International Journal of Robotics Research
000(00):1–13

Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI:doi number
http://mms.sagepub.com

Toward Efficient Navigation in Uncertain Gyre-Like Flows

Christoffer R. Heckman* **and Ira B. Schwartz**

Nonlinear Systems Dynamics Section, Plasma Physics Division, Code 6792, U.S. Naval Research Laboratory, Washington, DC 20375, USA

M. Ani Hsieh†

SAS Laboratory, Mechanical Engineering & Mechanics Department, Drexel University, Philadelphia, PA 19104, USA

Abstract

We present the development and experimental validation of an autonomous surface/underwater vehicle (ASV/AUV) control strategy that leverages the environmental dynamics and uncertainty to navigate in a stochastic fluidic environment. We assume the workspace is composed of the union of a collection of disjoint regions, each bounded by Lagrangian coherent structures (LCS). LCS are dynamical features in the flow field that behave like invariant manifolds in general time-invariant dynamical systems and delineate the boundaries of attraction basins. We analyze a passive particle's noise-induced transition between adjacent LCS-bounded regions and show how most probable escape trajectories with respect to the transition probability between adjacent LCS-bounded regions can be determined. Additionally, we show how the likelihood of transition can be controlled through minimal actuation. The result is an energy efficient navigation strategy that leverages the inherent dynamics of the surrounding flow field for mobile sensors operating in a noisy fluidic environment. We experimentally validate the proposed vehicle control strategy and analyze its theoretical properties. Our results show that the single vehicle control parameter exhibits a predictable exponential scaling with respect to the escape times and is effective even in situations where the structure of the flow is not fully known and control effort is costly.

Keywords
Underwater robots, navigation, Lagrangian coherent structures

* Current address: ARPG Laboratory, Department of Computer Science, University of Colorado, Boulder, CO 80309-0430
† Corresponding author; e-mail: mhsieh1@drexel.edu

1. Introduction

To better understand various physical, chemical, and biological processes in the ocean, researchers are employing autonomous underwater and surface vehicles (AUVs and ASVs) to monitor ocean salinity and temperature profiles (Lynch et al., 2008; Wu and Zhang, 2011; Sydney and Paley, 2011), the dynamics of algae blooms (Zhang et al., 2007; Chen et al., 2008; Caron et al., 2008; Das et al., 2012), and the dispersion of hazardous contaminants. However, working in geophysical fluid environments like the ocean presents significant challenges. Underwater communication is severely limited which means AUVs not only have to actuate control in the high inertia environment, they must also operate with minimal or no human supervision. And since the ocean spans large physical scales, is stochastic, and has time-varying dynamics (Trulsen, 2006; Ochi, 1998), accounting for the dynamics of the surrounding fluidic environment is extremely challenging. As such, sensors operating in these time-dependent and stochastic environments will tend to escape from their monitoring regions of interest which limits the amount of useful data that can be collected.

More recently, researchers have shown that AUV/ASV motion planning and adaptive sampling strategies can be improved by incorporating either historical ocean flow data (Smith et al., 2010a,b, 2012) or multi-layer partial differential equation (PDE) models of the ocean (Wang et al., 2009; Lolla et al., 2012). However, accessibility to and the overall quality of the flow data and/or numerical models is highly dependent on how well a given region of interest is instrumented. This is because numerical PDE models are often derived through a combination of theoretical and field observations. Current ocean current hindcasts, nowcasts, and forecasts provided by Naval Coastal Ocean Model (NCOM) databases (SCRIPPS, 2014) and regional ocean model systems (ROMS) (Smith et al., 2010b) are assimilated from satellite and field observations in conjunction with predictions from numerical PDE models (Shchepetkin and McWilliams, 1998, 2005). Furthermore, despite the U.S. Navy's and various public and private organizations' efforts in the last thirty to forty years to deploy a combination of stationary, surface, and at-depth sampling technologies, existing data sets that describe geophysical flows are still mostly finite-time and of low spatio-temporal resolution. While researchers have developed various analytical and numerical models of geophysical flows that exhibit specific dynamical features, these models are often scale invariant. As such, any attempts to incorporate historical and/or forecasted ocean flow data into existing AUV/ASV motion planning and control strategies becomes an extremely challenging endeavor.

In the last two decades, Lagrangian pictures of geophysical flows have led to new dynamical systems tools, including the study of transport by coherent structures Provenzale (1999). These studies have shown that these ubiquitous structures play a key role in transport and give insight into the dynamics of the fluidic environment. More interestingly, minimum energy paths in the ocean for AUVs/ASVs can coincide with a specific class of coherent structures called Lagrangian coherent structures (LCS) (Inanc et al., 2005; Senatore and Ross, 2008). Topologically, LCS are extensions of stable and unstable manifolds to general time-dependent flows and divide the flow into dynamically distinct regions (Haller and Yuan, 2000). For two-dimensional (2D) flows, LCS are analogous to ridges defined by local instability, and can be quantified by Finite-Time Lyapunov Exponents (FTLE) which provide a local measure of how fast a region in the flow field is expanding (Shadden et al., 2005). Since these structures are robust features along which organized and persistent flows exist for a relatively long time (D'Asaro et al., 2011; D'Asaro and et al, 1996), it makes sense that minimum energy or energy optimal paths for AUVs/ASVs would tend to trace or follow the LCS boundaries. Coupled with recent work showing the relative robustness of strong and coherent LCS to uncertainty (Lermusiaux and Lekien, 2005; Lermusiaux et al., 2006; Lermusiaux, 2006), these results suggest that LCS have the potential to provide a reduced-order representation of the surrounding geophysical fluid dynamics. Based on these insights, we hypothesize that it is possible to achieve significant improvements by incorporating LCS information into AUV/ASV motion planning and control strategies at a significantly reduced computational burden when compared to existing strategies that rely on PDE ocean model, NCOM, and/or ROMs data.

In this work, we consider the development of ASV/AUV control strategies that leverages the environmental dynamics and uncertainty to navigate in a stochastic fluidic environment. We build on our previous work in collaborative tracking of

LCS using teams of autonomous robots (Hsieh et al., May; Michini et al., 2014) and distributed allocation of autonomous sensing resources in complex flows (Mallory et al., 2013b), and consider the synthesis of energy efficient actuation strategies for mobile sensors operating in a collection of disjoint LCS bounded regions. We draw upon ideas presented in (Inanc et al., 2005; Senatore and Ross, 2008; Schwartz et al., 2011, 2010) on energy optimal paths in the ocean and optimal paths in stochastic systems to show how it is possible to increase/reduce uncertainty near the LCS boundaries through controls. The present work focuses on the synthesis of energy efficient control strategies as opposed to energy optimal strategies for autonomous mobile sensors operating in complex flow fields. This distinction is important since the computation of energy optimal strategies can be expensive, especially for resource constrained vehicles. As such, we focus on suboptimal strategies that can be implemented through simple actuation strategies. The result is an energy efficient control strategy that enables resource constrained mobile sensors to actively change the likelihood of their escape, or break out, from their monitoring regions and prolonging their operational lifespan.

Recently, new mathematical methods have been developed to elucidate and harness the effects of noise on dynamical switching behavior (Forgoston et al., 2011; Bollt et al., 2002; Billings and Schwartz, 2008; Froyland and Padberg, 2009; Billings et al., 2008). These methods accurately predict the expected switching time of particles between distinct basins of attraction. Moreover, the techniques predict the most probable (or optimal) escape path from a region resulting from a large fluctuation due to the underlying noise. Building on these techniques, we identify and analyze the most likely routes of transport between two adjacent LCS-bounded regions for passive vehicles, *i.e.*, vehicles under the influence of the fluid flow alone and with negligible inertia. We analyze the escape times via these most likely routes of transport and show that they correspond to fuel efficient paths since the environmental dynamics and noise is leveraged for navigation. From these results, we design a simple control strategy and show how it can effectively manipulate a vehicle's escape times from an LCS-bounded region and achieve comparable or improved fuel efficiency. The result is a control effort strategy that can be implemented with limited knowledge of the flow field to allow a vehicle to navigate from one basin of attraction to another by leveraging the surrounding environmental dynamics.

The paper is organized as follows: We formulate the problem and outline key assumptions in Section 2. The identification and analysis of the most likely paths for escape and the development of our control strategy is presented in Section 3. We briefly describe our experimental methodology in Section 4 and present our experimental results in Section 5. These experimental results demonstrate that the theory we have developed for particles applies to macroscopic systems of autonomous vehicles operating in noisy environments. This has implications for the design of controllers in environments where the structure of the flow is not necessarily known, significant stochastic fluctuations are present, and control effort is costly. We conclude with a discussion of our results and experimental insights in Section 6 and directions for future work in Section 7.

2. Problem Statement

We consider the deployment of a planar autonomous underwater or surface vehicle (AUV or ASV) within a two-dimensional (2D) obstacle free fluidic environment. We are interested in the deployment of mobile sensors with minimal actuation capabilities for environmental monitoring and sampling applications. The objective is the synthesize control strategies that will allow the vehicle to leverage the surrounding fluid dynamics to maneuver within the workspace. We assume the following 2D stochastic differential equation representing a kinematic model with multiplicative noise for the AUV:

$$\dot{\boldsymbol{q}}(t) = \boldsymbol{u}(t) + \boldsymbol{F}(\boldsymbol{q}(t)) + \boldsymbol{G}(\boldsymbol{q}(t))\boldsymbol{\eta}(t), \tag{1}$$

where $\boldsymbol{q} = (x, y)$ denotes the vehicle's state, $\boldsymbol{u}$ denotes the control input, $\boldsymbol{F}(\boldsymbol{q}) = [F_1(\boldsymbol{q}), F_2(\boldsymbol{q})]^T$ is a 2D planar vector field that describes the surrounding fluid dynamics, and $\boldsymbol{\eta}$ is a zero-mean δ-correlated white noise term with standard deviation $\sigma = \sqrt{2D}$ for a given noise intensity D. Here $\boldsymbol{G}(\boldsymbol{q}(t))$ represents a multiplicative noise array that accounts for spatial effects.

In this work $\boldsymbol{G}(\boldsymbol{q})\boldsymbol{\eta}$ represents the inherent noise in the environment but can also reflect a vehicle's measurement and/or actuation uncertainty and any vehicle and/or flow field model uncertainties. For simplicity of exposition, we will assume $\boldsymbol{G}(\boldsymbol{q}(t)) = \boldsymbol{I}$, where $\boldsymbol{I}$ is the identity matrix. That is, we will consider the additive noise case in this paper.

Let $\mathscr{W}$ denote the connected obstacle-free workspace. In this work, we assume $\mathscr{W} = \cup_{i=1}^{M} G_i$ where G_i denotes a convex Lagrangian coherent structure (LCS) bounded region. In 2D flows, LCS are one-dimensional (1D) boundaries that exhibit maximum Finite-Time Lyapunov Exponents (FTLE) measures (Shadden et al., 2005; Haller, 2011). Figure 1 shows an example 2D flow field and its corresponding FTLE field. In the time-independent case, LCS correspond to stable and unstable manifolds of saddle points in the system where the manifolds can also be characterized by maximum FTLE ridges[1]. A tessellation of the workspace along LCS, or boundaries characterized by maximum FTLE ridges, makes sense since LCS: 1) delineate dynamically distinct regions in a 2D flow field, and 2) consequently have been shown to correspond to regions in the flow field where more escape events occur (Forgoston et al., 2011). Therefore, by definition a passive particle's escape from G_i, or switching between two adjacent G_i's is influenced both by measurement uncertainty in the flow field as well as external noise.

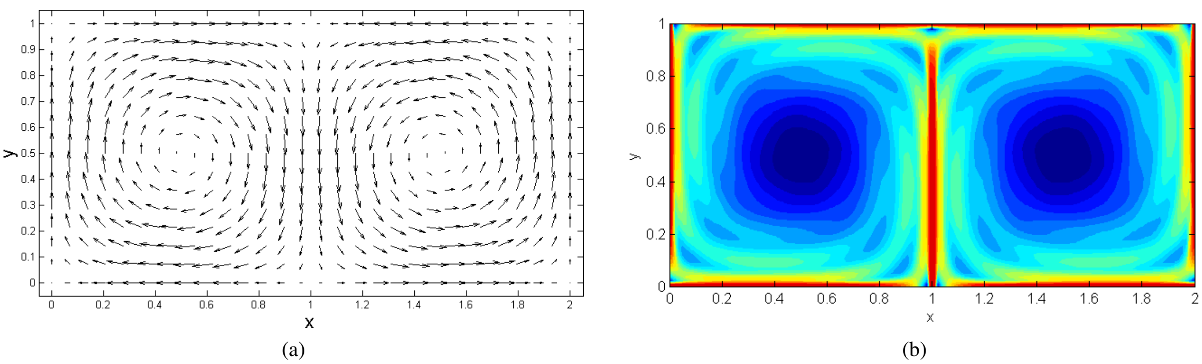

Fig. 1. Phase portrait (a) and corresponding FTLE field (b) of the model given by Eq. (4) with $A = 1$, $\mu = 0$, and $s = 1$. At each of the corners of the regions there exist saddle equilibria, and in the centers of the regions are weakly attracting equilibria.

The objective is to synthesize a control strategy that enables a AUV or robot to navigate between physically adjacent G_i's by leveraging the surrounding fluid dynamics and the inherent environmental noise to minimize the vehicle's control effort. Specifically, we are interested in control strategies that can be implemented on mobile sensors with minimal actuation capabilities. In this work, we assume the vehicle has the ability to localize within $\mathscr{W}$ and the ability to measure the flow velocity at its current position, *i.e.*, $\boldsymbol{F}(\boldsymbol{q})$ at $\boldsymbol{q}$. Additionally, we assume the tessellation of $\mathscr{W}$ into the individual G_i is given and the boundary locations of each G_i are known. As such, based on the LCS-bounded cell decomposition of $\mathscr{W}$, let $\mathscr{G} = (\mathscr{V}, \mathscr{E})$ define an directed graph with vertex set $\mathscr{V} = \{V_1, \ldots, V_M\}$ such that every V_i represents a corresponding G_i. An edge e_{ij} between V_i and V_j exists in the set $\mathscr{E}$ if the cells G_i, G_j are adjacent, *i.e.*, share a common LCS boundary, and escapes from V_i to V_j are possible.

3. Methodology

Given an AUV with kinematics described by Eq. (1), to better understand the impact of the surrounding environmental dynamics on the vehicle's ability to transition from one LCS bounded region to another, we first summarize the analysis of a vehicle's escape trajectories in the absence of controls from an LCS bounded region. The analysis will then motivate the proposed vehicle controller. The use of Eq. (1) as a model for AUV kinematics is justified when the flow is strong enough to consider the inertial term of the vehicle negligible. While the theory is general enough to handle vehicle inertial effects

[1]The FTLE are computed based on a backward (attracting structures) or forward (repelling structures) integration in time.

and non-Gaussian noise sources, we limit the following discussion to negligible vehicle inertia and Gaussian noise for the sake of simplicity.

3.1. Most Probable Escape Paths in Uncertain Flows

Under the influence of noise, the dynamical behavior of the system is determined by its stationary probability density. In particular, all equilibria are now peaks or troughs in a probability landscape describing where a particle is likely to be located. Although there exist many paths that transport a particle from one LCS bounded region to another, in the presence of small noise in the sense of (Feynman and Hibbs, 1965) there are "most likely transition paths" in the sense that they will lie along a local peak in the probability density. The fact that there exist most likely transition paths is not new, and in fact is a frequently-used result in statistical mechanics. However, our application of the idea to formally predicting such events from a controls perspective is unique. We characterize this path to transition from one LCS bounded region to another by considering the most likely paths between the two states. We briefly summarize the analysis for the sake of completeness and refer the interested reader to (Heckman et al., 2014) for the specific details.

The probability of escape from an attractor under the influence of small white noise in the sense of Schuss and Spivak (1998) scales exponentially (Feynman and Hibbs, 1965) as the action

$$\mathscr{P}_{\boldsymbol{\eta}}(\boldsymbol{q}) \asymp \exp(-\mathscr{R}(\boldsymbol{q}))/D, \tag{2}$$

where $\mathscr{R}$ represents the dynamical quantity known as the *action*. If we continue the thinking of (Feynman and Hibbs, 1965) and assume the noise is white, then for any realization of noise we get the action defined in (Freidlin and Wentzell, 1984) for Gaussian noise. We take a general Hamiltonian approach which also may be extended to escape induced by any non-Gaussian noise (Schwartz et al., 2009).

We characterize the paths that require the minimum action to cause the transition. It is important to note, that the term "action" here is taken in the mechanics sense. It is a dynamical quantity motivated by the Hamilton-Jacobi equations which, when stationary, results in Hamilton's equations of motion. Following this formulation, the background flow field $\boldsymbol{F}(\boldsymbol{q})$ in Eq. (1) now becomes a constraint. The action functional for the noise is given by (Dykman, 1990):

$$\mathscr{R}[\boldsymbol{q},\boldsymbol{\eta},\boldsymbol{\lambda}] = \frac{1}{2}\int \boldsymbol{\eta}(t)\cdot\boldsymbol{\eta}(t)dt + \int \boldsymbol{\lambda}(t)\cdot(\dot{\boldsymbol{q}}(t) - \boldsymbol{u}(t) - \boldsymbol{F}(\boldsymbol{q}(t)) - \boldsymbol{\eta}(t))dt, \tag{3}$$

where $\boldsymbol{\lambda}(t)$ is the time dependent Lagrange multiplier, and is the conjugate variable to the particle's path $\boldsymbol{q}(t)$. We first assume the control is turned off; *i.e.* $\boldsymbol{u} \equiv \boldsymbol{0}$. The action in Eq. (3) results in a scalar when evaluated, and the bounds of the integral are over the times at the end points of the trajectory $\boldsymbol{q}(t)$. When evaluating the action in Eq. (3), we compute $R = \min \mathscr{R}[\boldsymbol{q},\boldsymbol{\eta},\boldsymbol{\lambda}]$, where the minimum is taken over the functions $[\boldsymbol{q},\boldsymbol{\eta},\boldsymbol{\lambda}]$. Following the procedure for determining the equations of motion using Hamilton's principle, we set the first variation of the functional in Eq. (3) to zero. This yields a system of differential equations that identify which solutions extremize the action in terms of the path and the noise function necessary to realize the path. The solution in $\boldsymbol{q}$ is the most probable path. The resulting $\boldsymbol{\eta}$ is the "optimal" noise function that would realize such a transition. The function $\boldsymbol{\lambda}$ has is well-known in the optimal control theory for systems without noise, in particular in linear quadratic optimal control problems where it determines the optimal controller.

It is important to note that there exist many paths which may cause escape. When examining the final trajectory of those path realizations that escape a gyre, there is a well-defined maximum of a probability density that is global in phase space. The theory presented here explains this global structure, and the above variational procedure is analogous to the deterministic optimal trajectory generation problem where Eq. (3) is the objective function with $\boldsymbol{\eta}(t)$ as a control input. However, the fundamental difference in this work is that the problem is one of minimizing the exponent of a functional probability distribution of paths, which is fully stochastic. The end result yields a path and a noise function that determine

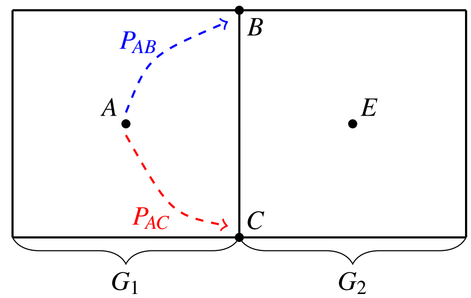

Fig. 2. Arrangement of the two-gyre system with diagram of boundary points for the integral in Eq. (3). P_{AB} and P_{AC} represent the probabilities of observing a realization of noise $\boldsymbol{\eta}(t)$ that brings a particle from A to B or C respectively. A and E represent weakly attracting foci, and B,C are boundary saddles. These probabilities may be calculated using knowledge of the vector field, noise statistics, and the theory of large fluctuations.

the most probable paths going from one part of phase space to another. It is important to note that this is not the minimum noise path, rather it is the *most likely* path in the presence of small noise.

In this work, we model the external flow field $\boldsymbol{F}(\boldsymbol{q})$ using the wind-driven double gyre flow model which is often used to describe large scale recirculation in the ocean (Veronis, 1966) and is given by:

$$\begin{aligned} F_1 &= -\pi A \sin(\pi x)\cos(\pi y/s) - \mu x \equiv f_1 - \mu x, \\ F_2 &= \pi A \cos(\pi x)\sin(\pi y/s) - \mu y \equiv f_2 - \mu y, \end{aligned} \tag{4}$$

where f_i represent the conservative part of the vector field. The parameter μ is a damping coefficient, s is a scaling dimension for the gyres, and A corresponds to the strength of the gyre flow. Figure 1(a) shows the phase portrait of Eq. (4) and Figure 1(b) shows the corresponding FTLE ridges. In this model, the workspace consists of a grid of gyres where each gyre defines an LCS-bounded region. As shown in Figure 1(a), flows in adjacent gyres circulate in opposite directions. Each gyre has an attractor located in the center of G_i and, since the system is time-invariant, the boundaries of each G_i consist of the stable and unstable manifolds of the saddle points located at the four corners of each G_i.

To fix ideas of the general theory, we consider the example of a gyre-driven flow with dynamics given by Eq. (4) in a 1×2 arrangement of G_i as in Figure 2. In the small noise limit, most likely paths connect the center of G_i to its boundary saddles. In considering the transition from $G_1 \to G_2$ there are two "most likely paths" that can lead to such a transition; the first leads a particle from $A \to B$ and the second from $A \to C$. To "transition" is to leave a neighborhood of the center of one region G_i and transit to the neighborhood of the center of another region G_j. The time to approach the stable equilibrium E is negligible once in its basin of attraction, and the return path is exponentially small in probability. In the case of the $A \to B$ path, upon arriving at B the particle has an equal chance of either transitioning to G_2 or returning to G_1 due to the occurrence of a random event at the boundary saddle, where the deterministic drift is weak; the same is true for the $A \to C$ path. Therefore, to transition from $G_1 \to G_2$, the probability is $P_{12} = \frac{P_{AB}}{2} + \frac{P_{AC}}{2}$. Since these paths are equally likely, we set $P_{AB} = P_{AC} = \alpha$ and $P_{12} = \alpha$. This is the probability of witnessing one of the two most likely paths that direct a particle out of G_1 and may be calculated using Eq. (2).

The mean switching time is the average amount of time necessary before a transition occurs; in this case, it is:

$$T_S = \underbrace{b \exp\left(\frac{R}{D}\right)}_{A\to(B\text{ or }C)} \tag{5}$$

where b is a prefactor determined through numerical simulation or experiment. In order to calculate the switching time in Eq. (5), we must find the most probable trajectory and noise in the sense of Eq. (3) to induce a transition. Computing the variational derivatives of the action and setting them to zero in the case of Eq. (4) results in the following differential equations:

$$\dot{x} = F_1(x,y) + \lambda_1 \tag{6}$$

$$\dot{y} = F_2(x,y) + \lambda_2 \tag{7}$$

$$\dot{\lambda_1} = -\frac{\partial F_1}{\partial x}\lambda_1 - \frac{\partial F_2}{\partial x}\lambda_2 \tag{8}$$

$$\dot{\lambda_2} = -\frac{\partial F_1}{\partial y}\lambda_1 - \frac{\partial F_2}{\partial y}\lambda_2 \tag{9}$$

where $\lambda_i, i = 1,2$ represent the conjugate momenta to x, y. Note that this is a conservative system which may be derived from a Hamiltonian given by

$$\mathscr{H}(\boldsymbol{q},\boldsymbol{\lambda}) = \frac{||\boldsymbol{\lambda}||^2}{2} + \boldsymbol{\lambda}\cdot\boldsymbol{F}(\boldsymbol{q}). \tag{10}$$

The *most probable* path to transition between adjacent LCS-bounded regions is determined by solving the system of Eqs. (6)–(9) subject to boundary conditions describing the equilibria in G_i. Let the coordinate of A be given by $\boldsymbol{q}_A$ and, selecting B as the boundary point of interest without loss of generality, let $\boldsymbol{q}_B$ denote the coordinate of B. Then the boundary conditions for the most probable path are, for $\boldsymbol{F}(\boldsymbol{q}_A) = \boldsymbol{F}(\boldsymbol{q}_B) = \boldsymbol{0}$, $\boldsymbol{q}(t \to -\infty) = \boldsymbol{q}_A$, $\boldsymbol{q}(t \to \infty) = \boldsymbol{q}_B$, and $\boldsymbol{\lambda}(t \to -\infty) = \boldsymbol{\lambda}(t \to \infty) = \boldsymbol{0}$. Determining the most probable path connecting $\boldsymbol{q}_A$ and $\boldsymbol{q}_B$ is equivalent to solving a two point boundary value problem in four dimensions.

It is important to note that in the case of infinitesimal noise, escape paths must pass through saddle points (B or C) because passage along the boundary does not satisfy Hamilton's principle. In the case of non-infinitesimal noise this is not the case, but if the noise is small enough then we use the infinitesimal noise case as a first order solution off which we perturb (Schuss and Spivak, 1998). The point A is a stable equilibrium in the original formulation, which means with the addition of small noise, a random walk will tend to bring a particle at the edges of the basin of attraction to the center. This statistically-expected behavior has as exceptions these "rare events" corresponding to transitioning, which are on average expected after time T_S but is exponential in distribution. In the limit of small noise, we assume that the relaxation time for a particle in the flow (that is, the time required for a particle in the basin of attraction of A to settle to the neighborhood of A) is much less than the switching time.

In our analysis, we set $\boldsymbol{q}_A$ to be the gyre attractor and $\boldsymbol{q}_B$ as one of the boundary saddle through which escape occurs[2]. Since both are saddle equilibria in the full set of the equations of motion, finding the *most probable* path mathematically requires identifying a heteroclinic orbit connecting the two points. To solve for the path numerically, we implement an algorithm known as the iterative action minimizing method (IAMM) (Lindley and Schwartz, 2013). We then employ continuation using Auto's HomCont (Doedel et al., 2008) to increase the accuracy of the approximation and to study the behavior of the path for different parameter values.

Figure 3(a) shows a simulation-derived probability density of paths to escape imposed over the most likely path as predicted by the theory. This theoretical analysis suggests that it is potentially possible to devise energy efficient control

[2]In the presence of noise, the likelihood of escape for any particle in G_i is dependent on the particle's proximity to the gyre boundaries and the noise intensity. Near the boundary saddle the vector field becomes very weak, resulting in noise dominating the dynamics. This results in high instability and corresponds to high escape likelihoods in the neighborhood of the boundary saddle.

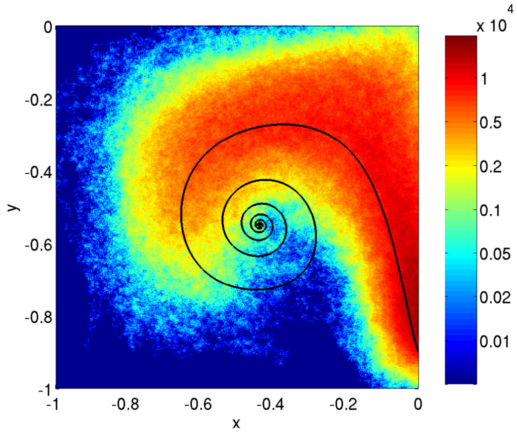

(a)

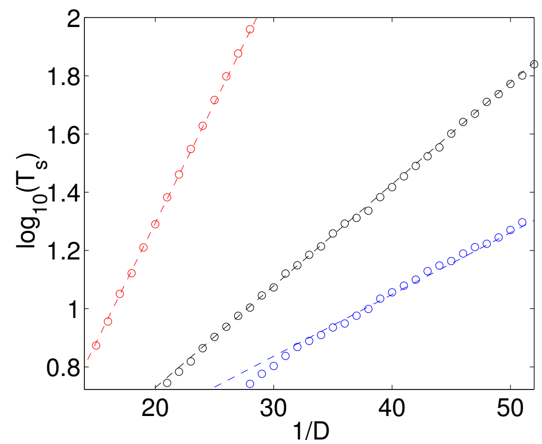

(b)

Fig. 3. (a) Comparison of the escape trajectories (color) and the most probable path to escape as calculated using theory (line) for the case of no control. Red indicates a higher incidence of occurrence while blue shows lower occurrence. The results are shown on a logarithmic scale to accentuate the ridge of higher probability. (b) Comparison of the switching time T_S as predicted by large-fluctuation theory in Eq. (5) (dashed lines) and computed from an average over many stochastic trials (points). Three regimes were examined: black represents no control, red represents $c \approx -0.31$ and blue represents $c \approx 0.1$. The parameters used to for both figures were $A = 1$, $s = 1$ and $\mu = 1$.

strategies for autonomous vehicles by leveraging the surrounding environmental. However, the described approach requires complete, although not perfect, knowledge of $\boldsymbol{F}(\boldsymbol{q})$.

In general, it is unreasonable to assume $\boldsymbol{F}(\boldsymbol{q})$ can be perfectly known for any given workspace $\mathscr{W}$, even more so by small resource constrained vehicles. Since LCS delineate regions in the flow that exhibit distinct dynamics, it suggests that it is possible to provide a reduced-order representation of $\boldsymbol{F}(\boldsymbol{q})$ using only the locations of the LCS boundaries. The locations of LCS boundaries can be approximated using historical and ocean model data. Recent work has even shown that LCS boundaries can also be tracked, potentially in real-time, by mobile robot teams using the strategy described in (Michini et al., 2014). As such, rather than assume $\boldsymbol{F}(\boldsymbol{q})$ is known, we assume $\boldsymbol{F}(\boldsymbol{q})$ can potentially be inferred from the LCS boundaries. In fact, our theoretical analysis shows that "optimal" escape paths from G_i move the particle/passive vehicle along the direction of the flow field towards the boundaries of G_i. This suggests that it is possible to achieve comparable performance using only knowledge of the LCS boundary locations, *i.e.*, partial knowledge of $\boldsymbol{F}(\boldsymbol{q})$, to devise a simple control strategy that will enable an autonomous vehicle to effectively navigate from one LCS-bounded region to another. The result would be an energy efficient control strategy that leverages the surrounding environmental dynamics without requiring explicit computation of "optimal" escape paths and subsequently full knowledge of the vector field. We present our controller synthesis in the following section.

3.2. Controller Synthesis

Consider a planar AUV with kinematics given by Eq. (1) operating in a workspace $\mathscr{W} = \cup_{i=1}^{M} G_i$ where each G_i denotes a convex region whose boundaries are defined by LCS such that the flow within each G_i is gyre-like. Consider the following controller acting in gyre G_i:

$$\boldsymbol{u} = \boldsymbol{\omega} \times c \frac{[f_1, f_2, 0]}{||[f_1, f_2, 0]||}, \tag{11}$$

where c denotes the controller gain, $[f_1, f_2, 0]$ represents the conservative part of the vector field $\boldsymbol{F}$ with an augmented third dimension, and $\boldsymbol{\omega} = [0, 0, 1]^T$. The above control strategy was first introduced in (Mallory et al., 2013a) and further analyzed in (Heckman et al., 2014). When $c > 0$ the controller effectively pushes the vehicle towards the boundary of G_i in a direction that is perpendicular to the flow at the vehicle's location. Similarly, when $c < 0$ the controller pushes the vehicle towards the center of G_i.

A significant advantage of the proposed control strategy in Eq. (11) over the computation of most likely escape paths is that it does not require full knowledge of the flow field. Furthermore, when $c > 0$, a vehicle is more likely to escape a given G_i as c increases. Figure 3(b) demonstrates this relationship by comparing the theoretically predicted (based on the optimization framework described in Section 3.1) and numerically simulated (Monte-Carlo simulations) switching time from one G_i to any adjacent G_i as a function of the noise intensity for three values of control gain c. Note that escape times *exponentially* decrease when $c > 0$ while escape times increase when $c < 0$ for different noise intensities, as predicted by Eq. (5).

These results suggest that the controller given by Eq. (11) effectively aids or abates the probability to transition between basins and affects the switching time. Furthermore, similar to the "optimal" escape paths, the proposed control strategy requires minimal control effort expenditure of the vehicle by leveraging the surrounding environmental dynamics. As such, the proposed strategy should result in more energy-efficient trajectories compared to shortest distance paths that do not account for the external flow field.

It is important to note that for robots to employ the proposed controller they must have knowledge of $\mathscr{W}$, the locations of the boundaries of G_i, and the ability to localize within each G_i. While it may be unreasonable to expect resource constrained AUVs to be able to determine the LCS locations in real-time, it is possible to compute approximate LCS boundary locations using historical and ocean model data obtained *a priori*. Furthermore, recent work has shown the possibility of tracking these boundaries using robot teams (Michini et al., 2014). As such, these assumptions similar to autonomous ground or aerial vehicle having a map of the environment.

3.3. Controller Analysis

The proposed control strategy should enable robots to effectively navigate in $\mathscr{W}$. The analysis thus far has focused on the stochastic transitions between adjacent G_i, but in the case where two regions are not adjacent, the robot will need to maneuver through multiple gyres before ultimately reaching its goal location. We begin with a statement of our theoretical result and follow an example to show the process of calculating transition probabilities for multiple transition systems.

Proposition 1 *Given the connected set $\mathscr{W}$ and its associated graph $\mathscr{G}$, V_j is reachable by any vehicle with kinematics given by Eq.* (1) *and control strategy* (11) *initially in V_i for any $i, j \in \{1, \ldots, M\}$. Furthermore, average transition times between two cells V_i, V_j may be calculated using Eq.* (5) *by applying the theory of most probable paths.*

The sketch of the reachability proof is provided in Appendix A. The result is clear to those familiar with Markov chain analysis, but we carefully define the probability concepts as they relate to the workspace $\mathscr{W}$ in order to prove the approximate switching times between individual gyres. To demonstrate how average transition times between arbitrary gyres can be calculated, consider the example of a gyre-driven flow with dynamics given by Eq. (4) in a 3×2 arrangement of G_i as in Figure 4. Specifically, consider the calculation of the expected switching time between two states that do not share a boundary saddle, *e.g.*, P_{16} in Figure 4. Following the procedure illustrated in (Lawler, 2006), we modify the transition matrix $\boldsymbol{T}$ (see Eq. (13) in Appendix A) such that G_6 is an absorbing state, *i.e.* we establish an absorbing boundary to all transitions to G_6 and find the expected number of iterations of the matrix necessary until absorption. That is, initialize the matrix:

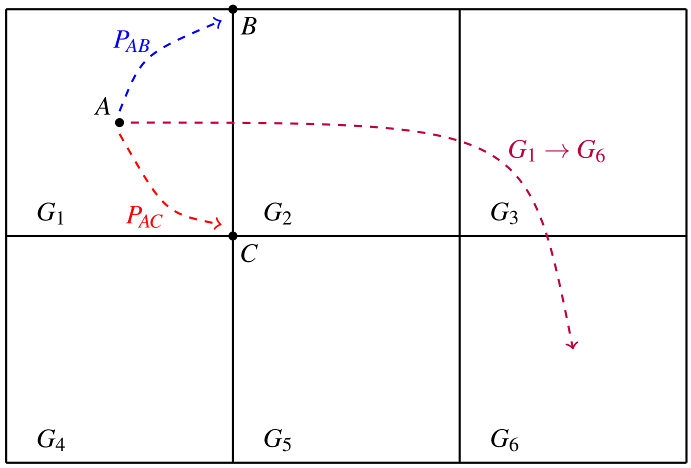

Fig. 4. Arrangement of the six-gyre system with diagram of boundary points for the integral in Eq. (3). Note that in this workspace, P_{AC} contributes to the probability of transitioning from G_1 to G_2 as well as G_4 and G_5.

$$\boldsymbol{Q} = \begin{bmatrix} 1-\frac{7\alpha}{4} & \frac{3\alpha}{4} & 0 & \frac{3\alpha}{4} & \frac{\alpha}{4} \\ \frac{3\alpha}{4} & 1-\frac{11\alpha}{4} & \frac{3\alpha}{4} & \frac{\alpha}{4} & \frac{3\alpha}{4} \\ 0 & \frac{3\alpha}{4} & 1-\frac{7\alpha}{4} & 0 & \frac{\alpha}{4} \\ \frac{3\alpha}{4} & \frac{\alpha}{4} & 0 & 1-\frac{7\alpha}{4} & \frac{3\alpha}{4} \\ \frac{\alpha}{4} & \frac{3\alpha}{4} & \frac{\alpha}{4} & \frac{3\alpha}{4} & 1-\frac{11\alpha}{4} \end{bmatrix}. \tag{12}$$

Note that $\boldsymbol{Q}$ is a substochastic matrix representing transitions between transient states. The procedure calls for calculating the matrix $\boldsymbol{M} = (\boldsymbol{I}-\boldsymbol{Q})^{-1}$ whose (j,i) entry represents the expected number of visits to G_i starting at G_j. By summing $\boldsymbol{M}^T$ over the transient states G_i assuming we will obtain the expected number of transitions to reach state G_i from state G_j. We find the expected number of transitions to reach state G_6 from G_1 to be $\frac{47}{11\alpha}$; the mean time to make this transit therefore is $\frac{47}{11\alpha} \cdot T_S$. The expected number of transitions if the particle starts in a different region is provided in Table 1.

Table 1: Expected number of transitions required to reach region G_6 from the initial region G_i. The number of transitions increases considerably if there is no shared boundary between the initial region and G_6. Note that the fractions are not reduced in order to compare the expectation value directly.

G_1	G_2	G_3	G_4	G_5
$\frac{235}{11\alpha}$	$\frac{188}{11\alpha}$	$\frac{136}{11\alpha}$	$\frac{231}{11\alpha}$	$\frac{168}{11\alpha}$

4. Experimental Methodology

To experimentally evaluate the control strategy using an experimental robotic platform, we will employ our multi-robot Coherent Structure Testbed (mCoSTe) (Larkin et al., 2014). The mCoSTe is an indoor laboratory experimental testbed that consists of several flow tanks and a fleet of micro-autonomous surface vehicles (mASVs). The mASVs are differential drive surface vehicles equipped with a micro-controller board, XBee radio module, and an inertial measurement unit (IMU). The vehicles are approximately 12 cm long and have a mass of about 45 g each (see Figure 5(b)). Localization for the mASVs is provided by an external motion capture system.

We employ the mCoSTe's Multi-Robot (MR) Tank, which is $3 \times 3 \times 1$ m^3 in size (see Figure 5(a)) in our experiments. The flows in the MR tank are patterned after the gyre flow model given by Eq. (4). We refer the interested reader to (Larkin et al., 2014) for a detailed description of the design and validation of the mCoSTe.

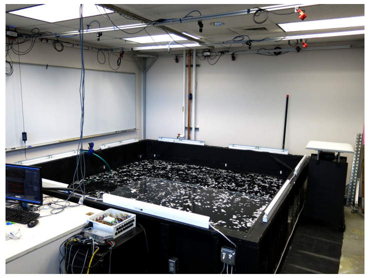

(a)

(b)

Fig. 5. Photos of (a) multi-robot (MR) tank and (b) three mASVs in the MR tank.

We first experimentally validate the exponential scaling for the noise-induced transitions between LCS-bounded regions as shown in Figure 3(b) with and without the controller given by Eq. (11). Additionally, we show how most probable escape paths as predicted by minimizing the action functional given by Eq. (3) correspond to more energy efficient trajectories out of given G_i.

4.1. Stochastic Escape Times

Using a simulated time-invariant flow field given by Eq. (4), we operated the mASVs in still water in the MR tank and examine whether the vehicles eventually leave the specified region or gyre, G_i. This setup was employed to allow for the inclusion of different amplitudes of noise in the vector field, which otherwise would be difficult to quantify. Further, establishing a steady state flow on the scale of the MR tank that would exhibit approximately the same deterministic flow behavior as Eq. (4) is a challenge in this novel testbed. The mASVs were controlled using a waypoint controller with the following waypoint-update scheme:

$$\boldsymbol{q}_{n+1} = \boldsymbol{q}_n + h(\boldsymbol{u}_n + \boldsymbol{F}(\boldsymbol{q}_n)) + \boldsymbol{\eta}_n$$

where we set $h = 0.1$ to create a sufficiently dense set of waypoints. In our experiments, the parameters in Eq. (4) were chosen to be $A = 1$, $s = 1$, $\mu = 1$ such that no new waypoint is greater than 20 cm away from the vehicle's current position. Each G_i was a single $1.2 \times 1.2\,\mathrm{m}^2$ gyre with surface flows and LCS boundaries similar to those shown Figure 1(b). We consider a stochastic flow to measure escape times for control constants $c = \{-0.5, 0, 0.5\}$. In this case $\boldsymbol{\eta}_n = \sqrt{h}\sigma\mathscr{N}(0,1)$ where $\mathscr{N}(0,1)$ represents a vector of random numbers drawn from a Gaussian distribution with mean 0 and standard deviation 1. Note that as h scales the step size, the numerical noise intensity, or diffusion, scales as $1/h$. We set the noise intensity to be $D = 0.1$ resulting in $\sigma \approx 0.45$; this results in a random distribution of waypoints around the position $\boldsymbol{q}_n$ with drifting in the direction of the gyre flow. Since the mASV is non-holonomic, its waypoint controller is formulated with this restriction in mind.

4.2. Minimum Action Paths

To determine the energy efficiency of the minimum action paths, we consider the average control effort expenditure of a vehicle executing the optimal escape trajectory predicted by the minimization of the action functional given by Eq. (3). To create a gyre-like flow in the MR tank, four rotating cylinders with diameter approximately 9 cm are placed in the MR tank at the vertices of a $1 \times 1\,\mathrm{m}^2$ square centered in the tank and set to rotate at approximately 20 Hz. The water depth is approximately 30 cm. The resulting flow in the MR tank is similar to that modeled by Eq. (4). In this environment, a single mASV is placed at a point along the calculated optimal path out of the flow and is instructed to exit the region by following two different paths. The first path is the minimum distance trajectory between the vehicle's starting position and the goal position, *i.e.*, the boundary saddle. This is the baseline case where the minimum distance trajectory is the optimal path without considering the external flow field. The second path is the minimal action path obtained by minimizing Eq. (3). The

Table 2: Comparison of escape times for the stochastic escape experiments for two different values of control and without control. $c = -0.5$ corresponds to the controller pointing inward, $c = 0$ no control and $c = +0.5$ the controller pointing outward.

c	Escape time (sec)	Number of trials
-0.5	977	52
0	363	55
$+0.5$	223	65

resulting desired trajectories are then approximated by a series of waypoints. In our experiments, the speed of the mASV along the trajectories was set such that the average time required to reach the waypoint is approximately 22 sec.

5. Experimental Results

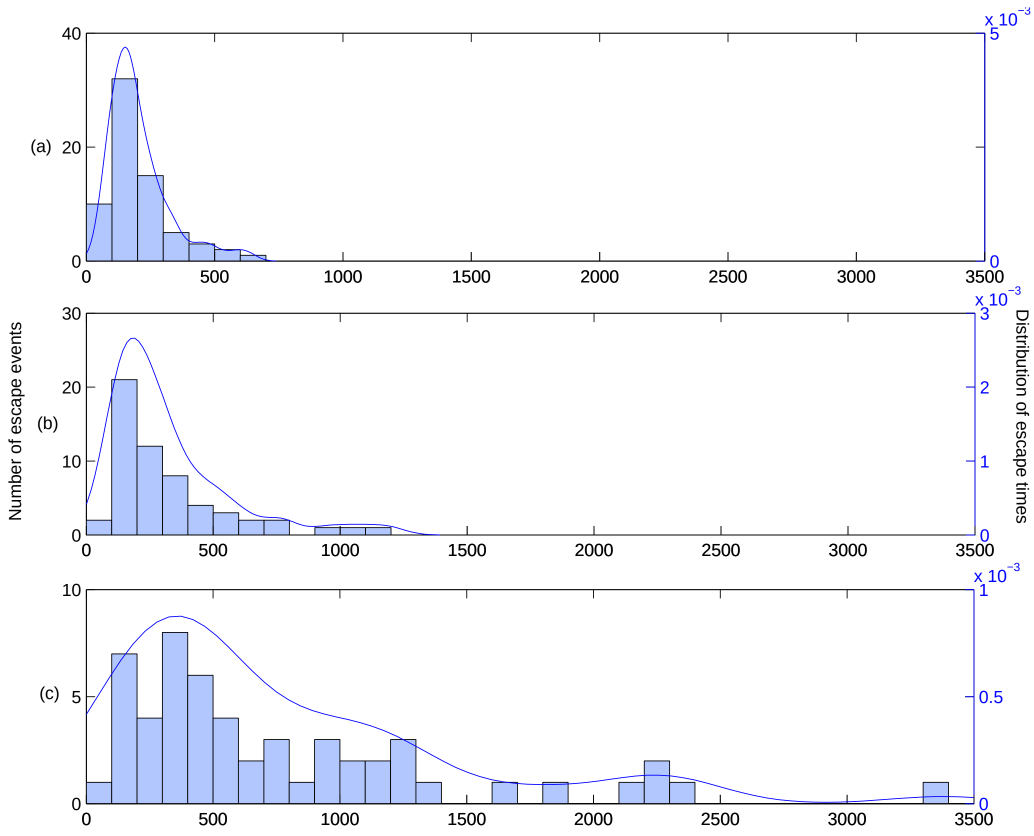

Fig. 6. Histograms of measured escape times from the simulated gyre with stochastic effects during the stochastic escape experiment, showing data for (a) $c = 0.5$, (b) $c = 0$ and (c) $c = -0.5$. The overlaid curves indicate the approximate pdf of escape times using a normal kernel for smoothening.

We begin with the results of the stochastic escape experiments. In the presence of noise, the mASV lingered in G_i for some time before stochastically escaping. We recorded the trajectories, waypoints and total time required for the mASV to leave the gyre for each of the three values of c. Figure 6 shows three histograms of measured escape times from the simulated gyre for the number of trials specified in each of the control regimes. Figure 8 shows the densities of the truncated prehistories for the mASVs as they escaped the simulated gyre for the three different gyroscopic control cases. Table 2 provides the calculated mean escape times for each value of c and the number of trials completed for each case.

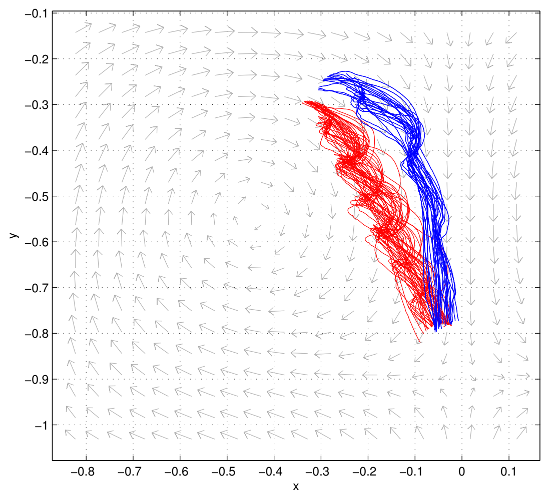

Fig. 7. The slope field of Eq. (4) overlaid with the actual paths followed by the mASV out of G_i. Blue lines correspond to the trials where the mASV was directed to follow the optimal path for $c = 0.097$ whereas red lines represent the straight-line path trials. The slope field provides only a theoretical representation of the flow lines and suggests that following the optimal path minimizes the current the mASV must overcome in order to reach its destination. The parameters used to generate the simulated flow field were $A = 1$, $s = 1$ and $\mu = 1$.

For the minimum action path experiments, we compare a series of desired trajectories and measure their performance based on control effort. The paths we compare are:

1. Straight-line path. Starting at an arbitrary point, follow waypoints along a desired path that minimizes the distance out of the current G_i.
2. Optimal path with $c = 0$.
3. Optimal path with $c = 0.097$. Changing values of c results in a homotopy of the path to escape.

Figure 7 compares the paths that the mASV followed as it exited G_i following both the straight-line path and optimal path as calculated using the flow field model in Eq. (4) with a c value that was witnessed to grant the least resistance with the imposed flow.

We use two metrics to determine the effort exerted by each vehicle. The first metric is the average RMS signal to the two motors for a given trial normalized by time-of-travel, and the second is normalized by total path length. The first metric gives a measure of the total required actuation while normalizing the data for differences in total time to execute the specified path; the second demonstrates whether the mASV leveraged the flow to reach its goal. Table 3 displays the data with 95% confidence intervals. Note that the signal to the motors ranged from $(-255, +255)$, so we use the root-mean-square of the signal to both motors to represent the instantaneous control effort. The data demonstrate that following the path obtained by solving the boundary value problem results in significantly lower control effort. The decrease in control effort can be reasonably inferred to be the result of not fighting against the flow and instead traveling along flow lines to reach the destination. Figure 7 shows that while the two paths are very similar in appearance, even a very small heading correction to follow the flow lines results in significantly lower control effort.

Table 3: Results from the path following experiments showing two metrics for average control effort to follow the path to escape.

c	Time-normalized signal	Path length-normalized signal	Number of trials
straight-line	2271 ± 118	79287 ± 4955	42
$+0.097$	1575 ± 141	54276 ± 5158	27
0	2227 ± 161	59005 ± 4418	27

6. Discussion and Experimental Insights

The escape trajectories taken by the mASV in the stochastic escape experiments and the optimal escape trajectories of a passive particle show good agreement. Both Figure 3(a) and the left column of Figure 8 show Monte Carlo simulations of a particle exiting its initial region and demonstrate a strong preference for the particle to approach the saddle point before exiting G_i. This is also exhibited in the stochastic escape experiment, shown in the right column of Figure 8. Note that the results for stochastic escape are rotationally invariant, *i.e.*, the escape paths were the same independent of which saddle was chosen as the exit point, in the double-gyre system. However as shown in Table 1, with a more complex workspace including many regions, predicting transition times between regions requires the additional step of considering non-advantageous traditions. This result suggests that a state-dependent controller that inhibits or encourages transitions has the potential to affect not only switching times out of a gyre but also bias transitions to more efficiently achieve a global spatial distribution of agents.

While the theory in this work has been established for a time-invariant vector field in which stable and unstable manifolds coincide with LCS, we have applied the tools to a real system in which such assumptions do not necessarily apply. While the flows created in the MR tank are mostly time-invariant, small variations over time exist nevertheless due to difficulties in maintaining ambient experimental parameters such as temperature and air flow at such physical scales. As a topic of future work, we would like to extend these theoretical results to regularly time-varying flows in which we employ actual LCS rather than their time-invariant counterparts. We also plan to establish a ground truth for the flow in the tank, in particular developing a quantitative understanding of the LCS.

Furthermore, the scaling predicted in Eq. (5) and illustrated with numerical comparison in Figure 3(b) suggests that controllers which make use of the underlying environmental dynamics and environmental stochasticity may be rigorously developed to manipulate the switching time of a robot operating in a noisy flow. This result is confirmed experimentally in Figure 6; by employing a simple control strategy, the mASV effectively manipulated its distribution of residence times in a simulated flow (see Table 2). Importantly, the data in Figure 3 does show exponential behavior since the logarithm of the escape times fall almost entirely on a straight line. The result is only predictive up to a prefactor in Eq. (5). This prefactor may be calculated in simulation and is generally approximated by a constant. While the explicit dependence of the mean transition time between gyres on the parameter c is difficult to derive analytically for most nonlinear vector fields. However, one may interpolate between c values for which escape data can be empirically estimated to predict the appropriate control input to obtain a desired mean escape time. We are currently seeking to better understand how different values of c impact the bounds of the mean escape time under different flow conditions.

Finally, we have confirmed that the most likely path as predicted by the theory not only lies along the peak of the observed escape trajectories as shown in Figure 3(a), but also requires less control effort than following a straight-line path to exit the initial G_i. The last column in Table 3 demonstrates that the total control effort by the mASV following the optimal path is considerably lower than the control effort exerted by the vehicle executing the straight-line path. This is not surprising since the mASV travels with the flow and is pushed along by the current and only actuates when necessary to correct course. More interestingly is that the proposed control strategy achieves comparable or slightly better performance from a control effort expenditure perspective. This is significant since the proposed controller can be implemented with only limited knowledge of the flow field.

This work demonstrates that the theory of large fluctuations applies to macroscopic systems of autonomous vehicles operating in noisy environments. The concepts of most likely paths and rare events explained in our methodology were originally developed for microscopic systems undergoing switching due to thermal fluctuations. However, our experiments suggest that a macroscopic agent operating in a massive stochastic environment is subject to the same phenomena as its microscopic counterparts. This has implications for the design of controllers in environments where the structure of the flow is not necessarily known, significant stochastic fluctuations are present, and control effort is costly.

7. Future Work

To the authors' knowledge, this is one of the first attempts to employ knowledge of coherent structures in designing more energy efficient navigation strategies for autonomous vehicles operating in uncertain fluidic environments. The main contribution lies in the synthesis of minimal control effort strategies using knowledge of *most likely* escape paths arising from noise-induced large fluctuations. The result is a strategy that leverages the surrounding environmental dynamics and the inherent environmental noise to minimize the overall control effort of the to navigate from one LCS-bounded region to another. The proposed strategy was evaluated using a novel testbed capable of creating controlled Lagrangian coherent structures in a laboratory setting. This is the first experimental study of most likely escape trajectories and minimal control effort escape trajectories for unmanned underwater/surface vehicles operating in realistic complex flows.

One future application is the extension of the proposed AUV/ASV control strategy to domains consisting of geometrically complex LCS-bounded regions. For instance, the LCS may be time-dependent which would more closely resemble ocean flows or the LCS-bounded regions may be non-convex. We would also like to revisit the assumption of additive noise drawn from a Gaussian distribution and the use of the time-invariant gyre-driven flow model. Most noise is in fact drawn from a colored distribution and arises in modeling as multiplicative noise; both of these factors have an effect on the shape of the most likely paths and the distribution of rare events. Also, the time-invariant gyre-driven flow model used in Eq. (4) follows the position of a particle in a two-dimensional vector field whereas actual field deployments require taking into consideration the third dimension and the actual size and geometry of the vehicle. We would like to extend the existing framework to properly account for these factors.

8. Acknowledgements

This research was performed while CRH held a National Research Council Research Associateship Award at the U.S. Naval Research Laboratory. IBS was supported by Office of Naval Research (ONR) Award Nos. F1ATA01098G001 and N0001412WX-20083, and by Naval Research Base Program N0001412WX30002. MAH and the mCoSTe are also supported by ONR Award Nos. N000141211019 and N0001413-10731. We would like to thank Matt Michini from the SAS Lab for his assistance in conducting the experiments.

A. Transition Probabilities Between Cells

To prove the connectivity and recurrence of the graph of gyres and the calculate probabilities of switching between them, we conduct an analysis which is based on discrete-time Markov chain models, provided for the sake of completeness and following a standard analysis of such chains. The region in which the agent is located corresponds to states in the chain, and the likelihood of switching between regions within a fixed time T_S is the transition probability. Recall T_S represented the mean time required to switch between two adjacent gyres.

First consider the case where the two regions V_i, V_j are adjacent and possibly also adjacent to a leave state. The graph for this workspace is $\mathscr{G} = (\{V_i, V_j, V_L\}, \{e_{ij}, e_{i,\text{leave}}, e_{j,\text{leave}}\})$. The dynamics on the graph are described by a discrete Markov chain in which the states are vertices on the graph and transitions between the states are along edges and occur according to the theory of rare events described in Section 3.1 because deterministically there are no transitions. The Markov chain dictates that there is a nonzero probability of transitioning along edge e_{12} to the adjacent region.

Now consider the case where V_i and V_j are not adjacent. Since all vertices in the graph are adjacent to at least one other vertex, then by induction the graph must be traversable to a node adjacent to V_j, *i.e.* there exists a path between any one state and another. The probability of following any particular path may be calculated by multiplying the transition probabilities between states in the Markov chain. This probability is certainly nonzero since all transition probabilities between adjacent regions are nonzero.

Moreover, the controller in Eq. (11) has already been shown to directly manipulate the transition times via directly modifying the action functional in Eq. (3). This proves that not only are all states reachable, but that we can formally evaluate the expected time required to reach a given state and increase or decrease that time with the controller in (11).

In this work, we assume $\mathscr{W}$ is composed of a union of convex LCS bounded regions denoted by G_i. The algorithm asymmetrically extends to more general convex environments by calculating the action R_{AB}, R_{AC}, etc. along heteroclinic connections that exist between the deterministically stable equilibrium and saddles along the basin boundary. The mean time to observe a given trajectory will then also be given by Eq. (5). In general, the convexity assumption for each G_i can be relaxed to accommodate star convex geometries. However, more topological insight is needed to extend these results to general non-convex geometries and their impact on transition likelihoods between neighboring cells to ensure that vehicles are not trapped in an infinite sequence of transitions between two adjacent LCS bounded regions.

Consider the example of a gyre-driven flow with dynamics given by Eq. (4) in a 3×2 arrangement of G_i as in Figure 4; we choose the boundaries to be reflecting so there is no absorbing state, but the process can also be considered for the case where a particle may leave the workspace.

We have already described the probabilistically-preferred paths in Section 3.1. In the case of the $A \to C$ path, upon arriving at C the particle has an equal chance of either transitioning to G_2 or returning to G_1 as before. In the $A \to D$ case however, upon arriving at D the particle is equally likely to transition into G_2, G_4 or G_5 or return to G_1. Therefore, to transition from $G_1 \to G_2$, the probability is $P_{12} = \frac{P_{AC}}{2} + \frac{P_{AD}}{4}$. Since the paths are equally likely, then $P_{12} = \frac{3\alpha}{4}$.

The transition matrix for the system in Figure 4 may be constructed using the transition rates in Figure 9 and therefore used to ascertain any arbitrary probabilities for transitioning between regions, *e.g.* the transition from $G_1 \to G_6$ as illustrated. It may also be used to show that all states of the chain are recurrent, and therefore any region is accessible from another arbitrary region within a certain average length of time. The transition matrix entries are given by P_{ij}, where i is the initial region and j is the region to which a switch may occur with probability P_{ij}. In consideration of the Markov diagram in Figure 9, the matrix is:

$$\boldsymbol{T} = \begin{bmatrix} 1-\frac{7\alpha}{4} & \frac{3\alpha}{4} & 0 & \frac{3\alpha}{4} & \frac{\alpha}{4} & 0 \\ \frac{3\alpha}{4} & 1-\frac{11\alpha}{4} & \frac{3\alpha}{4} & \frac{\alpha}{4} & \frac{3\alpha}{4} & \frac{\alpha}{4} \\ 0 & \frac{3\alpha}{4} & 1-\frac{7\alpha}{4} & 0 & \frac{\alpha}{4} & \frac{3\alpha}{4} \\ \frac{3\alpha}{4} & \frac{\alpha}{4} & 0 & 1-\frac{7\alpha}{4} & \frac{3\alpha}{4} & 0 \\ \frac{\alpha}{4} & \frac{3\alpha}{4} & \frac{\alpha}{4} & \frac{3\alpha}{4} & 1-\frac{11\alpha}{4} & \frac{3\alpha}{4} \\ 0 & \frac{\alpha}{4} & \frac{3\alpha}{4} & 0 & \frac{3\alpha}{4} & 1-\frac{7\alpha}{4} \end{bmatrix} \tag{13}$$

For the Markov chain to be irreducible, the matrix $\boldsymbol{T}$ must have a simple eigenvalue of unity whose left eigenvector has only positive entries. In fact $\boldsymbol{T}$ satisfies even stronger conditions and represents an *ergodic* Markov chain because it is positive recurrent and all states are aperiodic; however, the chain's irreducibility and lack of absorbing states proves that all states in the chain are accessible from any arbitrary state.

References

L. Billings and I. B. Schwartz. Identifying almost invariant sets in stochastic dynamical systems. *Chaos*, 18:023122, 2008.

L. Billings, M. I. Dykman, and I. B. Schwartz. Thermally activated switching in the presence of non-gaussian noise. *Phys. Rev. E*, 78 (5):051122, November 2008.

E. M. Bollt, L. Billings, and I. B. Schwartz. A manifold independent approach to understanding transport in stochastic dynamical systems. *Physica D*, 173:153–177, 2002.

D. Caron, B. Stauffer, S. Moorthi, A. Singh, M. Batalin, E. Graham, M. Hansen, W. Kaiser, J. Das, A. de Menezes Pereira, B. Zhang A. Dhariwal, C. Oberg, and G. Sukhatme. Macro- to fine-scale spatial and temporal distributions and dynamics of phytoplankton and their environmental driving forces in a small subalpine lake in southern California, USA. *Limnol. Oceanogr.*, 53(5):2333–2349, 2008.

V. Chen, M. Batalin, W. Kaiser, and G. Sukhatme. Towards spatial and semantic mapping in aquatic environments. In *IEEE Int. Conf. Robot. Automat.*, pages 629–636, Pasadena, CA, 2008.

J. Das, F. Py, T. Maughan, T. O'Reilly, M. Messié, J. Ryan, G. Sukhatme, and K. Rajan. Coordinated sampling of dynamic oceanographic features with underwater vehicles and drifters. *Int. J. Robot. Res.*, 31(5):626–646, 2012.

E.A. D'Asaro and et al. A lagrangian float. *Journal of Atmospheric and Oceanic Technology*, 13:1230–1246, 1996.

Eric D'Asaro, Craig Lee, Luc Rainville, Ramsey Harcourt, and Leif Thomas. Enhanced turbulence and energy dissipation at ocean fronts. *Science*, 332(6027):318–322, 2011. . URL `http://www.sciencemag.org/content/332/6027/318.abstract`.

Eusebius J. Doedel, Alan R. Champneys, Fabio Dercole, Thomas Fairgrieve, Yuri Kuznetsov, Bart Oldeman, Randy Paffenroth, Björn Sandstede, Xianjun Wang, and Chenghai Zhang. *AUTO-07P: Continuation and Bifurcation Software for Ordinary Differential Equations*, February 2008. URL `http://sourceforge.net/projects/auto-07p/`.

M. I. Dykman. Large fluctuations and fluctuational transitions in systems driven by coloured gaussian noise: A high-frequency noise. *Phys. Rev. A*, 42:2020–2029, 1990.

R. P. Feynman and A. R. Hibbs. *Quantum Mechanics and Path Integrals*. McGraw-Hill, Inc., 1965.

Eric Forgoston, Lora Billings, Phillip Yecko, and Ira B. Schwartz. Set-based corral control in stochastic dynamical systems: Making almost invariant sets more invariant. *Chaos*, 21(013116), 2011.

M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems*. Springer-Verlag, 1984.

G. Froyland and K. Padberg. Almost-invariant sets and invariant manifolds – connecting probabilistic and geometric descriptions of coherent structures in flows. *Physica D*, 238:1507–1523, 2009.

G. Haller. A variational theory of hyperbolic lagrangian coherent structures. *Physica D*, 240(7):574–598, 2011.

G. Haller and G. Yuan. Lagrangian coherent structures and mixing in two-dimensional turbulence. *Phys. D*, 147:352–370, December 2000. ISSN 0167-2789. . URL `http://dl.acm.org/citation.cfm?id=366463.366505`.

Christoffer R. Heckman, M. Ani Hsieh, and Ira B. Schwartz. Going with the flow: Enhancing switching rates in multi-gyre systems. Accepted to the ASME Journal of Dynamic Systems, Measurement and Control, 2014.

M.A. Hsieh, E. Forgoston, T.W. Mather, and I.B. Schwartz. Robotic manifold tracking of coherent structures in flows. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 4242–4247, May. .

T. Inanc, S.C. Shadden, and J.E. Marsden. Optimal trajectory generation in ocean flows. In *American Control Conference, 2005. Proceedings of the 2005*, pages 674 – 679, 8-10, 2005. .

D. Larkin, M. Michini, A. Abad, S. Teleski, and M. A. Hsieh. Design of the multi-robot coherent structure testbed (mcoste) for distributed tracking of geophysical fluid dynamics. In *Submitted to ASME International Design Engineering Technical Conferences (IDETC)*, Buffalo, NY USA, Aug 2014.

Gregory F. Lawler. *Introduction to Stochastic Processes*. Chapamn & Hall, 2006.

P.F.J. Lermusiaux. Uncertainty estimation and prediction for interdisciplinary ocean dynamics. *J. Comp. Phys.*, 217:176–199, 2006. .

P.F.J. Lermusiaux and F. Lekien. Dynamics and lagrangian coherent structures in the ocean and their uncertainty. In *Dynamical System Methods in Fluid Dynamics, Oberwolfach Workshop*. Mathematisches Forschungsinstitut Oberwolfach, 2005.

P.F.J. Lermusiaux, C.-S. Chiu, G.G. Gawarkiewicz, P. Abbot, A.R. Robinson, R.N. Miller, P.J. Haley, W.G. Leslie, S.J. Majumdar, A. Pang, and F. Lekien. Quantifying uncertainties in ocean predictions. *Oceanography*, 19:92–105, 2006. .

Brandon S. Lindley and Ira B. Schwartz. An iterative action minimizing method for computing optimal paths in stochastic dynamical systems. *Physica D*, 255:22–30, 2013.

T. Lolla, M. P. Ueckermann, P. Haley, and P. F. J. Lermusiaux. Path planning in time dependent flow fields using level set methods. In *Submitted to IEEE International Conference on Robotics and Automation*, Minneapolis, MN, 2012.

K. M. Lynch, P. Schwartz, I. B. Yang, and R. A. Freeman. Decentralized environmental modeling by mobile sensor networks. *IEEE Trans. Robotics*, 24(3):710–724, 2008.

K. Mallory, M. A. Hsieh, E. Forgoston, and I. B. Schwartz. Distributed allocation of mobile sensing swarms in gyre flows. *Nonlin. Processes Geophys.*, 20:657–668, 2013a.

K. Mallory, M. A. Hsieh, E. Forgoston, and I. B. Schwartz. Distributed allocation of mobile sensing swarms in gyre flows. *Nonlin. Processes Geophys.*, 20(5):657–668, 2013b.

M. Michini, M. A. Hsieh, E. Forgoston, and I. B. Schwartz. Robotic tracking of coherent structures in flows. *IEEE Trans. on Robotics*, 2014.

Michel K Ochi. *Ocean waves : the stochastic approach*. Cambridge, U.K. ; New York : Cambridge University Press, 1998. ISBN 9780521017671 (pbk.). Includes bibliographical references and index.

A. Provenzale. Transport by coherent barotropic vortices. *Ann Rev Fluid Mech*, 31:55–93, 1999.

Z. Schuss and A. Spivak. Where is the exit point? *J. Chem. Phys.*, 235:227–242, 1998.

I. B. Schwartz, L. Billings, M. Dykman, and A. Landsman. Predicting extinction rates in stochastic epidemic models. *J. Stat. Mech.-Theory E.*, page P01005, 2009.

Ira B. Schwartz, Eric Forgoston, Simone Bianco, and Leah B. Shaw. Evolving towards the optimal path to extinction. *Submitted*, 2010.

Ira B Schwartz, Eric Forgoston, Simone Bianco, and Leah B Shaw. Converging towards the optimal path to extinction. *Journal of The Royal Society Interface*, 8(65):1699–1707, 2011.

SCRIPPS. Naitonal HF RADAR network - surface currents, 2014. URL `http://cordc.ucsd.edu/projects/mapping/maps/`.

C. Senatore and S.D. Ross. Fuel-efficient navigation in complex flows. In *American Control Conference, 2008*, pages 1244 –1248, june 2008. .

Shawn C. Shadden, Francois Lekien, and Jerrold E. Marsden. Definition and properties of lagrangian coherent structures from finite-time lyapunov exponents in two-dimensional aperiodicï¿½flows. *Physica D: Nonlinear Phenomena*, 212(3-4):271 – 304, 2005. ISSN 0167-2789. . URL `http://www.sciencedirect.com/science/article/pii/S0167278905004446`.

A. Shchepetkin and J. McWilliams. Quasi-monotone advection schemes based on explicit locally adaptive dissipation. *Monthly Weather Review*, 126:1541–1580, 1998.

A. F. Shchepetkin and J. C. McWilliams. The regional oceanic modeling system (roms): a split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Modeling*, 9:347–404, 2005.

R. Smith, A. M. Pereira, Y. Chao, P.P. Li, D. A. Caron, B. H. Jones, and G. Sukhatme. Autonomous underwater vehicle trajectory design coupled with predictive ocean models: A case study. In *Proc. of the 2010 IEEE International Conference on Robotics and Automation*, pages 4770–4777, Anchorage, AK, 2010a.

R. Smith, J. Kelly, and G. Sukhatme. Towards improving mission execution for autonomous gliders with an ocean model and kalman filter. In *Proc. of the 2012 IEEE International Conference on Robotics and Automation*, Minneapolis, MN, 2012.

Ryan N. Smith, Yi Chao, Peggy P. Li, David A. Caron, Burton H. Jones, and Gaurav S. Sukhatme. Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model. *International Journal of Robotics Research*, 29(12):1475–1497, 2010b. .

N. Sydney and D. A. Paley. Multi-vehicle control and optimization for spatiotemporal sampling. In *IEEE Conf. Decis. Contr.*, pages 5607–5612, Orlando, FL, 2011.

Karsten Trulsen. Weakly nonlinear and stochastic properties of ocean wave fields. application to an extreme wave event. In John Grue and Karsten Trulsen, editors, *Waves in Geophysical Fluids*, volume 489 of *CISM International Centre for Mechanical Sciences*, pages 49–106. Springer Vienna, 2006. ISBN 978-3-211-37460-3. . URL `http://dx.doi.org/10.1007/978-3-211-69356-8_`

2.

G. Veronis. Wind-driven ocean circulation, part i and part ii. *Deep-Sea Res.*, 13(31), 1966.

D. Wang, P.F.J. Lermusiaux, P.J. Haley, D. Eickstedt, W.G. Leslie, and H. Schmidt. Acoustically focused adaptive sampling and on-board routing for marine rapid environmental assessment. *Journal of Marine Systems*, 78:393–407, 2009.

W. Wu and F. Zhang. Cooperative exploration of level surfaces of three dimensional scalar fields. *Automatica, the IFAC Journall*, 47(9): 2044–2051, 2011.

F. Zhang, D. M. Fratantoni, D. Paley, J. Lund, and N. E. Leonard. Control of coordinated patterns for ocean sampling. *International Journal of Control*, 80(7):1186–1199, 2007.

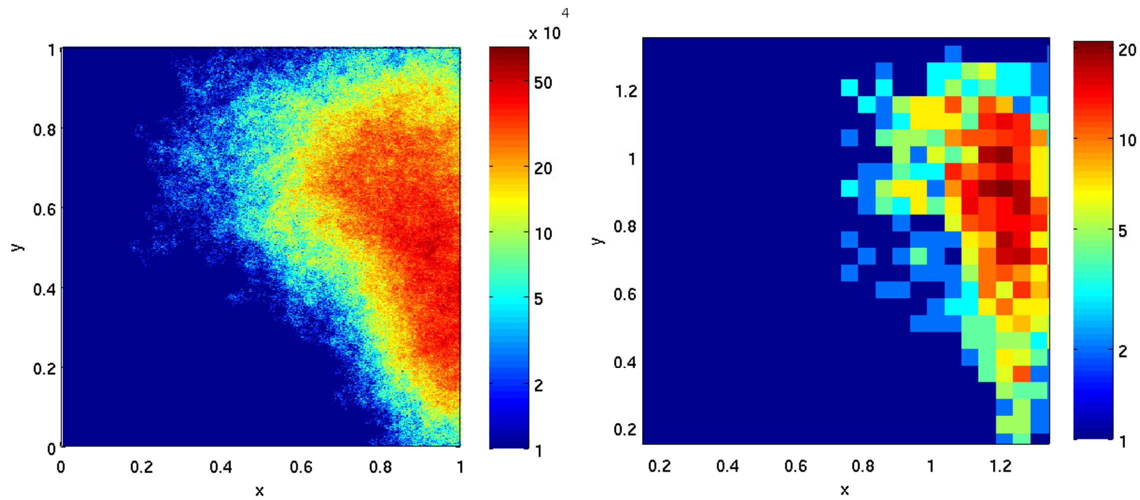

(a)

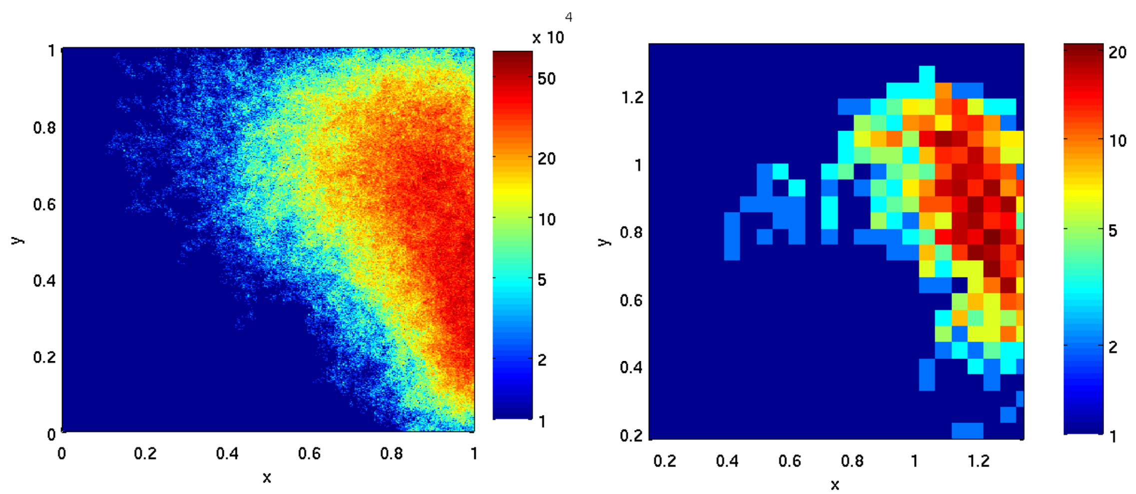

(b)

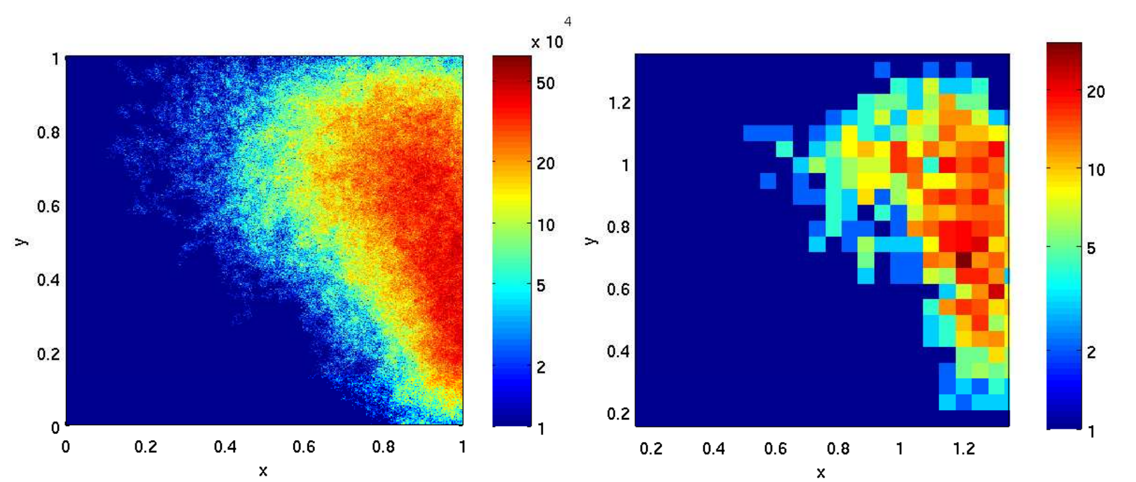

(c)

Fig. 8. Prehistories of stochastic escape from the initial region for controller gains (a) $c = -0.5$, (b) $c = 0$ and (c) $c = 0.5$. The left column shows Monte Carlo results for 1,000 trials at each controller gain, and the corresponding image in the right column shows the history of waypoints from the stochastic escape experiments for the same gain. The scaling for colors in both sets of figures is exponential and was truncated such that the approach to escape is accentuated compared with the time spent near the center of the region. The parameters used to generate the simulated flow fields were $A = 1$, $s = 1$ and $\mu = 1$.

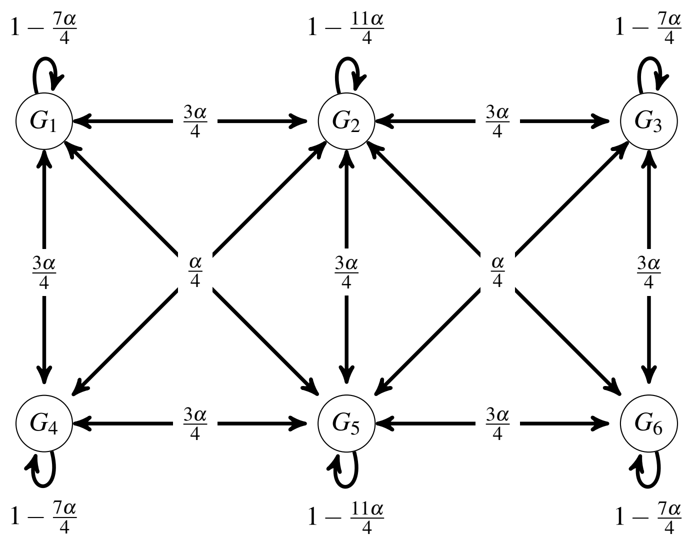

Fig. 9. A Markov diagram of the transitions between the workspace modeled in Figure 4. Since the transitions regions all have the same dynamics, transitions between them are symmetric. However, the method need not require symmetry of transition probability to construct the Markov chain.