IEEE Control Systems Letters paper presented at
2023 American Control Conference (ACC)
San Diego, CA, USA. May 31 - June 2, 2023

Predictive Velocity Trajectory Control for a Persistently Operating Solar-Powered Autonomous Surface Vessel

Kavin Govindarajan, Ben Haydon, and Chris Vermillion[1]

Abstract— **The Gulf Stream represents a major potential resource for renewable energy but is presently only sparsely characterized via radar, buoys, gliders, and intermittently operating human-operated research vessels. Dramatically greater resolution is possible through the use of persistently operating autonomous surface vessels (ASVs), which can be powered by wind, wave, or solar resources. Optimizing the control of these ASVs, taking into account the device and environmental properties, is crucial to obtaining good data. An ASV's path and velocity profile along that path both significantly influence the amount of a mission domain that can be covered and, ultimately, the scientific quality of the mission. While our previous work focused on optimizing the path of a solar-powered ASV with fixed speed, the present work represents the complement: optimizing the speed for a given path, accounting for the ASV dynamics, flow resource, and solar resource. We perform this optimization through a model predictive controller that maximizes the projected distance traversed, with a terminal incentive that captures the estimated additional long-duration range that is achievable from a given terminal battery state of charge. We present simulation results based on the SeaTrac SP-48 ASV, Mid-Atlantic Bight/South-Atlantic Bight Regional Ocean Model, and European Centre for Medium-Range Weather Forecasts solar model. Our results show improved performance relative to simpler heuristic controllers that aim to maintain constant speed or constant state of charge. However, we also show that the design of the MPC terminal incentive and design of the heuristic comparison controller can significantly impact the achieved performance; by examining underlying simulation results for different designs, we are able to identify likely causes of performance discrepancies.**

I. Introduction

The portion of the Gulf Stream adjacent to North Carolina's Outer Banks has been estimated to possess 160 TWh/year of energy [1], making it a strong candidate for deploying devices for harvesting renewable energy. However, existing resource characterizations are limited by spatial and temporal sparsity. Specifically, to date, these characterizations have been made through buoy-mounted acoustic Doppler current profilers (ADCPs), gliders, radar measurements of surface currents, and ADCP-equipped research vessels. Buoy-mounted ADCPs [2] allow for long-term data collection (months) for an entire water column but are restricted to a single location. Glider-mounted ADCPs (see [3], [4]) allow for measurements to be taken across a full path; however, the mission duration is limited by the flow environment and the battery capacity. Research vessels [5] similarly allow for measurements taken across an entire path, but persistence is precluded by staffing and fuel constraints. Finally, radar [6] can measure surface flow across a wide area but cannot measure sub-surface flow speeds.

To obtain a sufficiently spatially granular, long-duration data set suitable for performing site studies for energy-harvesting systems, the aforementioned sparse and/or short-duration measurements must be supplemented by granular, persistent, and autonomous observations. Renewably powered marine mobile robots, such as sailing drones [7] and solar-powered autonomous surface vessels (ASVs) [8], aim to satisfy those needs. These systems can operate persistent missions, measure flow speeds at depth using ADCPs, and move around in their environment.

To realize the full potential of renewably powered ASVs, each ASV's path and velocity profile must be controlled in a way that maximizes the quantity and utility of the collected scientific data. These control requirements are challenging for two significant reasons. First, both the renewable resource and the flow environment (which impact the energetic and dynamic performance of an ASV) vary in space and time. This is particularly challenging in the Gulf Stream, where the maximum surface currents can exceed 2 m/s, which approaches the maximum achievable transit speeds of existing sailing drones and solar-powered ASVs. Second, paths and velocity profiles for *persistent* missions should ideally be optimized over an *infinite* horizon, which stands in contrast with most existing information-driven planning studies aimed at oceanographic missions (See [9], [10]). Because infinite-horizon optimizations are generally not computationally achievable, effective and tractable optimizations for persistent missions require some estimate of the infinite-horizon scientific value associated with a given battery state of charge at the end of a finite horizon. Such a quantification has not been studied to-date.

This paper addresses the velocity trajectory planning element of persistent control. Specifically, for a prescribed path, we present an MPC formulation that incentivizes distance travelled over a finite prediction horizon, plus a terminal incentive that acts as an estimate of the additional long-horizon distance that can be traversed as a function of the terminal state of charge (relative to the traversable long-horizon distance under zero terminal state of charge). To study the problem in a concrete manner, we focus on the solar-powered SeaTrac SP-48 ASV, operating in a prescribed mission domain within the U.S. Gulf Stream. We present an energetic and mobility model for the ASV, along with a tunable formulation for the MPC optimization's terminal

[1]Kavin Govindarajan, Ben Haydon, and Chris Vermillion serve in the roles of Undergraduate Research Assistant, Graduate Research Assistant, and Associate Professor, respectively, with the Department of Mechanical and Aerospace Engineering, North Carolina State University, Raleigh, NC 27606, USA.

979-8-3503-2806-6/$31.00 ©2023 AACC

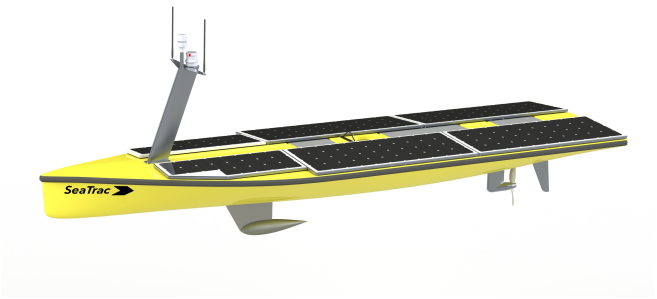

Fig. 1: SeaTrac ASV considered in this work [8] – Image used with permission.

incentive. Based on this formulation, we compare the results of the MPC-based velocity trajectory controller against two families of closed-form benchmarks: (i) controllers that maintain a prescribed constant velocity whenever possible (subject to a minimum allowable state of charge and the requirement of fighting current to stay on the prescribed path) and (ii) controllers that maintain prescribed time-of-day-dependent battery state of charge whenever possible (subject to the requirement of fighting current to stay on the prescribed path). Our results show that the best MPC approach outperforms the best implementable heuristic benchmarks; however, tuning has a significant effect on performance in each case. By examining key simulation variables, such as the ASV velocity profile and state of charge profile, we are able to gain some insight into the reasons why certain approaches outperform others. It is worth noting that while the present paper's focus is restricted to the problem of velocity trajectory control, the algorithm and results presented herein serve as a complement to our earlier work in [11], which addresses path planning under the assumption of a prescribed (un-optimized) velocity profile.

II. Modeling of the ASV and Environment

For this work, we consider the SeaTrac SP-48, a solar-powered ASV [8] shown in Fig. 1, with key parameters given in Table I.

We model the ASV's translational dynamics based on the two-state model below:

$$\begin{aligned} \dot{x} &= v_b \cos\theta_b + v_f(x,y,t)\cos\theta_f(x,y,t) \\ \dot{y} &= v_b \sin\theta_b + v_f(x,y,t)\sin\theta_f(x,y,t) \end{aligned} \tag{1}$$

where x and y are the two spatial states of the system, v_b and θ_b are the boat speed *relative to the moving water* and boat heading (which serve as control variables), respectively, and $v_f(x,y,t)$ and $\theta_f(x,y,t)$ are the spatiotemporally varying flow speed and direction at the x,y location of the boat. We also model the state of charge of the on-board battery, denoted by b, according to:

$$\dot{b} = \eta_s A_s I(t) - \eta_m P_D(v_b) - P_e \tag{2}$$

where η_s is the solar panel efficiency, A_s is the solar panel area, $I(t)$ is the solar shortwave radiation (in units of W/m^2), η_m is the motor efficiency, $P_D(v_b)$ is the power loss due to hydrodynamic drag, and P_e is the electrical power consumption of the on-board electronic systems. The power loss from drag, $P_D(v_b)$, is based on data made available by SeaTrac, a fit for which is shown in Fig. 2.

TABLE I: Speed Control Parameters

Variable Name	Symbol	Value	Units
Solar Panel Area	A_s	4.17	m^2
Average Solar Shortwave Radiation	I_{avg}	175	W/m^2
On-board Electronics	P_e	100	W
Wetted Area	A_w	5.82	m^2
Solar Panel Efficiency	η_s	0.18	-
Motor + Propeller Efficiency	η_{mp}	0.25	-
Drag Coefficient	C_D	0.0030	-

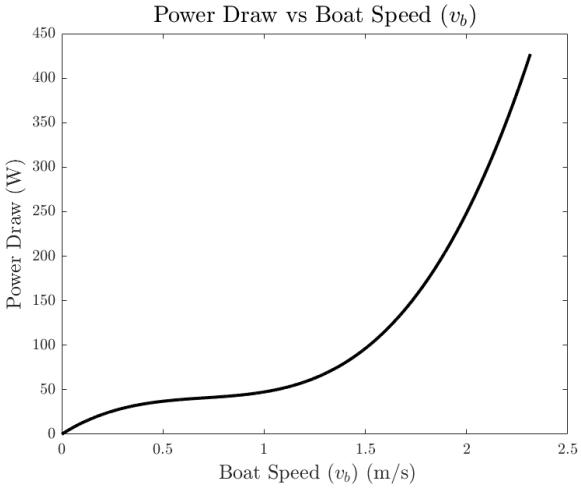

Fig. 2: ASV power draw vs. boat speed

In this paper, θ_b and v_b are treated as control variables, under the assumption that lower-level controllers quickly regulate the ASV's heading and speed to their target values. In practice, the SeaTrac ASV does in fact come equipped with a closed-loop heading controller, and the motor command to the ASV is monotonically related to v_b. To obtain values of $v_f(x,y,t)$ and $I(t)$ used for the simulations in this work, we obtained Gulf Stream flow data from the MABSAB-ROMS hindcast model [12] and solar irradiance data from the ERA-Interim model [13]. This work considers a rectangular mission domain in the U.S. Gulf Stream, adjacent to North Carolina's Outer Banks and illustrated in Fig. 3. This selected region lies at the confluence of three critical water masses, namely the Mid-Atlantic Bight, South-Atlantic Bight, and Gulf Stream, making it an extremely interesting target for oceanographic observation. In addition to showing the mission domain, Fig. 3 shows two snapshots of the flow speed contours, demonstrating the spatiotemporal variability of this disturbance. Fig. 4 shows the temporally varying solar resource, which exhibits a predictable diurnal pattern with inconsistent peaks from day to day. Because of the size of the mission domain, spatial variation in the solar irradiance is neglected.

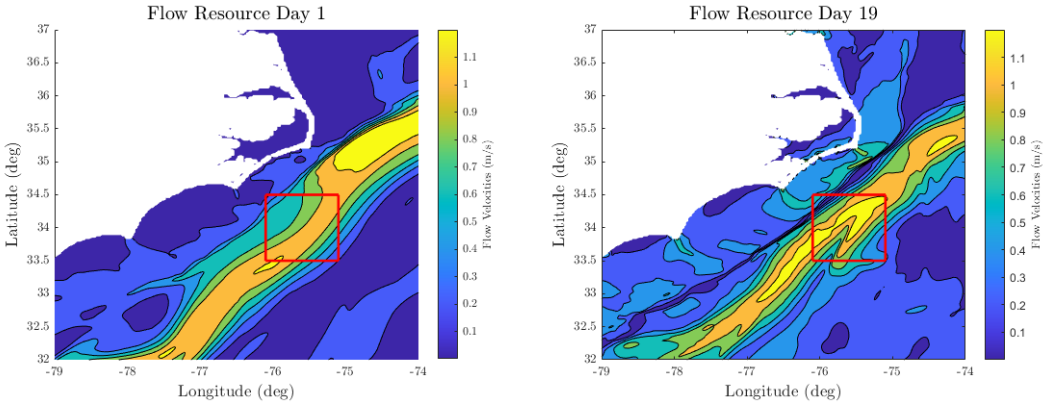

Fig. 3: Mission domain and flow resource at two selected times, separated by 18 days.

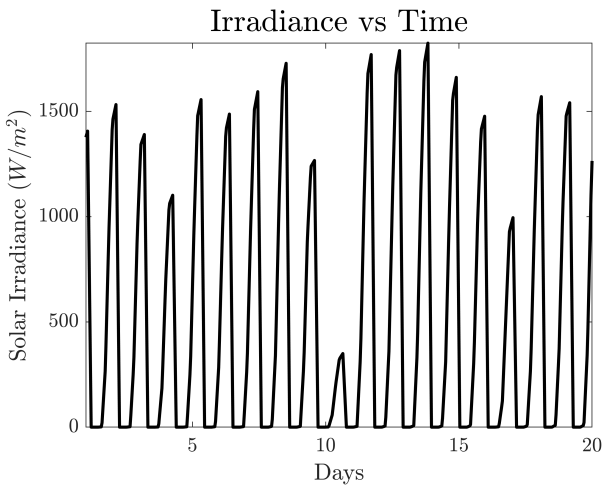

Fig. 4: Solar irradiance vs. time

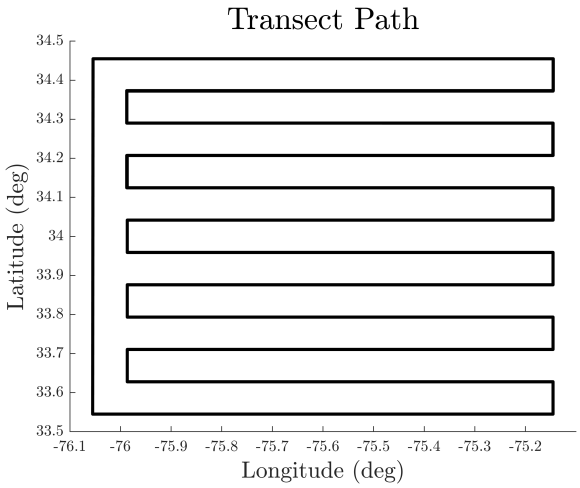

Fig. 5: Prescribed transect path

III. Candidate Control Strategies

The ASV controller must ultimately set the values for the heading (θ_b) and speed relative to the moving water (v_b). In this work, we consider a prescribed mission path, shown in Fig. 5. The pre-specification of the path makes heading control relatively simple, but still non-trivial due to the fact that the Gulf Stream flow, which can approach 2 m/s at particular locations and times, will often act to push the ASV off its path. Because of the pre-specification of the path in this work, we have elected to work with a relatively simple closed-form heading controller, which is used with each candidate velocity controller. Velocity control represents a critical subject matter in this work. In particular, we have examined several velocity control strategies – both predictive and heuristic – which are described in this section.

A. Heading Control

A common closed-form heading control strategy is considered throughout the results of this work. Written plainly, the heading control strategy is to choose a heading such that the ASV's true velocity points towards the next waypoint, where the waypoints are determined by the prescribed path as shown in Fig. 5. If the relative velocity chosen by the controller is insufficient to overcome the flow, the heading is taken directly into the flow in order to minimize deviation from the path. If the ASV leaves the path, it will continue in a straight line to the next waypoint, at which point it will rejoin the path. We chose the waypoints to be each of the points at which the path changes direction; however, including more waypoints would force the ASV to more closely adhere to the path. Written mathematically, the heading controller chooses:

$$\theta_b = \begin{cases} \pi + \theta_f & \text{if } v_b < v_f \sin(\theta_{fw}) \\ \theta_w + \sin^{-1}\left(\frac{v_f}{v_b}\sin(\theta_{fw})\right) & \text{if } v_b \geq v_f \sin(\theta_{fw}) \end{cases} \quad (3)$$

where θ_f and θ_w are the flow direction and heading to the next waypoint, respectively, and $\theta_{fw} = \theta_f - \theta_w$.

B. Model Predictive Velocity Control (MPC)

For this work, we consider multiple MPC variants, although all of them share the same general framework. In

particular, at each discrete time step, denoted as k, an $N \times 1$ sequence of speeds $\mathbf{v}_b(k)$ is optimized based on the maximization of a prescribed reward function, subject to hard constraints and a discrete-time approximation of the dynamic model. The elements of this sequence (for the optimization performed at step k) will be denoted as:

$$\mathbf{v}_b(k) = [\ v_b(k|k) \quad \dots \quad v_b(k+N-1|k)\ ]^T. \tag{4}$$

The first element of that N-step optimized trajectory is implemented, and the optimization is performed again at the next time step. Mathematically, $\mathbf{v}_b(k)$ is the solution to the following optimization problem:

$$\begin{aligned} \max_{\mathbf{v}_b(k)} \quad & J(\mathbf{v}_b(k); x(k), y(k), b(k)) = \sum_{i=k}^{k+N-1} d(v_b(i|k)) \\ & + \tilde{J}^*_\infty(b(k+N|k)) \\ \text{s.t.} \quad & b_{\min} \le b(i|k) \le b_{\max} \quad i = k+1 \dots N \\ & v_{\min} \le \mathbf{v}_b(i|k) \le v_{\max} \quad i = k+1 \dots k+N-1 \\ & \text{Equations (1), (2), and (3)} \end{aligned} \tag{5}$$

where $d(v_b(i|k))$ is the projected distance traveled (projected onto the path) during time step i, accounting for the flow disturbance, $b(i|k)$ is the projected battery state of charge at time step i, and $\tilde{J}^*_\infty(b(k+N|k))$ is the estimate of the optimal infinite horizon reward to go. This is the term that accounts for the fact that the mission is designed to be persistent whereas the optimization only accounts for a fixed amount of time (the prediction horizon). The purpose of this term is to estimate the amount of extra "value" that exists at the end of the prediction horizon based on contextual factors (i.e. how much extra distance the remaining battery is worth). The result of the MPC optimization is a control sequence, $\mathbf{v}^*_b(k)$, for which the first step is implemented as follows:

$$v_b(k) = \mathbf{v}^*_b(k|k). \tag{6}$$

The differences between the MPC variants considered in this work lies in the formulations for the reward to go terms. We now review the candidate mathematical structures (and rationale for choosing said structures) for this reward to go.

1) Baseline Reward to Go: For the first reward to go formulation, we define $v(P)$ as the inverse of the power curve in figure 2. This represents the steady-state, or "break-even," speed associated with power level P in calm water. Based on this characterization, the terminal reward in this formulation characterizes the additional distance that can be traversed with state of charge b, compared with a state of charge equal to zero, over a finite future period of time, T. Due to the fact that saving battery for multiple days is infeasible due to the overwhelming sunlight at midday, we take this period of time to be $T = 6$ hours, in this work. To make this calculation, we must predict an amount of available solar irradiance over that time window. For this terminal reward variant, we assume the the average solar irradiance for the mission domain, denoted by P_{avg}, will be available. With these ingredients in place, $\tilde{J}^*_\infty(b)$ is given by:

$$\tilde{J}^*_\infty(b) = T\left(v\left(P_{\text{avg}} + \frac{b}{T}\right) - v(P_{\text{avg}})\right) \tag{7}$$

2) Scheduled Reward to Go: The aforementioned formulation makes the assumption that solar power is constant when in reality, it strongly depends on time of day. For the second MPC implementation, we measure the average power expenditure as a function of time of day from a simulation of the previous version of MPC, $P_{avg}(t)$, where t represents time of day. Our new reward to go is then:

$$\tilde{J}^*_\infty(b) = T\left(v\left(P_{\text{avg}}(t) + \frac{b}{T}\right) - v(P_{\text{avg}}(t))\right) \tag{8}$$

3) Terminal Cost: For this MPC implementation, we utilize the knowledge that solar power depends on time of day, as explored in the previous implementation. However, while the scheduled reward to go formulation estimates future performance based upon estimates of power consumption, this implementation takes advantage of the fact that the cyclic nature of the solar resource incentivizes a similarly cyclic state of charge. We represent a typical state of charge trajectory in Fig. 6. This MPC formulation uses that state of charge trajectory as a target and penalizes deviation from it at the end of the MPC horizon. Thus, we have that our terminal cost is given by:

$$\tilde{J}^*_\infty(b) = -k_p\left(b(t) - b_{des}(t)\right) \tag{9}$$

where k_p represents a gain tuned heuristically, $b(t)$ is the battery state of charge at the end of the prediction horizon, and $b_{des}(t)$ is the prescribed battery state of charge at the end of the prediction horizon. For this implementation, we use $k_p = 0.75$.

C. Closed-Form Heuristic Benchmarks

For comparison, we consider several closed-form heuristic benchmark controllers. These controllers not only act as benchmarks to compare the performance of MPC against; they also serve as dramatically simpler control structures that in some cases can be tuned to approximate the performance of MPC. It is also worth noting that some of the benchmarks depend on omniscient knowledge that is not in fact available for real-time control; in these circumstances, the benchmarks result in upper bounds on the performance of any control strategy.

Each benchmark controller operates slightly differently, but there are a few key concepts that hold for each. Most importantly is the concept that wasting energy is inherently suboptimal (and therefore should be minimized for a reasonable benchmark). Specifically, energy harvested when at a full battery is wasted due to battery size constraints. Because a large amount of solar power is available during the middle of the day, often exceeding the maximum power draw from the motor, some increase in state of charge during those hours is unavoidable. For this reason, the general trend for any reasonable velocity trajectory will involve moving quickly during the daytime (in order to minimize energy losses from overcharging the battery) and moving slower at night.

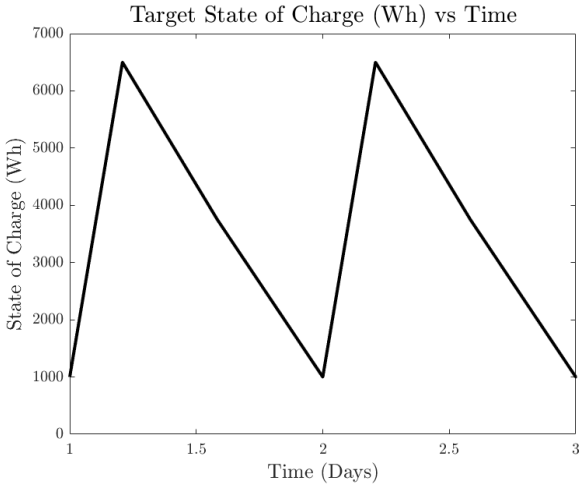

Fig. 6: Prescribed state of charge vs. time. Note that, based on the way that the data is formatted, the beginning of each day corresponds to noon, which also corresponds to the largest solar irradiance. Hence, it is advantageous for the battery to be at a low state of charge at this time.

Furthermore, the battery level just before peak irradiance should be somewhat low to account for the unavoidable energy gain.

The following sections detail the formulations for four benchmarks that are designed to follow the above guidance in slightly different ways.

1) Constant State of Charge Benchmark: The first benchmark strategy attempts to avoid overcharging by keeping the battery state of charge as close as possible to a constant, prescribed value during the day by increasing the power given to the motor to match the power generated by the solar panels. However, at peak irradiance, this strategy still generates energy in excess of what the motor is able to expend. For this reason, a minimum speed is enforced, enabling the ASV to use the excess energy to travel at a slower speed during the nighttime when there is no solar energy incoming. Mathematically, this strategy is represented as follows:

$$v_{b,\min} = \begin{cases} ||\vec{v}_{\min} - \vec{v}_f|| & \text{if } \hat{\vec{v}}_{\min} \cdot \vec{v}_f \leq ||\vec{v}_{\min}|| \\ v_f \sin(\theta_f - \theta_b) & \text{if } \hat{\vec{v}}_{\min} \cdot \vec{v}_f > ||\vec{v}_{\min}|| \end{cases} \tag{10}$$

$$v_b = \min(v_{b,\max}, \max(v(P_b), v_{b,\min})) \tag{11}$$

where $v(P_b)$ is the velocity required to draw a power equivalent to the power generated from the solar irradiance, $\vec{v}_{\min}$ is the vector representing traveling at the specified minimum speed towards the next waypoint, and $\vec{v}_f$ is the flow velocity at the current point. For this strategy, we defined the minimum true velocity as $||\vec{v}_{\min}|| = 0.5$m/s.

The switching condition is in place to account for circumstances where the flow has a component towards the next waypoint with a magnitude greater than $||\vec{v}_{\min}||$, in which case moving at the speed of that flow component takes less energy than moving at the specified minimum speed.

2) Prescribed State of Charge Setpoint Benchmark: While the constant state of charge strategy makes an explicit effort to avoid overcharging (or discharging) the battery, it does not strategically leverage the battery in dealing with the inevitable diurnal variations in the solar resource. In reality, it is well-known that the solar resource will increase during the morning, peak at midday, and decrease in the evening. This can be capitalized upon by strategically charging the battery during the day and discharging during the night. To do so, we define a time-varying setpoint for the state of charge shown in Fig. 6, and denoted by $b_{\text{des}}(t)$. These values were tuned heuristically in order to minimize energy loss from overcharging the battery and repeat on a diurnal cycle. Given this state of charge setpoint, the ASV velocity is computed according to the following proportional controller:

$$v_b(t) = v(P_b(t)) + k_p(b(t) - b_{\text{des}}(t)) \tag{12}$$

where k_p is a tunable gain affecting how aggressively this strategy attempts to force the state of charge to track b_{des}. For this strategy, we took $k_p = 0.75$.

3) Constant Relative Velocity Benchmark: This strategy is predicated on the fact that moving at a consistent speed would be optimal in the absence of environmental complications from the flow and the solar resource (due to the convex relationship between relative velocity and power expended over the relevant speeds). For this reason, we choose three speeds: v_{target} is the prescribed constant speed, $v_{b,\min}$ is the minimum speed required to maintain position (see equation (10) with $||\vec{v}_{\min}|| = 0$), and $v_{b,\max}$ is the maximum allowed relative velocity based on the limits of the ASV. In implementing this strategy, we choose which speed to command based on the following law:

$$v_b = \begin{cases} v_{b,\max} & \text{if } b \geq b_{\text{thresh}} \\ \max(v_{\text{target}}, v_{b,\min}) & \text{if } b < b_{\text{thresh}} \end{cases} \tag{13}$$

where b_{thresh} is a battery threshold above which the ASV is commanded to go the maximum velocity in order to reduce the possibility of overcharging the battery. For this strategy, we chose this value to be 90% of the total battery capacity.

4) Constant Speed Infinite Battery Benchmark: For our final benchmark, we consider a nonphysical upper bound on the constant-velocity performance of the system by removing battery constraints (essentially modeling an infinitely large battery). We then tune a speed v_{const} such that at the end of a simulation, the terminal battery level is the same as the initial battery level for an ASV moving at v_{const} for the entire simulation. For this strategy, we computed $v_{\text{const}} = 2.2$m/s. The control law is then:

$$v_b = v_{\text{const}} \tag{14}$$

While this isn't physically possible, it provides a benchmark for the best possible handling of the solar resource. Performing better than this benchmark is only possible with control intelligently accounting for the flow environment through a variable-velocity strategy.

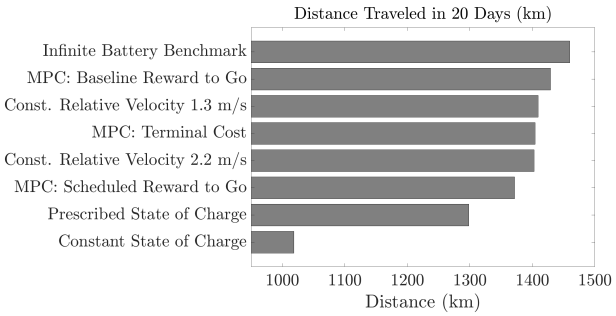

Fig. 7: Total distance traveled. Note that the Infinite Battery Benchmark Strategy represents a non-tight upper bound on achievable performance, whereas the Const. Relative Velocity 1.3 m/s requires precise knowledge of the flow field.

IV. Results and Discussion

Total distance traveled and state of charge trajectories from simulations of each strategy operating a 20-day long mission are displayed in Fig. 7 and Fig. 8 respectively. The simulation was conducted with the boat and environment model updated on 1-minute time-steps.

As can be seen from the results, the strategies utilizing MPC outperform the benchmarks where the battery is constrained. This can be attributed to the fact that the MPC formulation allows for improved allocation of the battery resource to account for times where an area of high flow must be traversed and for times when there is a limited solar resource.

Another metric for the performance of the MPC-based strategies are their performance relative to the constant speed infinite battery benchmark. The constant speed infinite battery benchmark represents the upper bound on the constant-velocity performance of the system by removing the battery constraints, thus representing an optimal utilization of the battery resource across the mission period. From Fig. 7, we see that the MPC strategies are able to nearly match the performance of the benchmark. This shows that the MPC strategies are able to intelligently allocate their battery resource to maximize performance.

In addition to Fig. 7, Fig. 8 provides insight into the battery usage of each strategy. From this, we can see that the MPC strategies consistently utilize large amounts of the battery resource, as opposed to the benchmarks, which are more conservative in their battery usage. An interesting characteristic of each MPC strategy to note are their allocation of battery prior to day 10, which marks a period of low solar irradiance and high flow magnitude. Each MPC strategy is able to allocate enough battery to traverse this "adverse" region without sacrificing performance by allowing for a high state of charge once the adverse region has been traversed.

Fig. 8 further provides insight into the performance of each MPC strategy. We can see that the MPC formulation with a scheduled reward to go has a more conservative utilization of battery when compared to the baseline reward to go implementation of MPC. As a result, this strategy underperforms the other MPC strategies. We further see that the implementation of MPC utilizing a penalty on deviating from a prescribed state of charge more aggressively utilizes battery, which reduces operational effectiveness in adverse regions. As such, we find that the MPC implementation with the baseline reward to go effectively manages battery utilization so as to maximize distance traveled.

Finally, Figs. 9 and 10 provide a depiction of the flow speed and direction over the course of a simulation, for the MPC strategy with the baseline reward to-go. Note that the flow direction is taken relative to the ASV, with a 0-degree direction representing a following sea and a +/- 180 degree direction representing a head sea. The observed 180-degree switching behavior arises due to the transect pattern prescribed in Fig. 5. In these (and all) of the simulations, the flow variables are largely governed by spatial variations, as the ASV repeatedly crosses the Gulf Stream jet (i.e., the strongest portion of the Gulf Stream) several times over the duration of the simulation, whereas time scales of the Gulf Stream are on the order of several days.

V. Conclusions and Future Work

In this paper, we developed an MPC-based velocity trajectory controller for persistent Gulf Stream exploration using a solar-powered ASV. The proposed controllers were compared against various closed-form heuristic benchmarks. Simulation results, which were based on an ERA-Interim solar resource model and a MAB-SAB-ROM ocean current model, show that the proposed MPC-based velocity trajectory controllers enable the ASV to traverse a larger distance during the mission period as opposed to the benchmarks. This can be attributed to the MPC-based velocity trajectory controllers' ability to estimate long-term performance through an appropriately designed terminal reward function.

This work represents a complement to our previous work in [11]. While the previous work explored path optimization for a constant commanded velocity, this work assumes a prescribed path and optimizes the commanded velocity. Moving forward, we seek to combine the path and velocity optimization strategies presented in these works to present a unified persistent path and velocity planner for spatiotemporally varying environments.

References

[1] K. Haas, "Assessment of energy production potential from ocean currents along the United States coastline," Georgia Inst. of Technology, Atlanta, GA (United States), Tech. Rep., 2013.

[2] R. He, J. Bane, M. Muglia, S. Haines, C. Lowcher, Y. Gong, and P. Taylor, "Gulf Stream marine hydrokinetic energy resource characterization off Cape Hatteras, North Carolina USA," in *OCEANS 2016 - Shanghai*, 2016, pp. 1–4.

[3] G. Gawarkiewicz, R. E. Todd, W. Zhang, J. Partida, A. Gangopadhyay, M.-U.-H. Monim, P. Fratantoni, A. M. Mercer, and M. Dent, "The changing nature of shelf-break exchange revealed by the OOI Pioneer array," *Oceanography*, vol. 31, no. 1, pp. 60–70, 2018.

[4] R. E. Todd, "Export of Middle Atlantic Bight shelf waters near Cape Hatteras from two years of underwater glider observations," *Journal of Geophysical Research: Oceans*, vol. 125, no. 4, p. e2019JC016006, 2020.

[5] M. Muglia, H. Seim, and P. Taylor, "Gulf Stream marine hydrokinetic energy off Cape Hatteras, North Carolina," *Marine Technology Society Journal*, vol. 54, no. 6, pp. 24–36, 2020.

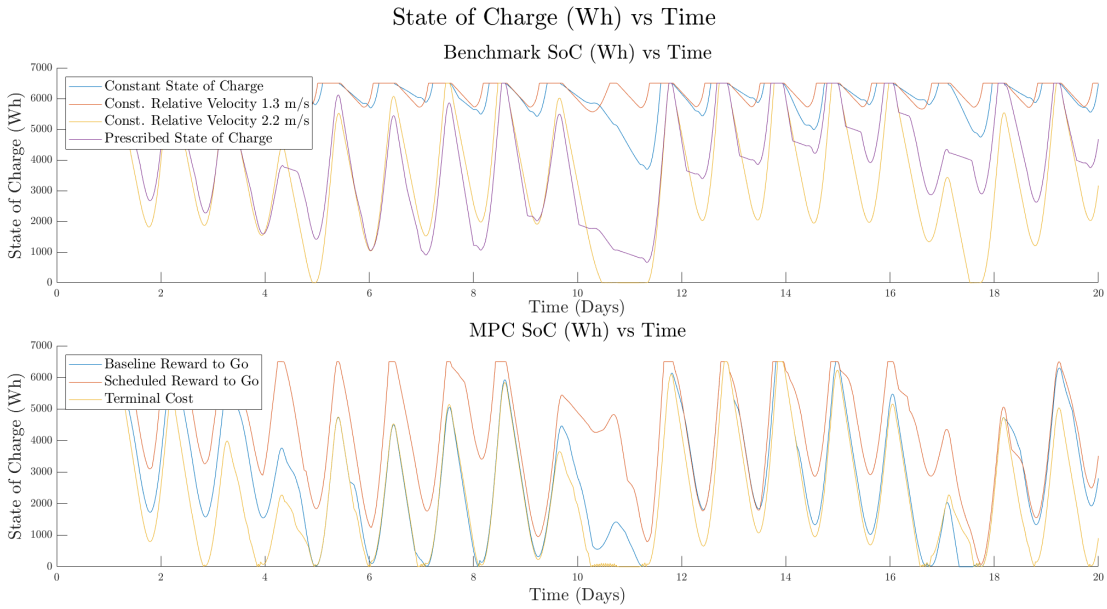

Fig. 8: State of charge vs. time

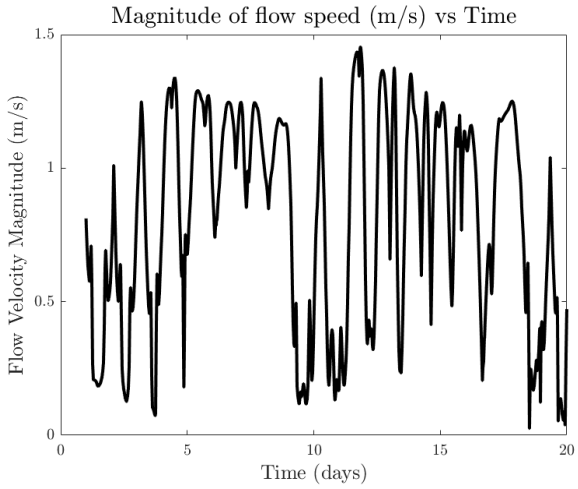

Fig. 9: Sample magnitude of flow velocity vs. time, taken from the Baseline Reward to Go MPC formulation.

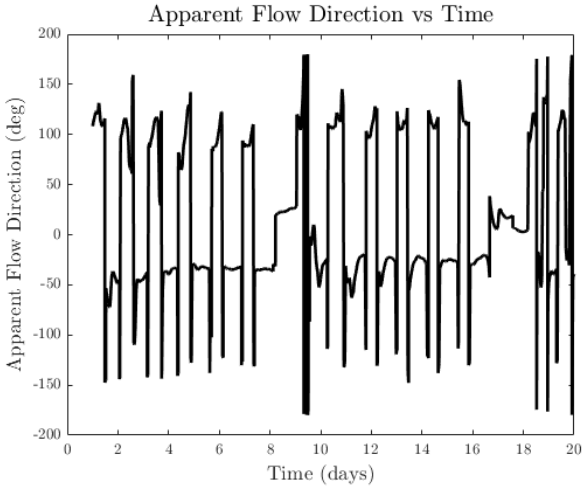

Fig. 10: Sample apparent flow direction vs. time, taken from the Baseline Reward to Go MPC formulation.

[6] S. Haines, H. Seim, and M. Muglia, "Implementing quality control of high-frequency radar estimates and application to Gulf Stream surface currents," *Journal of Atmospheric and Oceanic Technology*, vol. 34, no. 6, pp. 1207 – 1224, 2017. [Online]. Available: https://journals.ametsoc.org/view/journals/atot/34/6/jtech-d-16-0203.1.xml

[7] C. L. Gentemann, J. P. Scott, P. L. F. Mazzini, C. Pianca, S. Akella, P. J. Minnett, P. Cornillon, B. Fox-Kemper, I. Cetinić, T. M. Chin, J. Gomez-Valdes, J. Vazquez-Cuervo, V. Tsontos, L. Yu, R. Jenkins, S. D. Halleux, D. Peacock, and N. Cohen, "Saildrone: Adaptively sampling the marine environment," *Bulletin of the American Meteorological Society*, vol. 101, no. 6, pp. E744 – E762, 2020. [Online]. Available: https://journals.ametsoc.org/view/journals/bams/101/6/BAMS-D-19-0015.1.xml

[8] "SeaTrac." [Online]. Available: https://www.seatrac.com/

[9] D. Jones and G. A. Hollinger, "Planning energy-efficient trajectories in strong disturbances," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2080–2087, 2017.

[10] S. McCammon and G. A. Hollinger, "Topological path planning for autonomous information gathering," *Autonomous Robots*, vol. 45, no. 6, p. 821–842, 2021.

[11] K. Govindarajan, B. Haydon, K. Mishra, and C. Vermillion, "Coverage-maximizing solar-powered autonomous surface vehicle control for persistent gulf stream observation," in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 3675–3681.

[12] K. Chen and R. He, "Numerical investigation of the Middle Atlantic Bight shelfbreak frontal circulation using a high-resolution ocean hindcast model," *Journal of Physical Oceanography*, vol. 40, no. 5, pp. 949–964, 2010.

[13] D. P. Dee, S. M. Uppala, A. Simmons, P. Berrisford, P. Poli, S. Kobayashi, U. Andrae, M. Balmaseda, G. Balsamo, d. P. Bauer, *et al.*, "The era-interim reanalysis: Configuration and performance of the data assimilation system," *Quarterly Journal of the royal meteorological society*, vol. 137, no. 656, pp. 553–597, 2011.