dx_3}{dt'} = x_1 x_2 - x_3 \left(1 + K|x_1|^{1/3} \right) \quad (13c)$$

where the following linear transformations have been made to create dimensionless variables

$$\left. \begin{aligned} t' &= \frac{h_{w_0}}{\rho_0 c_p} t \\ x_1 &= \frac{\rho_0 c_p}{R h_{w_0}} u \\ x_2 &= \frac{\rho_0 c_p \gamma g}{R h_{w_0} f_{w_0}} S \\ x_3 &= \frac{\rho_0 c_p \gamma g}{R h_{w_0} f_{w_0}} \left(\frac{2}{\pi} \Delta T_w - C \right) \end{aligned} \right\} \quad (14)$$

The parameter $\alpha = \frac{1}{2}\rho_0 c_p f_{w_0}/h_{w_0}$ is comparable to the Prandtl number, the ratio of momentum diffusivity and thermal diffusivity. Similar to the Rayleigh number, the heating parameter

$$\beta = \frac{2}{\pi} \frac{\rho_0 c_p \gamma g}{R h_{w_0} f_{w_0}} \Delta T_w$$

determines the onset of convection as well as the transition to the chaotic regime. When $K = 0$, the EM system is analogous to the Lorenz system with $b = 1$. The lack of a geometric factor b is due to the toroidal geometry, which forces a single convection cell. Anticipating a discussion of state-augmented filtering in Section 3.2.3, the parameters can be represented with a vector $\mathbf{p} = (\alpha, \beta, K)^T$.

2.2.2 Equilibrium Solutions

As in the Lorenz system, the EM system exhibits two kinds of equilibrium solutions: the conductive steady state, located at the origin in state space and corresponding to no motion; and convective steady states, corresponding to constant velocity in the clockwise or counterclockwise directions. These are found from the positive real root ξ^* of the equation

$$(\beta - 1) + (\beta - 2)K\xi^{1/6} - K^2\xi^{1/3} - \xi = 0 \quad (15)$$

which gives the convecting equilibrium $(\pm\sqrt{\xi^*}, \pm\sqrt{\xi^*}, \xi^*/(1 + K\xi^{*1/6}))^T$.

2.3 FLUENT CFD

FLUENT is a commercial computational fluid dynamics (CFD) package capable of numerical simulation of the Navier-Stokes equations for arbitrary geometries. The preprocessing program GAMBIT is used to generate a discretization of the geometry. Given a fine enough mesh, the FLUENT solution is an excellent representation of the experimental thermosyphon dynamics.

The behavior of the thermosyphon was examined in a 2D (infinite concentric cylinders as opposed to the quasi-1D torus) FLUENT system. A realistic geometry was created with $R = 36$ cm and $r = 1.5$ cm, the same values that will be used to design the physical experiment. The mesh used for the simulation has $\mathcal{O}(10^4)$ elements [12].

As in Rayleigh-Bénard convection, where fluid is held between two infinite plates, the Rayleigh number determines the onset of convection in the toroidal geometry. Here it is defined as

$$\text{Ra} = \frac{g\gamma d^3 \Delta T_w}{\nu\kappa}$$

where g is the acceleration of gravity, γ again is the thermal expansion coefficient, $d = 2r$ is the characteristic length, ν is the kinematic viscosity, κ is the thermal diffusivity, and ΔT_w is the temperature difference applied to the walls. As the Rayleigh number is increased, the behavior




Fig. 4: A FLUENT data visualization of the temperature profile near the forcing discontinuity at the right end of the heating section is shown during a regime change. Initially, the flow is clockwise. As the fluid slows down, cool fluid mixes with warm fluid in the region due to a Kelvin-Helmholtz instability [12]. Heat is always passing from the wall to the fluid, the end result being an excess of hot fluid (e) where a “tongue” of cold (a) resided in the clockwise regime. This drives the flow in the counterclockwise direction. If the system pauses between regimes for an extended period, the fluid in the bottom becomes very hot, so when it starts to rotate it does so quickly and with large deviation from the convecting equilibrium. Image credit: Ridouane.

changes from a conducting steady state to a convecting steady state and then becomes chaotic. In the experiment ΔT_w is adjusted, but for the FLUENT system g is changed to the same effect. The value used in this experiment was $\text{Ra} = 1.5 \times 10^5$.

As the FLUENT model is integrated, monitors are set to output the mass flow rate $q \equiv \dot{m}$ as well as cross-sectional average temperatures at $\phi = \pm\pi/2$ every 10 seconds. This gives a time series of observations of the synthetic toy climate to be used with the EM model in a DA-forecasting scheme.

3 Data Assimilation

DA is the process by which observations of a dynamical system are combined with forecasts from a model to estimate error covariances and calculate a “best guess” for the current state of the system. In practice, this problem is difficult, because the forecaster uses an inexact forecasting model and never knows the exact state of the dynamical system. In NWP, the number of state variables is typically $\mathcal{O}(10^3)$ times the number of observations. Nevertheless, this best guess is used to produce a forecast, which is then used in the next assimilation cycle to fill in the blanks when new observations become available. Observations and forecasts are weighted depending on the confidence in each, given their error covariance matrices.

A variety of filters are capable of solving the basic DA problem for reasonably good models. The Kalman filter (KF) is the optimal algorithm for state estimation problems in a linear system. One of its first applications was to trajectory estimation and correction of missiles and rockets. We now describe some of the DA schemes for nonlinear estimation employed in this study. In 3D-variational DA (3D-Var), the background error covariance is estimated a single time, offline, prior to the data assimilation procedure. In the extended Kalman filter (EKF), it is evolved according to the linear tangent model, which approximates the evolution of small perturbations about the trajectory. Ensemble Kalman filters (EnKFs, of which the EnSRF is an example) use ensembles of forecasts to estimate the background error and can also be used for parameter estimation. A full derivation of the KF is now presented, followed by a description of the specifics of these nonlinear methods.




Fig. 5: A cartoon of the DA process: When new observations are available, the analysis state is generated by combining observations with past forecasts. The analysis gives the optimal ICs for future forecasts. After the initial spinup, allowing the filter to adjust, analyses and forecasts should closely follow the truth.

3.1 Derivation of the Analysis Equations

The KF is well-known and widely used in linear DA and control problems. Although the thermosyphon is highly nonlinear, the linear update equations are similar to those of the more sophisticated, nonlinear algorithms used for this experiment. The KF attempts to assimilate observations and forecasts for a process of the form

$$\mathbf{x}_k^t = \mathbf{W}\mathbf{x}_{k-1}^t. \quad (16)$$

In this case, \mathbf{x}^t is the true state, which advances in time according to the linear process \mathbf{W} , which is unknown but approximated by the model \mathbf{M} . Subscripts index the time step. Using the model,

the analysis from the previous time step is integrated to generate the background forecast for the current time step

$$\mathbf{x}_k^b = \mathbf{M}\mathbf{x}_{k-1}^a \quad (17)$$

where \mathbf{M} is the linear model, \mathbf{x}^a is the old analysis, and \mathbf{x}^b is the background. Because \mathbf{M} is only an approximation of \mathbf{W} , a perfect initial condition will not lead to a perfect forecast, so

$$\mathbf{x}_k^t = \mathbf{M}\mathbf{x}_{k-1}^t + \epsilon_k^q \quad (18)$$

where the model errors ϵ^q have covariance \mathbf{Q} (usually assumed to be constant in time) and are written on the right hand side for convenience. When deemed unnecessary, subscripts are left out in the rest of the derivation.

Given an observation \mathbf{y} and background forecast \mathbf{x}^b , the KF finds the optimal way to combine them into the *analysis* \mathbf{x}^a , the best guess of the current state. This becomes the IC when forecasting, Eqn. 17. In an operational context, we usually cannot observe every state variable. If $\mathbf{x} \in \mathbb{R}^N$ and $\mathbf{y} \in \mathbb{R}^M$, then $M < N$ (in NWP $M \ll N$), so we define the observation operator $\mathcal{H} : \mathbb{R}^N \rightarrow \mathbb{R}^M$ that takes the background forecast from the model state space into the observation space. This serves two purposes: first, it avoids extrapolation of observations to gridpoints in state space; and second, it enables us to interpret our forecasts by comparing them directly to observations. For the thermosyphon, \mathcal{H} is linear, so we write it as \mathbf{H} , but this is usually not the case for the observations in NWP, e.g. satellite radiances and radar reflectivities. The KF analysis is given by

$$\mathbf{x}_k^a = \mathbf{x}_k^b + \mathbf{K}_k(\mathbf{y}_k - \mathbf{H}\mathbf{x}_k^b). \quad (19)$$

The Kalman gain or weight matrix $\mathbf{K} \in \mathbb{R}^{N \times M}$ determines the importance of the difference between observations and background forecast $\mathbf{y} - \mathbf{H}\mathbf{x}^b$, known as the *innovation*.

The derivation of the optimal weight matrix assumes that the errors in the background, observations, and analysis are drawn from Gaussian distributions:

$$\begin{aligned} \mathbf{x}^b &= \mathbf{x}^t + \epsilon^b, & \epsilon^b &\sim N(0, \mathbf{B}) \\ \mathbf{y} &= \mathbf{H}\mathbf{x}^t + \epsilon^o, & \epsilon^o &\sim N(0, \mathbf{R}) \\ \mathbf{x}^a &= \mathbf{x}^t + \epsilon^a, & \epsilon^a &\sim N(0, \mathbf{A}) \end{aligned} \quad (20)$$

where $N(\mu, \mathbf{C})$ denotes a random vector drawn from a multivariate Gaussian distribution with vector mean μ and covariance \mathbf{C} . In each case, the errors are assumed to be unbiased, i.e. their means are zero. They are also assumed to be uncorrelated to each other. The diagonal elements of the analysis error covariance matrix \mathbf{A} are the variances in the analysis state variables. The KF wants to minimize the sum of those variances, which is the same as minimizing $\text{tr}(\mathbf{A})$ with respect to \mathbf{K} .

First, recognize that Eqn. (19) implies

$$\epsilon^a = \epsilon^b + \mathbf{K}(\epsilon^o - \mathbf{H}\epsilon^b) \quad (21)$$

and the matrices $\mathbf{A} \equiv \langle \epsilon^a \epsilon^{aT} \rangle$, $\mathbf{B} \equiv \langle \epsilon^b \epsilon^{bT} \rangle$, $\mathbf{R} \equiv \langle \epsilon^o \epsilon^{oT} \rangle$, where $\langle \cdot \rangle$ indicates time average or expectation. Therefore,

$$\begin{aligned} \mathbf{A} &= \langle (\epsilon^b + \mathbf{K}(\epsilon^o - \mathbf{H}\epsilon^b)) (\epsilon^b + \mathbf{K}(\epsilon^o - \mathbf{H}\epsilon^b))^T \rangle \\ &= (1 - \mathbf{K}\mathbf{H})\mathbf{B}(1 - \mathbf{K}\mathbf{H})^T + \mathbf{K}\mathbf{R}\mathbf{K}^T \\ &= \mathbf{B} - \mathbf{K}\mathbf{H}\mathbf{B} - (\mathbf{K}\mathbf{H}\mathbf{B})^T + \mathbf{K}(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})\mathbf{K}^T, \end{aligned} \quad (22)$$

where the symmetry of the covariance matrices has been used. Now, take the trace

$$\text{tr}(\mathbf{A}) = \text{tr}(\mathbf{B}) - 2\text{tr}(\mathbf{K}\mathbf{H}\mathbf{B}) + \text{tr}(\mathbf{K}(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})\mathbf{K}^T)$$

and note that we used $\text{tr}(\mathbf{KHB}) = \text{tr}((\mathbf{KHB})^T)$. Differentiating with respect to \mathbf{K} (see [16] for relevant identities), we find

$$\frac{\partial \text{tr}(\mathbf{A})}{\partial \mathbf{K}} = -2(\mathbf{HB})^T + 2\mathbf{K}(\mathbf{HBH}^T + \mathbf{R}) \quad (23)$$

Finding the minimum requires the derivative be zero, which gives the gain

$$\mathbf{K} = \mathbf{BH}^T(\mathbf{HBH}^T + \mathbf{R})^{-1}. \quad (24)$$

Finally, we substitute Eqn. (24) into (22) to obtain the analysis covariance update, $\mathbf{A} = (1 - \mathbf{KH})\mathbf{B}$.

The complete application of the KF consists of both a forecast step and an analysis step. The forecast step is

$$\mathbf{x}_k^b = \mathbf{M}\mathbf{x}_{k-1}^a \quad (25a)$$

$$\mathbf{B}_k = \mathbf{M}\mathbf{A}_{k-1}\mathbf{M}^T + \mathbf{Q} \quad (25b)$$

and the analysis step is

$$\mathbf{x}_k^a = \mathbf{x}_k^b + \mathbf{K}_k(\mathbf{y}_k - \mathbf{H}\mathbf{x}_k^b) \quad (26a)$$

$$\mathbf{A}_k = (1 - \mathbf{K}_k\mathbf{H})\mathbf{B}_k \quad (26b)$$

with \mathbf{K}_k given by Eqn. (24).

The forecast equations create the background forecast and update the background error covariance. The new background error covariance is the old analysis error integrated forward plus the model error. In the analysis step, this background forecast is incremented by the gain times the innovation to produce the analysis. The new analysis error is equal to the background error reduced by a factor of $(1 - \mathbf{KH})$. By finding the analysis, the filter has revealed the best possible starting point for the next background forecast.

3.2 Nonlinear Data Assimilation

3.2.1 3D-Var

Rather than minimize the analysis error variance, the analysis equations can also be derived by finding the analysis state \mathbf{x}^a that minimizes the following quadratic scalar cost function:

$$C(\mathbf{x}) = \frac{1}{2} [(\mathbf{x} - \mathbf{x}^b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) + (\mathbf{y} - \mathcal{H}\mathbf{x})^T \mathbf{R}^{-1}(\mathbf{y} - \mathcal{H}\mathbf{x})]. \quad (27)$$

The cost $C(\mathbf{x})$ has its minimum at $\mathbf{x} = \mathbf{x}^a$, where \mathbf{x}^a is given by Eqn. (26). This is called the three-dimensional variational (3D-Var) method since the minimization for NWP is with respect to a state vector embedded in a three-dimensional field (latitude, longitude, and height).

Formally, both 3D-Var and the KF yield the same solution [17]. However, in this case the control variable is the analysis, while in the KF the control variable is the weight matrix itself. In operational NWP, where there are $\mathcal{O}(10^9)$ state variables, the numerical implementations of 3D-Var and the nonlinear KF are drastically different. Because 3D-Var assumes the background error \mathbf{B} is fixed in time, the Kalman gain \mathbf{K} needs to be calculated only once. This is the most computationally prohibitive part of DA because it requires solving a linear system in $\mathcal{O}(10^9)$ variables. A constant \mathbf{K} thus makes the algorithm computationally simple; the most difficult part of implementing 3D-Var is finding the optimal \mathbf{B} .

However, a static \mathbf{B} is not realistic. From a dynamical systems standpoint, uncertainty is closely related to stability, which is clearly dependent on where the state of the system is in its attractor. With the Lorenz system, for example, the true background error is typically smaller near the unstable

convecting equilibria than near the unstable conducting equilibrium (see Fig. 2). Because 3D-Var is computationally cheap, the National Centers for Environmental Prediction (NCEP) employ it in their forecasts. But it cannot detect so-called “errors of the day”, state-dependent forecast errors which grow quickly but are not represented in the 3D-Var background error covariance matrix [17].

3.2.2 EKF

The extended Kalman filter (EKF) is the logical version of the KF for a nonlinear model. Given a nonlinear model \mathcal{M} , the error covariances are updated by the *linear tangent model*

$$\mathbf{M} \equiv \left. \frac{\partial \mathcal{M}}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}^b} \quad (28)$$

which takes the place of \mathbf{M} in Eqn. (25b). This model propagates small perturbations around the trajectory \mathbf{x}^b forward in time. To operate on the matrix \mathbf{A} with the linear tangent model, first take the Jacobian of F (the right hand side of the nonlinear differential equation $\dot{\mathbf{x}} = F(\mathbf{x})$) and evaluate it at the background point \mathbf{x}^b ; call this matrix \mathbf{J} . Each column \mathbf{a}_i of \mathbf{A} , which can be thought of as an error perturbation to the analysis state, is then integrated forward in time according to the linear ODE $\dot{\mathbf{a}}_i = \mathbf{J} \mathbf{a}_i$.

Also note that if the observation operator \mathcal{H} is nonlinear, it is replaced by a similar linear tangent model \mathbf{H} in the matrix equations (24) and (26). The transpose of these matrix functions are called *adjoint models*, which are used in sensitivity analyses of the state to perturbations.

Some modifications to the EKF algorithm are necessary to prevent the covariance matrices from becoming unbalanced. Because of its inherent approximations, the EKF tends to systematically underestimate the background error. As the background error shrinks to zero, the filter stops trusting the observations and follows the background trajectory. This is known as *filter divergence*. To combat this, a multiplicative inflation of $(1 + \Delta)$ was first applied to the background covariance matrix \mathbf{B} . Additive inflation by random numbers uniformly distributed between 0 and μ were then added to the diagonal elements of \mathbf{B} , as per Yang et al. [2]

To propagate the background covariance without the explicit adjoint model, \mathbf{B} was first decomposed with the Cholesky factorization [14] into the product of a lower and upper diagonal matrix before its columns were integrated forward with the linear tangent model \mathbf{M} .

$$\mathbf{B}_{k-1} = \mathbf{L}_{k-1} \mathbf{L}_{k-1}^T \quad (29)$$

$$\mathbf{T}_k = \mathbf{M}_{k-1} \mathbf{L}_{k-1} \quad (30)$$

$$\mathbf{A}_k = \mathbf{T}_k \mathbf{T}_k^T + \mathbf{Q} \quad (31)$$

This guarantees symmetry for the new analysis error covariance \mathbf{A} .

3.2.3 Ensemble Filtering: EnKF, EnSRF

The ensemble Kalman filter (EnKF) is a method that replaces a single forecast state with an ensemble of states. The spread of the ensemble about its mean gives an approximation of the background error covariance and forecast uncertainty, while the ensemble average gives the best guess of the forecast. The EnKF was first introduced by Evensen, and he gives a comprehensive overview of ensemble filters in [18]. It was shown that if the observation, which has random error with covariance \mathbf{R} , is perturbed with N_{ens} random errors (again with covariance \mathbf{R} ; N_{ens} is the number of ensemble members), to make an ensemble of independent observations $\{\mathbf{y}_i\}$, then the background error covariance can be written

$$\mathbf{B} \approx \frac{1}{N_{\text{ens}} - 1} \sum_{i=1}^{N_{\text{ens}}} (\mathbf{x}_i^b - \bar{\mathbf{x}}^b)(\mathbf{x}_i^b - \bar{\mathbf{x}}^b)^T \quad (32)$$

which is simply the ensemble average outer product of members' deviation from the mean. The background forecast of ensemble member i is denoted \mathbf{x}_i^b , and $\bar{\mathbf{x}}^b$ is the background forecast ensemble average. In this case, each ensemble member is updated according to the KF equations for their associated observation. The advantages of the EnKF are many: there is no linear tangent model to compute, the number of ensemble members can be small ($\mathcal{O}(10^2)$ for a global weather model) relative to the dimensionality of the state space, and prior knowledge about the structure of the forecast errors is not necessary.

The above approach adds noise to the observation to create linearly independent observations and is classified as a *perturbed observations* method [19]. This necessarily introduces additional sampling error into the forecast. For this reason, Whitaker and Hamill [20] introduced the ensemble square root filter (EnSRF) as an improved EnKF. In the EnSRF, the ensemble mean is updated with the traditional Kalman gain (Eqn. (24))

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^b + \mathbf{K}(\mathbf{y} - \mathcal{H}\bar{\mathbf{x}}^b) \quad (33)$$

and deviations from the mean are updated with a different gain

$$\mathbf{x}'^a = \mathbf{x}'^b + \tilde{\mathbf{K}}(\mathbf{y} - \mathcal{H}\mathbf{x}'^b) \quad (34)$$

where the primes denote the perturbations, e.g. $\mathbf{x}'^b = \mathbf{x}^b - \bar{\mathbf{x}}^b$, and

$$\tilde{\mathbf{K}} = \mathbf{B}\mathbf{H}^T \left[\left(\sqrt{\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R}} \right)^{-1} \right]^T \times \left[\sqrt{\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R}} + \sqrt{\mathbf{R}} \right]^{-1}. \quad (35)$$

When the observation is a scalar, it can be shown that

$$\tilde{\mathbf{K}} = \left(1 + \sqrt{\frac{\mathbf{R}}{\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R}}} \right)^{-1} \mathbf{K}. \quad (36)$$

If observation errors are uncorrelated (\mathbf{R} is diagonal), then Eqn. (36) can be used to process observations one at a time [20]. The updated analysis ensemble is then $\{\mathbf{x}_i^a\}$, where $\mathbf{x}_i^a = \bar{\mathbf{x}}^a + \mathbf{x}'_i^a$. Square root filters have better numerical stability and speed than their standard KF counterparts. For these reasons, the Potter square root filter was employed for navigation in the Lunar Module of the Apollo program.

The EnKF can also be modified slightly to estimate the model parameters by *state augmentation*. In state-augmented filtering, the vector of parameters \mathbf{p} is treated as part of the state vector, hence $\mathbf{x} \rightarrow \begin{pmatrix} \mathbf{x} \\ \mathbf{p} \end{pmatrix}$. The model equations are augmented with $\dot{\mathbf{p}} = 0$. State-augmented filtering is accessible with the EnKF because it requires no prior knowledge of the model covariances. Although state-augmented filtering has been shown to work for the Lorenz system with a perfect model and the initial ensemble of parameters centered on the correct values [21], the filter may not be able to converge to the correct parameters from a poor initial guess.

As with the EKF, ensemble filters tend to underestimate the background error, resulting in an ensemble spread which is typically less than the root mean square (RMS) error (together, RMSE) found when verifying. Multiplicative inflation of \mathbf{B} by $(1 + \Delta)$ is a common method, shown successful in [2, 18–21]. The most reliable way we found to prevent filter divergence was to use a hybrid EnKF/3D-Var scheme. The background error is written as a combination of the background error from the EnKF and from 3D-Var

$$\mathbf{B} = \theta\mathbf{B}_{\text{ens}} + (1 - \theta)\mathbf{B}_{\text{3D-Var}} \quad (37)$$

where θ is a tuning parameter. In this study, the optimal combination was found for $\theta \sim 0.5$, indicating that the ensemble estimate is only as reliable as 3D-Var. Currently, 4D-Var (like 3D-Var except taking into account older observations) and ensemble filters are the most promising candidates being considered to replace 3D-Var.

4 Methods and Results

Due to the variety of models and data assimilation algorithms, the scenarios are listed by “forecasting model” then “true system”: Lorenz forecasting Lorenz (LL) and EM forecasting FLUENT (EMF).

4.1 LL

Results similar to those of Yang et al. [2] were achieved with EKF and 3D-Var when assimilating a synthetic Lorenz truth with a Lorenz model. The same parameters were used in all cases. Gaussian noise with a standard deviation of $\sqrt{2}$ was added to each variable of the numerically generated truth every 8 time steps to create the observations. The analysis was initialized from a point far away in phase space. Performance was measured by assimilating observations and measuring the RMSE of the analysis versus the truth (without added noise). Because all state variables are known for both truth and forecast, the RMSE was calculated

$$\text{RMSE} = \sqrt{\langle \delta x^2 + \delta y^2 + \delta z^2 \rangle / 3} \quad (38)$$

where each δx_i is the difference between forecast and truth in the i th state variable, e.g. $\delta x = x^a - x^t$ for analysis error or $\delta x = x^b - x^t$ for background error. The results are shown in Table 1 and Fig. 6.

4.2 EMF

As described in 2.3, the mass flow rate q and cross-sectionally averaged temperature at $\phi = \pm\pi/2$ were reported by the FLUENT simulation. The mass flow rate gave the fluid velocity $u = q/(2r\rho_0)$, since $2r$ is the 2D cross-sectional area. The temperature measurements were used to calculate ΔT_{3-9} . The temperature data, however, was very noisy compared to the smooth trajectories typical of the model. This is due to the neglect of higher order temperature modes and the presence of a Kelvin-Helmholtz instability [12] at the boundary between heat source and sink, shown in Fig. 4. Assimilation of q alone proved sufficient. Gaussian noise with standard deviation equal 1% of the mass flux climatological mean $\sqrt{\langle q^2 \rangle}$, approximately 6×10^{-4} kg/s, was added to make the observations realistic. In this case, the observations y , which are of state variable q , provide the only validation, so the RMSE is calculated in observation space


$$\text{RMSE} = \sqrt{\langle \delta q^2 \rangle} \quad (39)$$

where $\delta q = \mathbf{H}\mathbf{x}^b - y$.


4.2.1 Parameter Estimation

Before any forecasts could be generated, the parameters matching the EM model to the FLUENT simulation needed to be determined. A state-augmented EnKF was applied to the problem but failed to reveal which EM parameters best described the FLUENT system. The method that ultimately succeeded was a straightforward search of parameter space.


The next parameter estimation attempt used a state-augmented EnKF. Starting with the parameters determined by Ehrhard and Müller, $\mathbf{p} = (15, 38, 0.35)^T$, and assimilating q , the EnSRF converged to a triple of parameters that produce periodic behavior. These parameters yield an acceptable forecast within one assimilation cycle, but the forecast soon becomes useless because the model can no longer predict chaotic behavior. The filter was also sensitive to the initial guess of parameters, with different initializations converging to different values. Changing the parameters of the model each assimilation cycle was also problematic because they affect the period and height of the EM model oscillations, which are used to empirically estimate the scalings to dimensionless velocity and time. As a result, the observation operator and the dimensionless time scale would



(a)



(b)



(c)

| | RMSE: | analysis | |
|---|-----------|----------|-----------|
| | Observed: | y | x, y, z |
| 3D-Var | | 0.80 | 0.64 |
| EKF ($\Delta = 0.05, \mu = 0.02$) | | 0.51 | 0.34 |
| EnSRF ($\Delta = 0.04, N_{\text{ens}} = 6$) | | 0.55 | 0.28 |
| ($\Delta = 0.04, N_{\text{ens}} = 3$) | | 0.77 | 0.65 |

Table 1: (Above) LL average analysis RMSE from assimilation cycles 500 to 2500. It compares the performance of 3D-Var, EKF, and EnSRF DA algorithms assimilating observations which have error $\sqrt{2}$ every 8 steps. The results for 3D-Var and EKF match very closely with those of Yang et al. [2]. The results for EnSRF with 6 members match what was found using an ensemble transform Kalman filter (ETKF, another EnKF variant) by Kalnay et al. [19], but the EnSRF with 3 members has roughly twice the RMSE found with the ETKF.


Fig. 6: (Left) Regime duration plotted against the amplitude of the last oscillation in the preceding regime for different state variables of the Lorenz system. Each “step” corresponds to an integer number of oscillations. The x (6(a)) and y (6(b)) variables show that duration and oscillation amplitude are positively correlated, while the z (6(c)) variable shows an inverse relationship. This is because x and y are directly proportional to the size of the temperature anomaly in the convection cell (see Section 5). On the other hand, larger z indicates that more of the anomaly is located in the bottom half of the loop, meaning that the state is near the highly unstable conducting equilibrium, so the new regime will not last. Each regime change is colored by the 8-step BV average growth rate at last extremum before regime change. The figure shows a clear gradient in these growth rates. This is because the local Lyapunov numbers decrease as the state moves closer to the connecting equilibria.

change as the parameters are varied. We continue to investigate whether state-augmented ensemble filtering can be used for parameter estimation in the EMF forecasting scheme.


The appropriate parameters were found by a simple search of the parameter space for $\alpha \in \{1, 2, \dots, 20\}$, $\beta \in \{20, 21, \dots, 55\}$, and $K \in \{0, 0.01, \dots, 0.5\}$. This required integrating the EM system 36720 times for 2500 time units. Because of the size of the search space, this was performed in parallel using the Vermont Advanced Computing Center (VACC) cluster. Note that for more complicated models, this brute force search of the parameter space becomes computationally prohibitive. The optimal EM parameters were chosen to be those that produce solutions which most closely match the climatology of the FLUENT system, as measured by three statistics. The first two of these statistics characterize regime change behavior: the average number of oscillations per regime and the maximum number of oscillations before a regime change. The third statistic measures the relative height of the velocity oscillations. The relative oscillation height was measured by computing of RMS distance from the convecting equilibrium and dividing by the magnitude of the convecting equilibrium itself. The three statistics were calculated first for the FLUENT q time series, then for the x_1 time series generated for each triple in the parameter search space. Only those parameter triples that produced a maximum of 7 oscillations before regime change, the same as was found from the q data, were considered. Minimizing the 2-norm difference, between model runs and FLUENT data, of the remaining two statistics gave the model parameters $\mathbf{p} = (7, 33, 0.07)^T$.

Given this set of parameters, the dimensionless time scale and observation operator were estimated empirically. The time scale constant of proportionality (t'/t) was given by τ_{EM}/τ_{FLUENT} , where the τ 's are the average oscillation periods in steps of the corresponding time series. The one nonzero entry in the observation operator connects dimensionless velocity to velocity. The velocity scale ($[\mathbf{H}]_{11} = x_1/u$) was calculated with the ratio of average velocity magnitudes $|x_1|/|u|$.


4.2.2 Figures



(a)



(b)



(c)

Fig. 7: Results of 3D-Var assimilating the same data are shown for three different assimilation windows. In (a), observations are made frequently enough to keep the forecast close to the truth, whereas in (b) DA is incapable of keeping the forecast in the correct regime. In (c), the filter has satisfactory overall performance (RMSE $\approx 35\%$ of $\sqrt{\langle q^2 \rangle}$); note the error spike around 8×10^3 s when the forecast and truth end up in different regimes. The largest errors tend to occur at regime changes.




Fig. 8: Regime duration plotted against the amplitude of x_1 at the last extremum before regime change. This amplitude is calculated from the EM model 3D-Var analysis of FLUENT observations, using a 30 s assimilation window. Inset, upper left: timeseries corresponding to a single point in the scatter plot, marked with a black cross. This figure contains almost 50 days of data and 2323 regime changes. Points are colored by the average BV growth rate over the preceding assimilation window, showing a BV growth rate gradient similar to what was found for the Lorenz system. Outliers with the largest deviations from convecting equilibrium result in shorter regime durations than expected. A possible explanation is that this is due to mixing that occurs when the temperature anomaly is very large, effectively setting an upper limit on the number of oscillations in a regime; this phenomenon would not occur in the Lorenz-like models. Inset, lower right: histogram showing the likelihood of different regime durations, given an x_1 amplitude between 12 and 13. The most probable duration is about 33 minutes, the middle “step” for that amplitude range.

5 Discussion

DA was shown to be an effective way of coupling a low-dimensional model to complex simulations of the toy climate. The DA algorithms were first implemented for a perfect model, the LL scenario, giving accuracies comparable to other studies. In the EMF scenario, background errors were always larger than observational noise, but the filters, whose individual performance varied, were still able to produce analyses within observational noise of the truth. Finally, the amplitude of the last oscillation in the current regime has been shown to correlate with the duration of the following regime in both LL and EMF scenarios (Figs. 6 and 8). Using oscillation amplitudes and BV growth rates, we should be able to predict the occurrence of regime changes and residencies. Because the EnKF performed poorly compared to the results of Kalnay et al. [19], we expect further tuning to improve background and analysis RMSE. In the future, DA using the EM model, and possibly a full FLUENT simulation, will be applied to data from the laboratory thermosyphon. The effectiveness of different DA approaches for these imperfect model experiments can inform meteorologists as to which approaches are best suited for global weather and climate models.

Implementing DA in the LL scenario was an important first step before EMF forecasting was

considered. Numerous investigations have examined the Lorenz system, so we could be confident the filter was properly coded once it performed as well as reported in other studies. The filters’ sensitivity to tuning parameters is high, and running simulations for a large range of tuning parameters is time-consuming, especially if the modeler does not know which regions of parameter space to sample. Previous studies gave us an idea of what to expect. Tuning the LL forecasting code helped the author develop an intuitive understanding of how DA schemes are optimized. Once the filters were properly tuned, the 3D-Var and EKF errors closely matched what was found in the study by Yang et al. [2]. A similar study, carried out by Kalnay et al. [19], included results for an ETKF ensemble filter, against which our EnSRF results are compared. We found comparable analysis errors with a 6 member ensemble, although when cut to 3 members the EnSRF had approximately twice the RMSE of the ETKF with the same ensemble size. It appears likely that this discrepancy is due to differences in the two ensemble DA algorithms.

Forecasting the Lorenz system with a perfect model (LL) was also useful in the analysis of regime changes. Knowledge of the true state was always available, so we were able to plot oscillation amplitude versus regime length for all three state variables (Fig. 6). The strong correlation between amplitudes and residencies prompts a discussion as to why the system behaves this way. It was counterintuitive to think that a large oscillation in one regime should result in smaller oscillations in the next. One would not think that being far from one convecting equilibrium brings the system close to the other convecting equilibrium, but that is exactly what happens.




Fig. 9: FLUENT simulation showing the temperature profile (units: K) of steady convecting flow, rotating counterclockwise. In the chaotic case, opposite anomalous regions of warmer and cooler fluid will be superimposed on this temperature profile. As these pass through the loop, the “tongues” of warm and cool fluid extending into the top and bottom halves of the loop will grow and shrink simultaneously until the hot tongue (visible near 2 o’clock) reaches the opposite side of the loop. The flow will then stall and reverse direction. Image credit: Ridouane.

viscous and thermal dissipations are in phase, so that an increase (decrease) in flow rate leads to an increase (decrease) in friction and a decrease (increase) in buoyancy, and such perturbations are damped out [4, 5].

However, not all regime changes occur when the system reaches the same oscillation amplitude, nor do all regimes last the same amount of time. In the Lorenz system, larger oscillations in the x and

The first explanation of regime changes was presented by Welander [4] and repeated by Creveling et al. [5]. Welander, who was also the first to discover that thermosyphons exhibit aperiodic oscillatory behavior, explained the instability of steady convecting flow by considering a thermal anomaly or “warm pocket” of fluid. With each cycle through the loop, the warm pocket is amplified due to out of phase viscous and thermal dissipations. Welander explained that when the warm pocket emerges from the heating section and enters the cooling section, it feels a greater buoyant force than the surrounding fluid and accelerates, exiting the cooling section quickly. Thus, it has less time to radiate away its energy. As the pocket moves into the heater, the buoyant force it experiences is again higher than normal, so now the pocket decelerates and passes slowly through the heating section, gaining more energy.

This positive feedback effect causes the pocket to grow hotter and larger, and it manifests as growing oscillations of x and y about the unstable convecting equilibrium. But these oscillations do not increase forever; eventually, the pocket becomes large and hot enough that its descent towards the heating section is stopped entirely by its own buoyancy. Without movement, the pocket dissipates, but its remnant heat biases new rotation in the opposite direction as before: a regime change. For low heating rates, the convecting equilibrium is stable. This is because vis-

y variables immediately before regime change lead to longer residences in the new regime (Figs. 6(a) and 6(b)). These oscillations reflect the temperature of the anomalous pocket described earlier. This means that a longer lasting new regime requires a large fluctuation from convective equilibrium in the velocity and horizontal temperature difference of the fluid before regime change. An inverse relation is true for z (Fig. 6(c)). This is because large z corresponds to higher temperatures in the bottom of the loop than in the top, a situation close to the highly unstable conducting equilibrium state, hence the system is destabilized as it enters the new regime. A lasting regime needs to begin nearby the corresponding convecting equilibrium (in state space), so that the initial oscillations in the state variables are small and have time to grow.

During a regime change, the fluid motion stalls after the center of the pocket enters the heating section. Hot fluid extending from the heating section to the cooling section in the direction of the earlier motion (see Fig. 9) will dissipate without any rotation. If the pocket is large (corresponding to a large deviation from convecting equilibrium in x, y), some of the anomaly will extend into the cooling section, creating a temperature profile very near that of steady convection, but in the new direction. The resulting regime lasts longer since the instability requires more time to grow before causing the next reversal. If the pocket is small, it will mostly dissipate, bringing the temperature profile close to that of conduction. This is a more unstable equilibrium, since the vertical temperature gradient builds until the fluid in the bottom is much hotter than the fluid above. When the fluid begins to rotate, it accelerates rapidly. The large amount of heat carried by the fluid creates a very hot anomalous pocket, which carries the system state far from the convecting equilibrium. Animations of the CFD-simulated temperature field during regime reversal, created in FLUENT by El Hassan Ridouane, are consistent with this explanation.

With DA functioning for the Lorenz system and a more complete understanding of the mechanism for regime changes, our research focus shifted to the problem of forecasting the realistic, CFD-simulated thermosyphon. The first difficulty we encountered was the estimation of appropriate EM model parameters, described in detail in section 4.2.1. Although the empirical method was successful, scalable methods of parameter estimation are needed for higher-dimensional models.

Model errors further complicate the implementation of DA in realistic forecasting scenarios. For instance, to test the effect of model errors on different DA schemes, Kalnay et al. [19] created a Lorenz truth with $r = 28$ but used a forecasting model with $r = 26$. Their analysis RMSE for a 3 member EnKF, observing x, y, z every 8 steps, went from 0.30 to 0.81, and the necessary multiplicative inflation factor increased from $\Delta \approx 0.08$ to 1.2. Model error is likely even more significant for EMF forecasts, since the EM model is a weaker approximation of the FLUENT system than the almost-perfect model used above. Within the KF update equations, model error appears as the \mathbf{Q} covariance matrix. We set $\mathbf{Q} = 0$ and counted on additive and multiplicative background covariance inflation capturing model error; better results might have been obtained with an optimized \mathbf{Q} .

Another difficulty of operational DA is incomplete knowledge of the system state. For this reason, errors were measured by transforming the state vector into observation space. Analysis error in observation space, which can be small even for large assimilation windows, is not an appropriate metric for assessing model performance. For example, 3D-Var in the EMF scenario with a 10 minute time window yields analysis and background errors of 0.000506 and 0.0703, respectively. The analysis error is less than the observational noise, but the background error is on the order of the climatological mean (RMS) mass flow rate $\sqrt{\langle q^2 \rangle}$. In other words, background forecasts are practically meaningless, but the filter “knows” this and weights the observations heavily over the background forecasts. The analyses generated are close to observations, but the unobserved model state variables are poor predictions.

The performance metric used instead was the relative background error with respect to the climatological mean. Following the S1 score convention of NWP [17], relative error above 70% is considered a “useless” forecast, while under 20% the forecast is “perfect”. Perfect EMF forecasts for 3D-Var and EnSRF were found up to and including 3.5 minute assimilation windows; the EKF performed perfectly with 4 minute windows (Fig. 10).




Fig. 10: Relative background RMSE with respect to $\sqrt{\langle q^2 \rangle}$ plotted for different DA algorithms against assimilation window. As the window becomes larger, the relative error increases until saturation (3D-Var, inset). The dashed line in the main figure shows the limit of a “perfect” forecast; in the inset, this is the lower dashed line, and the upper dashed line demarcates a “useless” forecast [17].

The next phase of this study will use DA and breeding to qualitatively predict regime changes and new regime duration. Evans et al. [15] set a BV growth rate threshold above which regime changes occurred with a 91.4% success rate. Besides informing when regime changes will occur, the evident gradient of bred vector growth rates up the “steps” in Figs. 6 and 8 should be useful in determining the duration of the upcoming regime. In the more distant future, the imperfect model experiment we have devised could be used to compare the relative performance of other DA algorithms (4D-Var, ETKF) or synchronization approaches (adaptive nudging, see [2]).

Although the thermosyphon is far from representing anything as complex and vast as Earth’s weather and climate, there are many things our model has in common with global atmospheric models. Both are, at best, only an approximate representation of the numerous processes that govern the Earth’s climate. Global models and the EM model both parameterize fine-scale processes that determine large-scale behavior. For example, clouds and precipitation are sub-grid-scale process in a global weather model, and the correlations for the heat transfer and friction coefficients are parameterizations of fluid behavior on a finer scale than can be dealt with in the derivation.

The methods we use to forecast the toy model are also similar to the methods used for global geophysical systems. Both require state estimation to find the IC from which to generate forecasts. Also, when we make forecasts in either system, we care more about the climatology than the specific behavior: In the case of the thermosyphon, this would be the occurrence and duration of regime changes. In the case climate, globally and regionally-averaged temperatures and their effects on rainfall, ice cover, etc. are the important unknown variables, due to the threat of global climate change. All of these are statistics that must be post-processed from the model output. To meet these global challenges, we need many techniques in our modeling toolbox. Toy models can provide us with insights that are applicable to the most important prediction problems of today.

References

- [1] Edward N. Lorenz. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2):130–141, 1963.
- [2] Shu-Chih Yang, Debra Baker, Hong Li, Katy Cordes, Morgan Huff, Geetika Nagpal, Ena Okereke, Josue Villafañe, Eugenia Kalnay, and Gregory S. Duane. Data assimilation as synchronization of truth and model. *Journal of the Atmospheric Sciences*, 63:2340–2354, 2006.
- [3] Joseph B. Keller. Periodic oscillations in a model of thermal convection. *Journal of Fluid Mechanics*, 26:599–606, 1966.
- [4] Pierre Welander. On the oscillatory instability of a differentially heated fluid loop. *Journal of Fluid Mechanics*, 29:17–30, 1967.
- [5] H.F. Creveling, J.F. de Paz, J.Y. Baladi, and R.J. Schoenhals. Stability characteristics of a single-phase free convection loop. *Journal of Fluid Mechanics*, 67:65–84, 1973.
- [6] M. Gorman and P.J. Widmann. Chaotic flow regimes in a convection loop. *Physical Review Letters*, 52:2241–2244, 1984.
- [7] M. Gorman, P. J. Widmann, and K. A. Robbins. Nonlinear dynamics of a convection loop: a quantitative comparison of experiment with theory. *Physica D*, 19:255–267, 1986.
- [8] Peter Ehrhard and Ulrich Müller. Dynamical behaviour of natural convection in a single-phase loop. *Journal of Fluid Mechanics*, 217:487–518, 1990.
- [9] Po Ki Yuen and Haim H. Bau. Optimal and adaptive control of chaotic convection — theory and experiments. *Physics of Fluids*, 11:1435–1448, 1999.
- [10] Y.Y. Jiang and M. Shoji. Spatial and temporal stabilities of flow in a natural circulation loop: Influences of thermal boundary condition. *Journal of Heat Transfer*, 125:612–623, 2003.
- [11] Gilles Desrayaud, Alberto Fichera, and Manuel Marcoux. Numerical investigation of natural circulation in a 2d-annular closed-loop thermosyphon. *International Journal of Heat and Fluid Flow*, 27:154–166, 2006.
- [12] El Hassan Ridouane, Christopher M. Danforth, and Darren L. Hitt. A numerical study of chaotic flow in a 2d natural convection loop. submitted, 2009.
- [13] Christopher M. Danforth and James A. Yorke. Making forecasts for chaotic physical processes. *Physical Review Letters*, 96(14):144102–4, 2006.
- [14] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 3rd edition, 1996.
- [15] Erin Evans, Nadia Bhatti, Jacki Kinney, Lisa Pann, Malaquias Peña, Shu-Chih Yang, Eugenia Kalnay, and James Hansen. RISE: Undergraduates find that regime changes in Lorenz’s model are predictable. *Bulletin of the American Meteorological Society*, pages 520–524, 2004.
- [16] Kaare Brandt Petersen and Michael Syskind Pedersen. The matrix cookbook. <http://matrixcookbook.com>, 2008.
- [17] Eugenia Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2002.
- [18] Geir Evensen. The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4):343, 2003.

- [19] Eugenia Kalnay, Hong Li, Takemasa Miyoshi, Shu-Chih Yang, and Joaquim Ballabrera-Poy. 4-d-var or ensemble Kalman filter? *Tellus*, 59A:758–773, 2007.
- [20] Jeffery S. Whitaker and Thomas M. Hamill. Ensemble data assimilation without perturbed observations. *Monthly Weather Review*, 130:1913–1924, 2002.
- [21] J.D. Annan and J.C. Hargreaves. Efficient parameter estimation for a highly chaotic system. *Tellus*, 56A:520–526, 2004.