

Clarity-Driven Ergodic Control for Persistent Tip-and-Cue Missions with Synchronized Rendezvous

David Li, Kevin M. Govindarajan, and Chris Vermillion

Abstract—This paper presents a persistent control methodology for a sustainably powered host/agent network that executes tip-and-cue oceanographic observing operations within a spatiotemporally evolving environment. Specifically, a renewably powered host vessel simultaneously serves as a recharging platform for an autonomous aerial vehicle (AAV), while also performing broad surveillance of an evolving mission domain. When the AAV is on board the host vessel, the mission trajectory (termed the “nominal” trajectory) is selected based on a clarity-driven ergodic planner. When a location of interest (termed a “tip” location) is detected by the host vessel, the AAV is dispatched to provide detailed observation of that location. This necessitates a replanning operation (of the “rendezvous” trajectory) wherein a rendezvous point is selected to maximize the mutual long-horizon benefit to the host and agent. Because the mutually beneficial rendezvous point will, in general, deviate from the original ergodic trajectory, another replanning operation (of the nominal trajectory) is completed on rendezvous. In this paper, we demonstrate the efficacy of the combination of the ergodic trajectory planner and rendezvous planner for a solar-powered host vessel (the SeaTrac SP-48 ASV) and a quadrotor (Agilicious) agent vehicle. In particular, the combined control system is shown to significantly outperform a line-transect strategy and an ergodic controller wherein the rendezvous point is constrained to lie on the nominal mission trajectory.

I. INTRODUCTION

Tip-and-cue frameworks are widely used for characterizing uncertain and evolving oceanographic and atmospheric environments [1]. These networks rely on broad observations to provide “tips” on areas of interest, after which mobile agents are “cued” for detailed observations. Traditionally, satellites have been used for detecting tip events [2]. However, occlusion due to clouds makes satellite observations impossible, and underwater observations (e.g., identification of aquatic life, subsurface sea temperature, or subsurface current velocity) are unobtainable through satellite imagery.

An alternative to the use of satellite imagery to determine tip events is to use a persistently operating autonomous surface vessel (ASV), examples of which include Saildrone [3], Liquid Robotics Wave Glider [4], and SeaTrac SP-48 solar-powered ASV [5]. ASVs can serve as mobile hosts for aerial and/or underwater agents. However, this introduces new challenges: the agent must rendezvous with the moving host

post-mission, and the host may need to adapt its trajectory to facilitate rendezvous. This creates a tradeoff between broad surveillance and detailed observations.

Several works have studied attributes of the combined host/agent control problem in isolation. For example, [6] and [7] consider the deployment of AAVs from a truck for surveillance, with the AAVs returning to the truck for recharging/battery swaps; however, these works consider task locations that are known a priori. As such, these works are limited to finite-duration missions that can be planned completely offline. In [1], ship trajectories are predicted for efficiently identifying anomalies, which can then be used for integration into satellite-based tip-and-cue services. However, successful prediction depends on the existence of a large number of high-quality datasets for training the neural network. In the context of a persistent mission with a large domain size, such datasets are frequently unavailable.

In this paper, we address the aforementioned unresolved challenges by building upon a *clarity-driven ergodic trajectory planning* approach that has been developed by our team in [8] and [9]. The concept of operations for the proposed approach is shown in Fig. 1. At the time that the host detects a tip event, a separate sub-mission planning operation is performed, which minimizes a cost function that balances the time the agent is able to service the tip event with the estimated deficit in long-term clarity that will result from the host revising its mission trajectory to achieve rendezvous with the agent upon the completion of the sub-mission. Upon completion of the sub-mission, the host/agent ergodic trajectory is re-optimized based on a clarity metric.

Our proposed mission planning/replanning strategy, termed a “collaborative ergodic” strategy, has been simulated based on the SeaTrac SP-48 ASV and the Agilicious quadrotor AAV, which are depicted in Fig. 1. The vehicles’ positional and energetic dynamics have been incorporated into the simulation using manufacturer data from [5] and [10]. In this work, we consider a persistent surveillance mission over a $2 \text{ km} \times 2 \text{ km}$ water domain. Such a mission (and associated scale) can be motivated by several applications, including long-term observation of (i) spatiotemporally varying river plumes [11], (ii) localized harmful algal blooms [12], and (iii) marine life [13]. The ASV host possesses a maximum speed of $v_{\text{host, max}} = 2.315 \text{ m/s}$, sensing radius of $r_{\text{sensing}} = 200 \text{ m}$, and battery capacity of $b_{\text{host, max}} = 6500 \text{ Wh}$. The AAV agent possesses a maximum speed of $v_{\text{agent, max}} = 30 \text{ m/s}$ and battery capacity of $b_{\text{agent, max}} = 44.4 \text{ Wh}$, and it is only able to perform detailed sensing when operating at the location of a tip event. The ultimate goals

Supported by NSF Award Number 2223844.

David Li is a Graduate Research Assistant with the Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109, USA. davidli@umich.edu

Kevin M. Govindarajan is a Graduate Research Assistant with the Department of Robotics, University of Michigan, Ann Arbor, MI 48109, USA. kmgovind@umich.edu

Chris Vermillion is an Associate Professor with the Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI, 48109, USA. cvermill@umich.edu




Fig. 1. Concept of operations for the tip and cue framework.

are to (i) identify and service as many tip events as possible while (ii) maximizing the time spent servicing tip events. We compare our framework against two benchmarks: (i) a “directive ergodic” strategy, which is forced to select rendezvous points on the original ergodic trajectory and (ii) a “directive transect” strategy, which is forced to select rendezvous points on a line-transect nominal trajectory. Both strategies place the rendezvous burden more heavily upon the agent than upon the host, which leads to a suboptimal tradeoff between the broad surveillance goals of the host and detailed observational goals of the agent.

II. PRELIMINARIES

A. Notation

Let $\mathbb{R}$ denote the set of real numbers, $\mathbb{R}_{>0}$ the positive reals, and $\mathbb{R}_{\geq 0}$ the non-negative reals. The mission domain is denoted by $\mathcal{D} \subset \mathbb{R}^2$. The variables $p \in \mathbb{R}^2$ and $b \in \mathbb{R}$ represent the position and battery state of charge, where subscripts are used to distinguish between the host, agent, and tip event (for p).

B. Tip Dynamics

A tip event occurs at location $p_{\text{tip}} \in \mathcal{D}$ over an interval $[t_1, t_2]$, where $t_2 - t_1 = \frac{1}{\lambda_{\text{tip}}}$ is the event duration. New tip events are generated every T_{tip} seconds, with locations sampled uniformly in $\mathcal{D}$. At steady state, the expected number of simultaneous tip events is $\frac{1}{T_{\text{tip}} \lambda_{\text{tip}}}$.

C. Host-Agent Network

The system comprises a solar-powered ASV (*host*) and a quadrotor AAV (*agent*). At any given time, the host detects

the present tip events within a sensing radius r_{sensing} , but only knows their locations, not start or end times. The agent detects an event only by physically reaching its location.

Upon tip detection, the rendezvous planner (Sec. III) determines whether the agent can travel to and service the event before returning to a rendezvous point with the host. If feasible, the agent gets deployed to service the tip (by hovering), returns, and is recharged by the host. If not, the host continues its nominal clarity-driven ergodic search.

A tip event is considered serviced if the agent spends any nonzero time at its location. Mission performance is measured by the total number of serviced events and total time spent servicing.

D. Vehicle Dynamics

The mission alternates between two phases:

Docked: The agent and host move together with dynamics:

$$\dot{p}_{\text{host}}(t) = u_{\text{host}}(t), \quad \|u_{\text{host}}(t)\| \leq v_{\text{host, max}}, \quad (1)$$

$$\dot{b}_{\text{host}} = P_{\text{in}} - k_1 \|u_{\text{host}}\|^3 - k_2 - \delta_{\text{charge}} k_c(t), \quad (2)$$

$$\dot{b}_{\text{agent}} = \delta_{\text{charge}} k_c(t), \quad (3)$$

where $\delta_{\text{charge}} = 1$ during charging and 0 otherwise.

Deployed: The host and agent move independently. The host position dynamics remain the same as in Eqn. (1), whereas the remaining dynamics are given by:

$$\dot{b}_{\text{host}} = P_{\text{in}} - k_1 \|u_{\text{host}}\|^3 - k_2, \quad (4)$$

$$\dot{p}_{\text{agent}} = u_{\text{agent}}, \quad \|u_{\text{agent}}\| \leq v_{\text{agent, max}}, \quad (5)$$

$$\dot{b}_{\text{agent}} = -k_3 \|u_{\text{agent}}\|^3 - k_4. \quad (6)$$

E. Clarity-Driven Ergodic Search

Clarity $q(t, p) \in [0, 1]$ quantifies the host's knowledge of tip events at location p and time t , with $q = 1$ indicating full knowledge. Initially, $q(t_0, p) = 0$ for all p . The clarity dynamics are: for all $p \in \mathcal{D}$,

$$\begin{cases} q(t, p) = 1 & \text{if } \|p - p_{\text{host}}(t)\| \leq r_{\text{sensing}}, \\ \frac{\partial q}{\partial t}(t, p) = -k_5 & \text{otherwise.} \end{cases} \quad (7)$$

The ergodic search [14], [15] generates trajectories that match a specified target information spatial distribution (TISD) ϕ_d . We define $\phi_d(p) = 1$ for uniform coverage. The planner uses the current clarity map $q(t, \cdot)$, host trajectory history, and host's current speed to plan future trajectories with constant-speed assumptions.

III. RENDEZVOUS PLANNING

In this section, we consider the problem of planning a deployment of an AAV agent to service a tip event and the subsequent rendezvous with the ASV host (i.e., planning the host's and agent's paths during a deployed phase of the mission). We refer to the deployment and rendezvous plans together simply as the *rendezvous plan*.

1) *Problem Statement*: We want to design a rendezvous plan that, when combined with the clarity-driven ergodic search for the nominal path planning, maximizes the average number of tip events serviced and time spent servicing those events. Mathematically, if $n_{\text{tips}}(t)$ denotes the total number of tip events serviced from the start of the mission (t_0) to t , and $t_{\text{tips}}(t)$ denotes the total amount of time the agent spends servicing tip events from t_0 to t (note that the mission from t_0 to t may consist of multiple deployed phases), then it is desirable to maximize both

$$\lim_{t \rightarrow \infty} \frac{n_{\text{tips}}(t)}{t} \text{ and } \lim_{t \rightarrow \infty} \frac{t_{\text{tips}}(t)}{t}. \quad (8)$$

2) *Proposed Solution*: We propose a solution that balances the amount of time the agent services the tip location with the estimated deficit in long-term clarity that results from the host revising its previous ergodic trajectory to achieve rendezvous with the agent.

a) Assumptions:

- Without loss of generality, we assume that the rendezvous sub-mission (i.e., the deployed phase) starts at time 0 (note that this means the start time t_0 of the overall mission is less than or equal to 0).
- We assume full knowledge of the constants related to the dynamics described in Sec. II-D and Sec. II-E: $k_1, k_2, k_3, k_4, b_{\text{agent, min}}, b_{\text{agent, max}}, b_{\text{host, min}}, b_{\text{host, max}}, k_c, \text{max}, v_{\text{agent, max}}, v_{\text{host, max}}, r_{\text{sensing}}$, and k_5 . We also assume knowledge of P_{in} .
- We assume full knowledge of the tip location $p_{\text{tip}} \in \mathcal{D}$, the initial position $p_0 \in \mathcal{D}$ of the host and agent ($p_{\text{host}}(0) = p_{\text{agent}}(0) = p_0$), the initial states of charge of the host and agent ($b_{\text{host}}(0)$ and $b_{\text{agent}}(0)$, respectively), the host's past trajectory $\text{traj} : [t_0, 0] \rightarrow \mathcal{D}$, and the initial clarity map $q(0, \cdot) : \mathcal{D} \rightarrow [0, 1]$.

b) *Decision Variables*: Our solution framework relies on the following decision variables:

- rendezvous point, $p_r \in \mathcal{D}$;
- state of charge $b_r \in [b_{\text{agent, min}}, b_{\text{agent, max}}]$, below which the agent must start returning to the rendezvous point;
- agent's speed, $v_{\text{agent}} \in [0, v_{\text{agent, max}}]$, when traveling to the tip location and returning to the rendezvous point.

c) *Derivable Quantities*: We designed our solution framework such that all the quantities related to the objective function may be derived from the three decision variables above. Specifically, as shown in Fig. 2, we assume straight-line paths between the initial position and the tip location, between the tip location and the rendezvous point, and between the initial position and the rendezvous point. Additionally, we assume that the host travels at a constant speed such that the host and agent arrive at the rendezvous point at the same time. These assumptions mean that given a candidate set of values for the decision variables, we can calculate the following quantities:

- time $t_{\text{agent, service}} \in \mathbb{R}_{>0}$ when the agent arrives at the tip location and starts servicing the tip event,
- time $t_{\text{agent, return}} \in \mathbb{R}_{\geq t_{\text{agent, service}}}$ when the agent starts returning to the rendezvous point from the tip location,
- time $t_{\text{rendezvous}} \in \mathbb{R}_{\geq t_{\text{agent, return}}}$ when both the agent and the host arrive at the rendezvous point,
- speed $v_{\text{host}} \in [0, v_{\text{host, max}}]$ of the host during the deployed phase,
- clarity map $q(t_{\text{rendezvous}}, \cdot) : \mathcal{D} \rightarrow [0, 1]$ at the time of rendezvous,
- clarity map $q^*(t_{\text{rendezvous}}, \cdot) : \mathcal{D} \rightarrow [0, 1]$ that would be obtained if the host were to use the ergodic nominal path planner to travel at the maximum speed (for maximum overall coverage of the domain) from time 0 to time $t_{\text{rendezvous}}$,
- total state of charge $b_{\text{host}}(t_{\text{rendezvous}}) + b_{\text{agent}}(t_{\text{rendezvous}}) \in \mathbb{R}_{\geq 0}$ of the system at the time of rendezvous, and
- time $t_{\text{final}} \in \mathbb{R}_{\geq t_{\text{rendezvous}}}$ when the system's total state of charge would be greater than or equal to the system's initial total state of charge $b_{\text{host}}(0) + b_{\text{agent}}(0)$ if the host were to continue to travel at a speed of v_{host} after reaching the rendezvous point.

d) *Objective Function*: We define the agent's reward function $R_{\text{agent}} : \mathcal{D} \times [b_{\text{agent, min}}, b_{\text{agent, max}}] \times [0, v_{\text{agent, max}}] \rightarrow [0, 1]$ to be given by

$$R_{\text{agent}}(p_r, b_r, v_{\text{agent}}) = \frac{t_{\text{agent, return}} - t_{\text{agent, service}}}{t_{\text{final}}} \quad (9)$$

We define the host's reward function $R_{\text{host}} : \mathcal{D} \times [b_{\text{agent, min}}, b_{\text{agent, max}}] \times [0, v_{\text{agent, max}}] \rightarrow [0, 1]$ to be given by

$$R_{\text{host}}(p_r, b_r, v_{\text{agent}}) = e^{q_{\text{avg, rendezvous}} - q_{\text{avg, rendezvous}}^*} \frac{t_{\text{rendezvous}}}{t_{\text{final}}} \quad (10)$$

where

$$q_{\text{avg, rendezvous}} = \frac{1}{\text{Area}(\mathcal{D})} \iint_{\mathcal{D}} q(t_{\text{rendezvous}}, [x, y]^T) dx dy \quad (11)$$

is the average clarity over the domain at time $t_{\text{rendezvous}}$ if the host were to follow the rendezvous plan and

$$q_{\text{avg, rendezvous}}^* = \frac{1}{\text{Area}(\mathcal{D})} \iint_{\mathcal{D}} q^*(t_{\text{rendezvous}}, [x, y]^T) dx dy \quad (12)$$

is the average clarity over the domain at time $t_{\text{rendezvous}}$ if the host were to use the ergodic nominal path planner to travel at the maximum speed from time 0 to time $t_{\text{rendezvous}}$. Note that the inclusion of t_{final} penalizes rendezvous solutions that consume large amounts of energy.

We design our objective function (to be maximized) $R : \mathcal{D} \times [b_{\text{agent, min}}, b_{\text{agent, max}}] \times [0, v_{\text{agent, max}}] \rightarrow \mathbb{R}$ to be a weighted sum of the two reward functions:

$$R(p_r, b_r, v_{\text{agent}}) = \omega R_{\text{agent}}(p_r, b_r, v_{\text{agent}}) + (1 - \omega) R_{\text{host}}(p_r, b_r, v_{\text{agent}}) \quad (13)$$

where $\omega \in [0, 1]$ is the weight parameter that adjusts the importance of the agent servicing the tip event relative to that of the host looking for tip events.

e) *Constraints*: In order for a candidate set of values for the decision variables $(p_r, b_r, v_{\text{agent}}) \in \mathcal{D} \times [b_{\text{agent, min}}, b_{\text{agent, max}}] \times [0, v_{\text{agent, max}}]$ to be feasible, it must satisfy the following constraints:

- $v_{\text{host}} \in [0, v_{\text{host, max}}]$
- $b_{\text{agent}}(t_{\text{rendezvous}}) \geq b_{\text{agent, min}}$
- $b_{\text{agent}}(t_{\text{agent, service}}) \geq b_r$
- $b_{\text{host}}(t_{\text{rendezvous}}) \in [b_{\text{host, min}}, b_{\text{host, max}}]$
- $b_{\text{host}}(t_{\text{final}}) \in [b_{\text{host, min}}, b_{\text{host, max}}]$

f) *Contingency Plan*: As described in Sec. II-C, the host has no access to the start or end times of tip events, so it is possible that once a feasible optimal solution is found and the agent is deployed to service the tip event, the tip event would end prior to $t_{\text{agent, return}}$. As such, we assume the agent stores enough information about the rendezvous plan prior to deployment so that when the tip event ends early (a “contingency”), it can calculate the new rendezvous position $p_{\text{rendezvous, new}}$ and new rendezvous time $t_{\text{rendezvous, new}}$ based on the host’s current position and its trajectory until $t_{\text{rendezvous}}$.

For mission efficiency (Sec. III-1), the agent will return early in the event of a contingency, which allows it to be redeployed quickly. The agent will return to the host as quickly as its speed and battery allow while the host continues towards the original rendezvous point p_r at its own, unchanged speed v_{host} . Once docked, the host will replan its trajectory for the docked phase.

IV. SIMULATION & RESULTS

A. Simulation

We have evaluated the proposed rendezvous algorithm using the parameters from Sec. I. The mission follows the

structure in Fig. 1, starting with the AAV agent docked to the ASV host. When a tip location is identified, the optimization in Sec. III is applied. If a feasible rendezvous solution exists, the AAV is deployed to service the event and then return to the rendezvous point.

In this work, we consider three different mission strategies, the first of which represents our proposed approach, the second and third of which represent benchmark variations:

1) *Collaborative-Ergodic Strategy*: The collaborative-ergodic strategy generates the nominal path using the clarity-driven ergodic search described in Sec. II-E. When solving the optimization problem for rendezvous planning, the candidate rendezvous point can be located anywhere within the mission domain $\mathcal{D}$. Following the rendezvous, the ASV host recomputes a new ergodic nominal path to follow.

2) *Directive-Ergodic Strategy*: The directive-ergodic strategy is identical to the collaborative-ergodic strategy except that the rendezvous point must lie on the original ergodic nominal path.

3) *Directive-Transect Strategy*: The directive-transect strategy generates straight-line paths that cover the domain in a back-and-forth line-transect manner (a “lawnmower pattern”). When solving the optimization problem for rendezvous planning, the candidate rendezvous point must lie on the original line-transect path.

B. Results

TABLE I
NUMBER OF TIP EVENTS SERVICED AND TOTAL SERVICING TIME FOR EACH STRATEGY

Strategy	Tips Serviced	Total Servicing Time (s)
Collaborative-Ergodic	8	1508.0
Directive-Ergodic	5	1092.0
Directive-Transect	3	913.75

In Table I, we list the number of tip events serviced and the total servicing time for each strategy. The results show that in terms of both the number of tip events serviced and the total time servicing them, the collaborative-ergodic strategy significantly outperforms the directive-ergodic and directive-transect strategies.

Fig. 3 shows that while the strategies differ in performance, all achieve similar averaged-over-domain clarity at steady state. Fig. 4 illustrates that the host’s path closely follows the agent’s path in the collaborative-ergodic strategy while deviations occur in the other two strategies. This aligns with the better performance of the collaborative-ergodic strategy, where the agent carries less rendezvous burden. Fig. 5 shows the ASV host’s state-of-charge and speed trajectories, indicating minimal energy loss even when traveling at maximum speed and charging the agent. Note that the host’s rendezvous speed (calculated based on the optimal values for the three decision variables) is typically slower than its maximum. Fig. 6 shows the AAV agent’s state-of-charge and speed trajectories, highlighting that optimal rendezvous plans drain the AAV’s battery as much as possible before




Fig. 2. The mission strategies compared in this work.




Fig. 3. Evolution of mean clarity for each strategy

reaching the rendezvous point, which allows the agent to service the tip event for as long as possible. The AAV agent's speed is also kept below its maximum to conserve energy, as slower speeds allow more time to service the tip event. Fig. 7 shows the agent's mission-state trajectories. We see that the collaborative-ergodic strategy has the shortest total duration for the idle-on-host and charging-on-host states combined, which leads to better performance.

V. CONCLUSIONS AND FUTURE WORK

In this work, we considered a persistent mission planning tip-and-cue problem involving a single host and agent. We proposed the collaborative-ergodic strategy, which combines a clarity-driven ergodic nominal path planner and a rendezvous planner whose solution can be characterized by three decision variables. We demonstrated its superior performance over the directive-ergodic and directive-transect strategies.

In future work, we will add flexibility to the collaborative-ergodic strategy, in addition to investigating the theoretical optimality and feasibility properties of the approach. Specifically, we will consider mechanisms for detecting and queuing tip events during a rendezvous deployment, allow agents to service multiple tip events during a single deployment, and consider the optimization of both the topology (i.e., the numbers of hosts and agents) and the mission plan.

REFERENCES

- [1] M. Rodger and R. Guida, "Ship trajectory prediction model for space-based maritime surveillance," in *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2024, pp. 1447–1450.
- [2] J. Sung, S. Kim, J. Lee, D. Yi, J. Ryu, and J. Lee, "Spaceborne sar antenna development status and plans suitable for tip and cue monitoring," in *2024 International Symposium on Antennas and Propagation (ISAP)*. IEEE, 2024, pp. 1–2.
- [3] C. Gentemann, J. P. Scott, P. L. Mazzini, C. Pianca, S. Akella, P. J. Minnett, P. Cornillon, B. Fox-Kemper, I. Cetinić, T. M. Chin et al., "Saildrone: Adaptively sampling the marine environment," *Bulletin of the American Meteorological Society*, vol. 101, no. 6, pp. E744–E762, 2020.
- [4] "The Wave Glider — Unmanned Surface Vehicle by Liquid Robotics." [Online]. Available: <https://www.liquid-robotics.com/wave-glider/overview/>
- [5] "SeaTrac." [Online]. Available: <https://www.seatrac.com/>
- [6] S. Ramasamy, J.-P. F. Reddinger, J. M. Dotterweich, M. A. Childers, and P. A. Bhounsule, "Coordinated route planning of multiple fuel-constrained unmanned aerial systems with recharging on an unmanned ground vehicle for mission coverage," *Journal of Intelligent & Robotic Systems*, vol. 106, no. 1, p. 30, 2022.
- [7] F. Zeng, Z. Chen, J.-P. Clarke, and D. Goldsman, "Nested vehicle routing problem: Optimizing drone-truck surveillance operations," *Transportation Research Part C: Emerging Technologies*, vol. 139, p. 103645, 2022.
- [8] D. R. Agrawal and D. Panagou, "Multi-agent clarity-aware dynamic coverage with gaussian processes," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*. IEEE, 2024, pp. 37–44.
- [9] K. Naveed, D. Agrawal, C. Vermillion, and D. Panagou, "Energy-aware clarity-driven ergodic search," in *IEEE International Conference on Robotics and Automation*. IEEE, 2024.
- [10] P. Foeht, E. Kaufmann, A. Romero, R. Penicka, S. Sun, L. Bauersfeld, T. Laengle, G. Cioffi, Y. Song, A. Loquercio et al., "Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight," *Science robotics*, vol. 7, no. 67, p. eabl6259, 2022.
- [11] G. Marmorino and C. Trump, "A salinity front and current rip near Cape Hatteras, North Carolina," *Journal of Geophysical Research*, vol. 99, no. C4, pp. 7627–7637, 1994.
- [12] M. Babin, C. Roesler, and J. Cullen, "Real-time Coastal Observing Systems for Marine Ecosystem Dynamics and Harmful Algal Blooms: Theory, Instrumentation, and Modelling," *Oceanographic Methodology series: UNESCO Publishing*, 2008.
- [13] A. Colefax, P. Butcher, and B. Kelaher, "The potential for unmanned aerial vehicles (UAVs) to conduct marine fauna surveys in place of manned aircraft," *ICES Journal of Marine Science*, vol. 75, no. 1, pp. 1–8, 2018.
- [14] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4–5, pp. 432–442, 2011.
- [15] L. Dressel and M. J. Kochenderfer, "Tutorial on the generation of ergodic trajectories with projection-based gradient descent," *IET Cyber-Physical Systems: Theory & Applications*, vol. 4, no. 2, pp. 89–100, 2019.




Fig. 4. Paths for the ASV host and AAV agent. Note that while the locations of all the tip events are shown, those tip events do not all start at the same time and thus do not all simultaneously exist throughout the mission. Also note that the initial position is (100 m, 100 m).




Fig. 5. State-of-charge and speed trajectories for ASV host. The colors correspond to the AAV agent's mission states in Fig. 7 (black = idle on host, green = charging on host, blue = traveling to tip event, purple = servicing tip event, red = returning to rendezvous point). Note that $b_{\text{host}, \text{min}} = 0$ Wh, $b_{\text{host}, \text{max}} = 6500$ Wh, and $v_{\text{host}, \text{max}} = 2.315$ m/s.




Fig. 6. State-of-charge and speed trajectories for AAV agent. The colors correspond to the AAV agent's mission states in Fig. 7 (black = idle on host, green = charging on host, blue = traveling to tip event, purple = servicing tip event, red = returning to rendezvous point). Note that $b_{\text{agent}, \text{min}} = 0$ Wh, $b_{\text{agent}, \text{max}} = 44.4$ Wh, and $v_{\text{agent}, \text{max}} = 30$ m/s.




Fig. 7. Mission-state trajectory for AAV agent. The mission states are described in Sec. II-D.