Fusion of Indirect Methods and Iterative Learning for Persistent Velocity Trajectory Optimization of a Sustainably Powered Autonomous Surface Vessel

Kavin M. Govindarajan, Devansh R Agrawal, Dimitra Panagou, and Chris Vermillion

***Abstract*— In this paper, we present the methodology and results for a real-time velocity trajectory optimization for a solar-powered autonomous surface vessel (ASV), where we combine indirect optimal control techniques with iterative learning. The ASV exhibits cyclic operation due to the nature of the solar profile, but weather patterns create inevitable disturbances in this profile. The nature of the problem results in a formulation where the satisfaction of pointwise-in-time state of charge constraints does not generally guarantee persistent feasibility, and the goal is to maximize information gathered over a very long (ultimately persistent) time duration. To address these challenges, we first use barrier functions to tighten pointwise-in-time state of charge constraints by the minimal amount necessary to achieve persistent feasibility. We then use indirect methods to derive a simple switching control law, where the optimal velocity is shown to be an undetermined constant value during each constraint-inactive time segment. To identify this optimal constant velocity (which can vary from one segment to the next), we employ an iterative learning approach. The result is a simple closed-form control law that does not require a solar forecast. We present simulation-based validation results, based on a model of the SeaTrac SP-48 ASV and solar data from the North Carolina coast. These simulation results show that the proposed methodology, which amounts to a closed-form controller and simple iterative learning update law, performs nearly as well as a model predictive control approach that requires an accurate future solar forecast and significantly greater computational capability.**

I. INTRODUCTION

The collection of oceanographic data supports numerous applications: ocean currents inform marine energy site selection [1], surface temperature data enhances weather forecasts [2], and salinity measurements improve understanding of climate processes [3]. Current observation methods include moored and boat-mounted sensors [4], [5], high-frequency radar [6], and undersea gliders [7], [8], but these platforms typically yield sparse or short-duration measurements.

To obtain spatially granular and long-duration oceanographic data, these traditional methods must be augmented with persistent and autonomous observation systems. Renewably powered marine robots, such as sailing drones [9] and solar-powered autonomous surface vessels (ASVs) [10], offer this capability by combining mobility with long-duration operation.

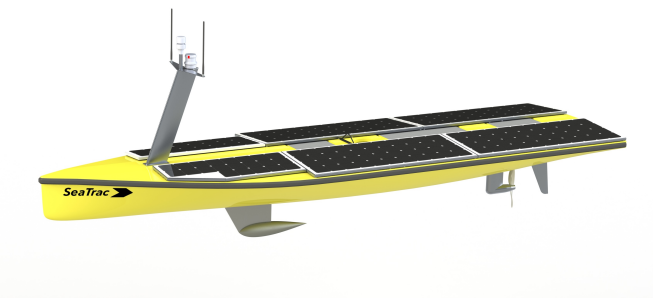

Fig. 1. SeaTrac ASV considered in this work [10]. Image used with permission.

In this work, we consider the long-duration deployment of a solar-powered ASV (shown in Fig. 1) in a dynamic ocean environment. We treat distance traveled as a proxy for information, and aim to compute a velocity trajectory that maximizes this distance over a mission horizon, subject to energetic constraints imposed by battery capacity and time-varying solar input. Our previous work addressed this using a model predictive control (MPC) framework [11], which approximated the infinite-horizon objective with a finite-horizon optimization and a heuristic terminal reward. However, tuning the reward function was largely trial-and-error, performance degraded with solar forecast errors, and the real-time computational load was significant, especially for stochastic MPC formulations.

To address these limitations, we adopt indirect methods from optimal control theory to derive necessary conditions for velocity trajectories that are provably optimal, even under time-varying solar resource availability and battery state-of-charge (SOC) constraints. This study has two main goals: (1) to use insight from Pontryagin's Minimum Principle (PMP) to interpret and benchmark MPC behavior, and (2) to derive a control strategy that approximates MPC performance while eliminating forecast dependence and significantly reducing tuning and computational demands.

Prior work in transportation systems, particularly in hybrid and electric vehicles, has explored finite-horizon energetic optimization [12], often using PMP [13]. However, these

Supported by NSF Award Number 2223844.

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

Devansh R Agrawal is a Graduate Research Assistant with the Department of Aerospace Engineering, University of Michigan. Dimitra Panagou is an Associate Professor with the Department of Robotics and the Department of Aerospace Engineering, University of Michigan, Ann-Arbor, MI, 48109, USA. devansh@umich.edu, dpanagou@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

methods generally result in boundary-value problems involving unknown costates that are difficult to resolve. For cyclically operating systems, iterative learning control (ILC) has been applied to estimate these parameters [14].

Given the diurnal periodicity of solar energy, a PMP+ILC approach is well-suited to ASV trajectory optimization. However, unlike cyclic systems in prior work, our problem features day-to-day variability in solar input and requires a notion of persistent feasibility, where satisfying SOC constraints at time t guarantees feasibility at all future times. This critical issue is not addressed in existing literature.

To this end, we first introduce minimally tightened SOC constraints using a barrier function approach that guarantees persistent feasibility across time. We then apply PMP to derive necessary conditions for optimality, revealing that the optimal velocity is piecewise constant whenever SOC constraints are inactive. Since PMP does not directly yield the value of this constant velocity, which can vary across constraint-inactive intervals, we propose an ILC-based method to iteratively estimate it.

Using a high-fidelity model of the SeaTrac SP-48 ASV, we benchmark our method against three strategies: (1) a constant velocity that ensures energy balance over the mission, (2) the MPC approach developed in [11], and (3) an upper-bound benchmark assuming no SOC constraints. Simulations using real solar irradiance data from North Carolina show that our ILC-based controller achieves comparable performance to MPC without relying on forecasts or incurring heavy computation.

In summary, our contributions are:

- A barrier function formulation to derive minimally tightened SOC constraints that guarantee persistent feasibility;
- An indirect optimal control formulation that yields a piecewise constant optimal control law;
- An iterative learning approach to estimate the optimal control input in unconstrained intervals;
- A comprehensive comparison with constant-velocity and MPC-based strategies in a realistic simulation environment.

II. MODEL & OPTIMIZATION PROBLEM

A. Optimal Control Formulation

The ASV velocity control problem to maximize distance traveled is posed as a fixed-final-state, free-final-time optimal control problem:

$$\min_{u(t)} J = \int_0^{t_f} -u(t)\, dt \tag{1}$$

subject to:

$$\dot{b}(t) = P_{in}(t) - k_h - k_m u(t)^3 \tag{2a}$$
$$b(t_f) = b(0) \tag{2b}$$
$$b_{min} \leq b(t) \leq b_{max} \qquad t \in [0, t_f] \tag{2c}$$
$$u_{min} \leq u(t) \leq u_{max} \qquad t \in [0, t_f] \tag{2d}$$

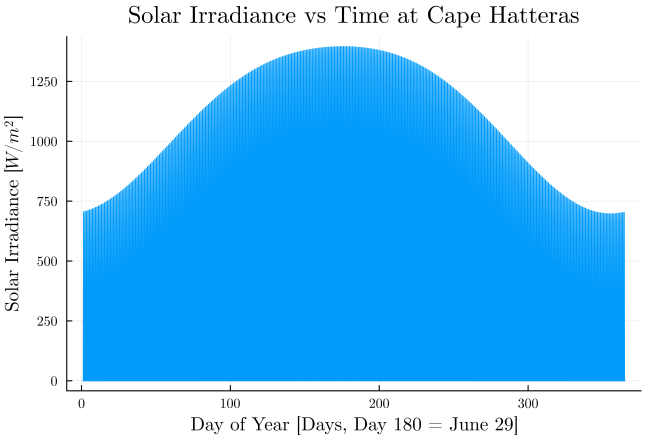

Fig. 2. Idealized solar irradiance profile from Eqn. 3

Here, $u(t)$ is the ASV velocity, $b(t)$ is the battery state of charge (SOC), and $P_{in}(t)$ is the incoming solar energy. The constants k_h and k_m represent hotel load and motor energy consumption, respectively. Limits b_{min}, b_{max} and u_{min}, u_{max} define feasible SOC and velocity ranges.

All simulations use parameters from the SeaTrac SP-48 vehicle (Table I).

TABLE I
SEATRAC SP-48 PARAMETERS

Variable Name	Symbol	Value	Units
Hotel Load	k_h	10	W
Motor Constant	k_m	83	$kg \cdot m^{-1}$
Minimum SOC	b_{min}	0	Wh
Maximum SOC	b_{max}	6500	Wh
Minimum Speed	u_{min}	0	$m \cdot s^{-1}$
Maximum Speed	u_{max}	2.315	$m \cdot s^{-1}$

B. Solar Irradiance Model

We evaluate performance under both idealized and real solar irradiance profiles. The idealized model enables analysis in a controlled setting, while the real data assesses robustness in real-world conditions.

The idealized irradiance model, adapted from [15], is:

$$P_{in}(t) = \max\left(0, D_0(t) + D_1(t)\cos\left(\frac{2\pi t}{T}\right)\right) \tag{3}$$

where $D_0(t)$ is the average irradiance, $D_1(t)$ is the oscillation amplitude, and T (typically 24 hours) is the period. Both $D_0(t)$ and $D_1(t)$ depend on latitude and time of year. For simplicity, we fix the ASV along the latitude of Cape Hatteras, North Carolina. A sample profile is shown in Fig. 2.

However, this model does not account for variability due to weather (e.g., cloud cover). To capture real-world disturbances, we also use ERA5 reanalysis data for solar irradiance at Cape Hatteras in 2022 [16], shown in Fig. 3.

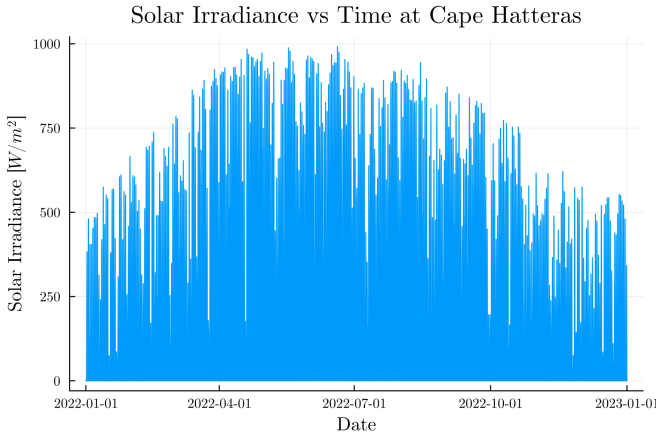

Fig. 3. Measured solar irradiance at Cape Hatteras in 2022

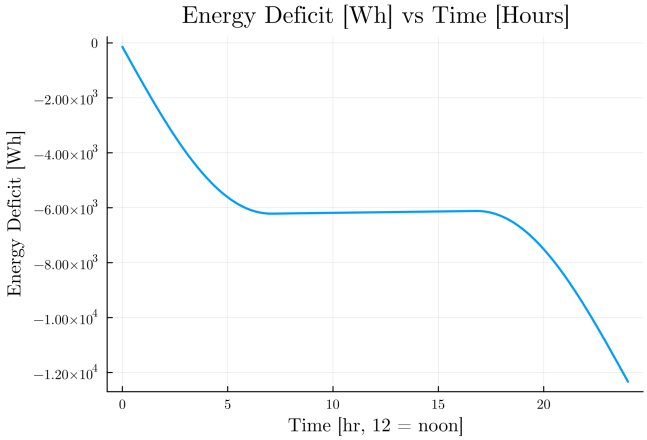

Fig. 4. Energy deficit curve

III. OPTIMAL CONTROL FORMULATION: BARRIER FUNCTIONS & INDIRECT METHODS

We derive a simple expression for the optimal velocity trajectory in two steps. First, we construct barrier functions that minimally tighten the SOC constraints to ensure persistent feasibility. Second, we apply indirect methods to derive necessary conditions showing that the optimal velocity is constant during each constraint-inactive interval.

A. Barrier Functions for Persistent Feasibility

The constraints in Eqns. 2a–2d do not guarantee that satisfying SOC constraints at time t ensures feasibility at later times. For example, if $b(t) = b_{min}$ and solar input drops to zero, the constant hotel load k_h will cause SOC to fall below b_{min}, violating the state constraint.

To guarantee *persistent feasibility*, we reformulate the state constraints into time-varying barrier functions.

Starting with the lower SOC bound, we define the *energy deficit* assuming $u(t) = 0$:

$$\epsilon_-(t) = \int_{t_0}^{t} (k_h - P_{in}(\tau))\, d\tau \tag{4}$$

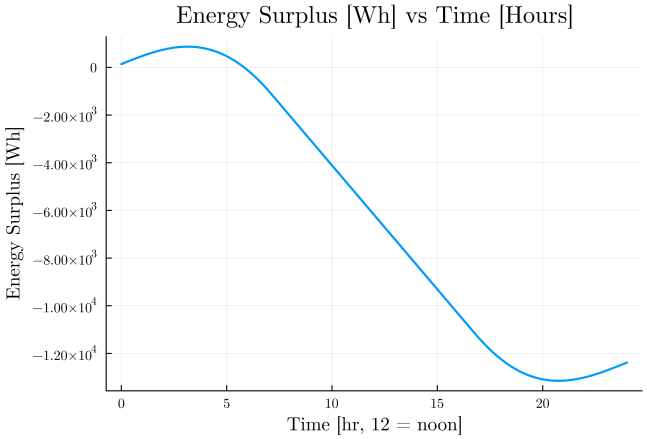

Fig. 5. Energy surplus curve

We then define the lower SOC barrier as:

$$\epsilon_-^{\dagger}(t) = \epsilon_-(t) - \epsilon_-(t_1) \tag{5}$$

$$b_l(t_1) = \begin{cases} \sup\limits_{t_2 \in [t_1, t_f]} \epsilon_-^{\dagger}(t_2), & \text{if } \sup > 0 \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

where t_1 is the current time at which we are evaluating the barrier function, representing the "starting point", and $t_2 \in [t_1, t_f]$ is a time representing some time in the future.

For the upper SOC bound, we define the *energy surplus* assuming $u(t) = u_{max}$:

$$\epsilon_+(t) = \int_{t_0}^{t} (P_{in}(t) - k_h - k_m u_{max}^3)\, dt \tag{7}$$

Then, the upper SOC barrier is:

$$\epsilon_+^{\dagger}(t) = \epsilon_+(t) - \epsilon_+(t_1) \tag{8}$$

$$b_u(t_1) = \begin{cases} \sup\limits_{t_2 \in [t_1, t_f]} \epsilon_+^{\dagger}(t_2), & \text{if } \sup > 0 \\ 0, & \text{otherwise} \end{cases} \tag{9}$$

From Eqns. 6 and 9, we define the time-varying SOC constraint:

$$b_l(t) \le b(t) \le b_u(t) \tag{10}$$

B. Indirect Methods for Deriving the Optimal Velocity Profile

With the reformulated constraints in place, we apply indirect optimal control to derive a key structural result.

Lemma 1: When $b_l(t) < b(t) < b_u(t)$ (i.e., constraints are inactive), the optimal velocity $u^*(t)$ is constant.

Proof: We augment the system to include a penalty state $x_2(t)$ for constraint violations:

$$\begin{aligned} \dot{\vec{x}} &= \begin{bmatrix} \dot{b}(t) \\ \dot{x}_2(t) \end{bmatrix} \\ &= \begin{bmatrix} P_{\text{in}}(t) - k_h - k_m u(t)^3 \\ b^2\, \mathbb{1}(-b) + (b_{\max} - b)^2\, \mathbb{1}(b_{\max} - b) \end{bmatrix} \end{aligned} \tag{11}$$

where $\mathbb{1}$ is the Heaviside step function.

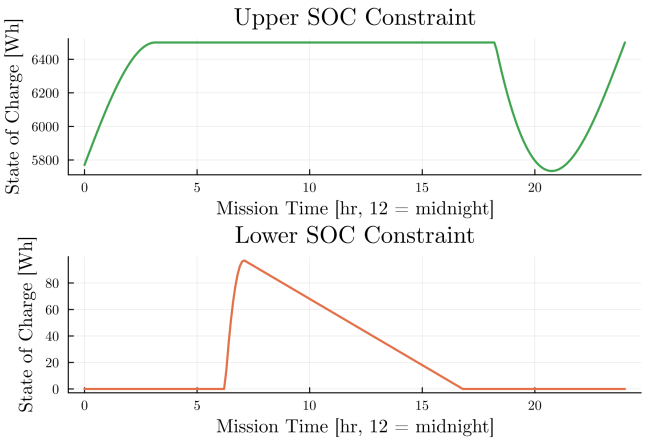

Fig. 6. SOC barrier functions define time-varying upper/lower SOC bounds guaranteeing persistent feasibility.

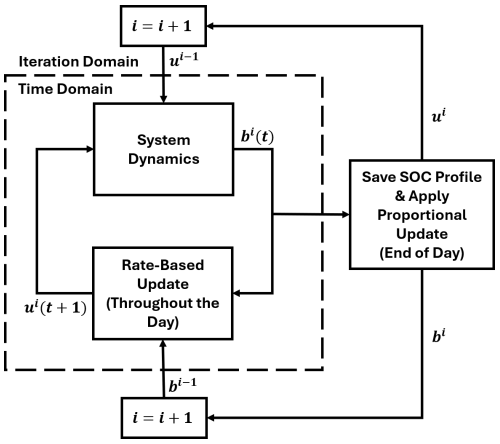

Fig. 7. Block diagram demonstrating the iterative approach to learning the optimal constant speed.

The corresponding Hamiltonian is:

$$H(t,x,u,p) = -u + p_1(P_{in} - k_h - k_m u^3) \quad (12)$$
$$+p_2((b_u - b)^2 \mathbb{1}(b - b_u) + (b - b_l)^2 \mathbb{1}(b_l - b))$$

where p_1, p_2 are the co-states. The co-state dynamics follow:

$$\dot{p}_1 = -2p_2(b_u - b)\mathbb{1}(b - b_u) + p_2(b_u - b)^2\delta(b - b_u)$$
$$+ 2p_2(b - b_l)\mathbb{1}(b_l - b) + p_2(b - b_l)^2\delta(b_l - b) \quad (13)$$
$$\dot{p}_2 = 0 \quad (14)$$

When constraints are inactive, all $\mathbb{1}$ and δ terms vanish, so $\dot{p}_1 = 0$ and p_1 is constant. Setting $\frac{\partial H}{\partial u} = 0$ yields:

$$u^* = \begin{cases} u_{max}, & b \geq b_u(t) \\ u_{min}, & b \leq b_l(t) \\ \sqrt{\frac{-1}{3k_m p_1}}, & \text{otherwise} \end{cases} \quad (15)$$

Since p_1 is constant when constraints are inactive, u^* must also be constant, completing the proof. ■

Eqn. 15 and Lemma 1 yield a compact characterization of the optimal control law. However, p_1 (and thus the optimal velocity) can vary between constraint-inactive intervals and is not directly known. In the next section, we describe an ILC-based method to estimate p_1.

IV. REAL-TIME REALIZATION OF THE OPTIMAL VELOCITY PROFILE THROUGH ITERATIVE LEARNING

The previous section has established that the optimal velocity control strategy is a switching one, where the optimal velocity is constant during each constraint-inactive interval. Two complications arise from the aforementioned analysis. First, the actual controller must be implemented in sampled (discrete) time, in the presence of sensor noise. In such a scenario, a switching controller is prone to significant chatter, which must be mitigated for the resulting controller to be practical. Secondly, while the previously derived theory establishes that the optimal velocity is constant during each constraint-inactive interval, it does not establish the optimal value of that constant velocity. This section addresses both of these complications.

A. Discretization

When we discretize time, the switching controller will chatter as the SOC barrier functions go from active to inactive (and vice versa). This chatter will be further complicated in actual experimental implementations where the SOC measurement is corrupted by noise. To address these complications, we modify the control solution to include the "buffer" shown in Eqn. 16:

$$u(t) = \begin{cases} \frac{b(t)-b_l(t)}{\delta} u^*(t) + \left(1 - \frac{b(t)-b_l(t)}{\delta}\right) u_{min} \\ \qquad \text{when } 0 < b(t) - b_l(t) < \delta \\ \frac{b_u(t)-b(t)}{\delta} u^*(t) + \left(1 - \frac{b_u(t)-b(t)}{\delta}\right) u_{max} \\ \qquad \text{when } 0 < b_u(t) - b(t) < \delta \end{cases} \quad (16)$$

where δ is a tunable parameter that represents the width of the buffer in Watt-hours and u^* represents the optimal unconstrained velocity.

B. Iterative Learning

To compute the optimal value of p_1 (which is known to be constant during each constraint-inactive interval but otherwise unknown), and thus the value of the commanded velocity, we utilize an iterative learning approach as depicted in Fig.7. To do this, we begin with an initial guess of p_1, then update the estimate for this co-state value as per the update laws defined in Eqns. 17 and 18. We define each iteration i to be a 24-hour period (one solar cycle). The update law is given in two parts. The first of these parts is a once-per-iteration update law given by:

$$u^{i+1} = u^i + k_p(b^i(t_f) - b^i_{des}) \quad (17)$$

where k_p is a proportional learning gain that acts upon the difference between the SOC at the end of iteration (represented by $b^i(t_f)$) and the desired terminal state of

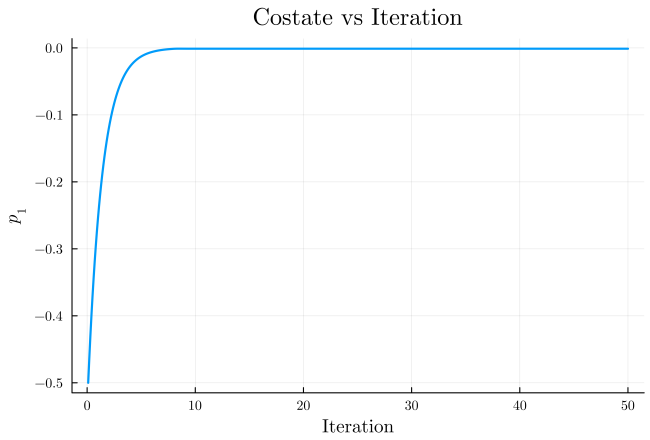

Fig. 8. Costate vs iteration under a consistent day-to-day solar irradiance profile. This simulation demonstrates convergence of the ILC algorithm under consistent solar conditions.

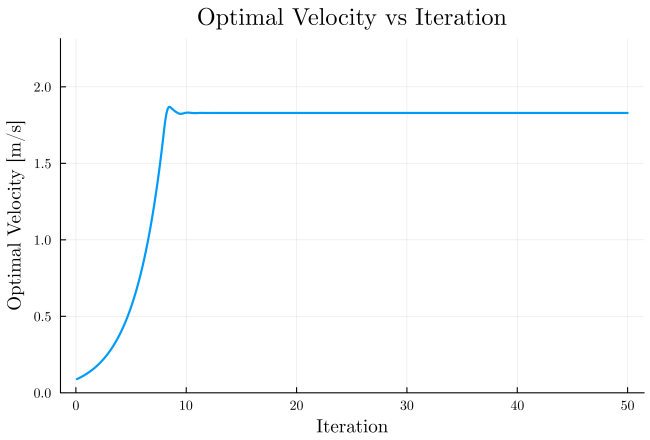

Fig. 9. Optimal Velocity vs iteration under a consistent day-to-day solar irradiance profile. This simulation demonstrates convergence of the ILC algorithm under consistent solar conditions.

charge (b^i_{des}). This update is applied at the end of each iteration (day).

The second part of the update law represents a continuous rate-based update law that is given by:

$$u^i(t+1) = u^i + k_d(b^i(t) - b^{i-1}(t)) \tag{18}$$

where k_d is a rate-based gain that acts upon the difference between the current SOC and the SOC at the same time in the previous iteration (day). This update is then applied to the velocity from the same time in the previous iteration. This update law is applied continually following the first iteration and serves as a "damping" term on the control updates.

To demonstrate the convergence of the above strategy to an optimal costate value p_1, we apply the strategy to a simulation where the environmental conditions are repeated across each iteration. The results of this simulation are shown in Fig. 8, and the velocities corresponding to the costate values are shown in Fig.9. From these figures, we can see that the costate converges to the optimal value in under 15 iterations for $k_p = 5\times10^{-5}$ and $k_d = 1\times10^{-5}$. This optimal costate value is $p_1 = -0.0012$, which corresponds to a speed of $u = 1.83\ m\cdot s^{-1}$.

In the following section, we demonstrate the performance of the proposed control strategy against other comparison strategies in simulation.

V. RESULTS

To benchmark the performance of the proposed controller, we compare it against the performance of other strategies in simulation.

A. Comparison Strategies

As an initial basis for comparison, we compute the constant velocity for which the energy expended over a one-year simulation duration is exactly equal to the solar energy available. To do this, we solve the below equation for u:

$$\int_0^{t_f} k_m u^3 dt = \int_0^{t_f} P_{in}(t)dt \tag{19}$$

where k_m represents a gain that represents the power consumed by the ASV's motor, u is the constant velocity of the ASV, and $P_{in}(t)$ is the energy produced by the ASV from the solar resource from times 0 to t_f. We allow the ASV to travel at this velocity without being subject to the SOC constraints to establish the upper limit on ASV performance. Note that this is not in fact a tight upper bound, as the ASV will be subject to SOC constraints in reality.

As a second comparison control strategy, we implement the aforementioned constant-velocity strategy only when the tightened SOC constraints are inactive. During periods in which the previously computed constant velocity would lead to violation of tightened SOC constraints, the velocity is set to the lower or upper limit in order to satisfy the constraint at equality.

Finally, our indirect methods + ILC strategy and the two comparison strategies described above are compared with an MPC strategy as implemented in [11].

To evaluate these strategies, we considered two simulation scenarios, where we seek to maximize the distance traveled by the ASV over a mission period of one year. In the first scenario, the idealized solar irradiance model was used as the energetic resource. In the second scenario, ERA5 data for Cape Hatteras was used [16].

B. Simulation Results

For the simulations performed under the idealized solar resource profile, from Fig. 10, we can see that the proposed learning controller outperforms the constant velocity benchmarks, while nearly matching the performance of MPC. It is noteworthy that the MPC implementation is based on a solar resource forecast, and this forecast is assumed to be perfect for the purpose of optimization. The ILC-based approach does not require any forecast and performs almost identically under the idealized solar resource model. The daily-averaged velocity and SOC are plotted against time in Figs. 11 and 12. Figs. 11 and 12 illustrate an increased level of volatility in the ILC approach, as ILC attempts to learn the optimal

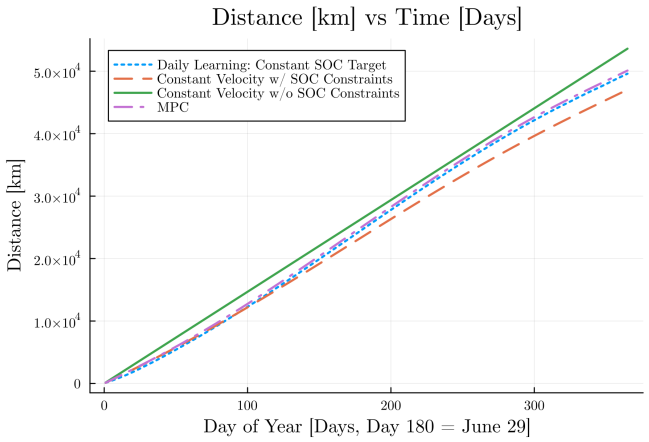

Fig. 10. Distance traveled vs. time under the idealized solar irradiance model.

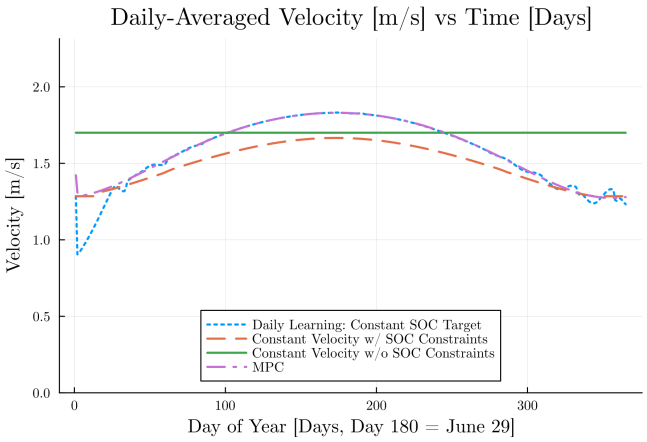

Fig. 11. Velocity vs. time under the idealized solar irradiance model.

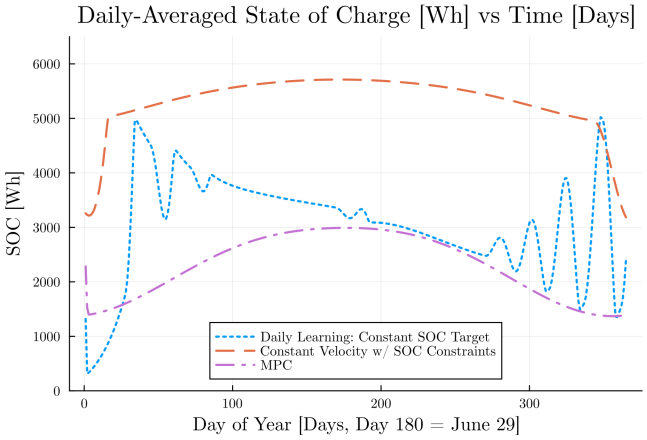

Fig. 12. SOC vs. time under the idealized solar irradiance model.

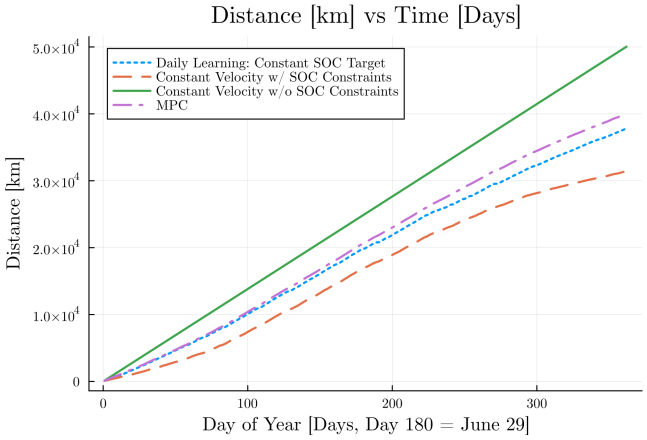

Fig. 13. Distance traveled vs. time under the actual solar irradiance data for Cape Hatteras.

velocity. This volatility, however, is not associated with any appreciable performance reduction.

Similar results can be seen when the algorithms are applied to the real-world solar irradiance profile. From Fig. 14, we can see that the ILC-based strategy is particularly sensitive to the fluctuations in solar irradiance that are present in the real-world data. This is attributable to the fact that the estimated optimal velocity value requires time to adjust to day-to-day variations in solar irradiance due to atmospheric conditions, whereas the MPC strategy has the benefit of accessing a perfect forecast. Accordingly, the state of charge is particularly volatile as it reacts to the changing velocity and solar profile as seen in Fig.15. Nevertheless, the ILC approach proposed in this work nearly matches the MPC performance, even in the presence of real-world solar fluctuations.

The ILC approach falls just short of the performance levels yielded by MPC, yet it possesses tremendous advantages. First, it requires no forecast of the solar resource. Secondly, the ILC approach does not require the heuristic tuning of a terminal reward function within a real-time optimization. Finally, the learning controller is simpler to implement and requires less computational power. This is due to the fact that the learning controller only has two gains that must be tuned for performance, whereas the MPC requires a detailed model.

The presented strategies were implemented in Julia [17] and simulated on a computer with an Intel®Xeon®W-2125 4.0 GHZ CPU and 32GB of RAM. For the same year-long simulation at a time step of 6 minutes, the ILC approach took roughly 20 seconds to simulate, whereas the MPC strategy (implemented using the JuMP Ipopt optimizer [18]) took 21 hours and 36 minutes. This demonstrates a significant improvement in computational efficiency using the presented ILC approach.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we considered the problem of long-horizon or persistent velocity trajectory optimization for a solar-powered autonomous surface vessel (ASV). Because of the

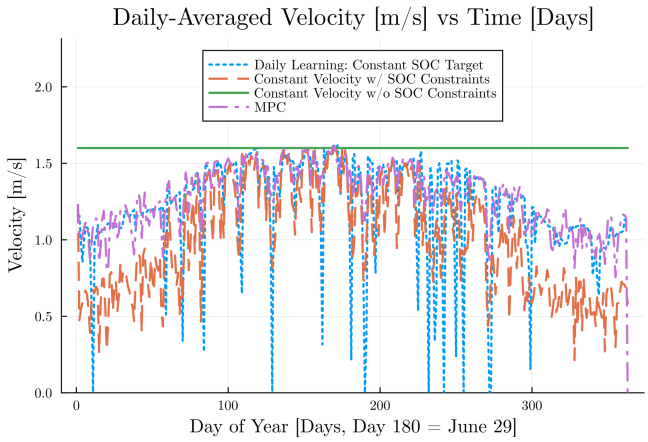

Fig. 14. Velocity vs. time under the actual solar irradiance data for Cape Hatteras.

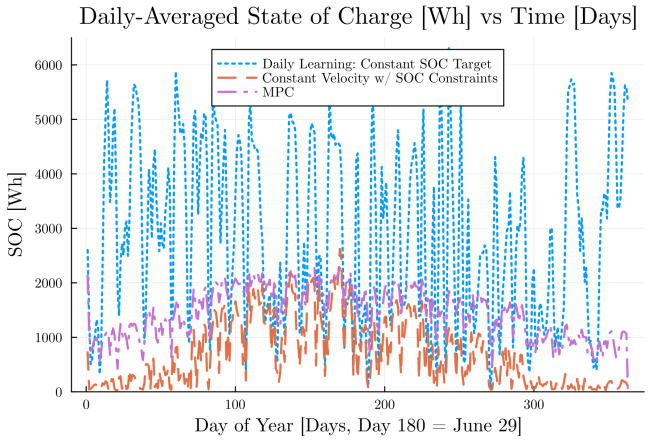

Fig. 15. SOC vs. time under the actual solar irradiance data for Cape Hatteras.

unique features of this system that do not immediately guarantee persistent feasibility of pointwise-in-time state constraints, we constructed barrier functions to tighten these constraints by the minimum amount required to ensure persistent feasibility. We then utilized indirect methods to arrive at a provably optimal switching control law, where we subsequently used iterative learning control (ILC) to obtain estimates of the optimal constant velocity during each constraint-inactive interval. Using real solar data and a model of the SeaTrac SP-48 ASV, we showed that this formulation outperforms simple benchmark strategies. Furthermore, we showed that the proposed approach nearly matches the performance of a more complex MPC strategy that requires a solar forecast while significantly reducing computational complexity.

In future work, we will focus on replacing our distance-maximization objective with an actual information-maximization mechanism. Furthermore, we will focus on combining velocity trajectory optimization with path planning, all with the goal of maximizing information gathered. Finally, we will focus on experimentally validating the developed control algorithms on the SP-48 ASV presently in our possession.

References

[1] K. Naik and C. Vermillion, "Integrated physical design, control design, and site selection for an underwater energy-harvesting kite system," Renewable Energy, vol. 220, p. 119687, 2024.

[2] M. N. Huda, "Machine learning for improvement of ocean data resolution for weather forecasting and climatological research," Ph.D. dissertation, Virginia Tech, 2023.

[3] H. E. Seim, D. Savidge, M. Andres, J. Bane, C. Edwards, G. Gawarkiewicz, R. He, R. E. Todd, M. Muglia, J. Zambon et al., "Overview of the processes driving exchange at cape hatteras program." Oceanography, vol. 35, no. 2, 2022.

[4] 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.

[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.

[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] 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.

[8] 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.

[9] 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., "Adaptively sampling the marine environment," 2020.

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

[11] K. Govindarajan, B. Haydon, and C. Vermillion, "Predictive velocity trajectory control for a persistently operating solar-powered autonomous surface vessel," in 2023 American Control Conference (ACC), 2023, pp. 2077–2083.

[12] M. Miyatake and K. Matsuda, "Energy saving speed and charge/discharge control of a railway vehicle with on-board energy storage by means of an optimization model," IEEJ Transactions on Electrical and Electronic Engineering, vol. 4, 2009. [Online]. Available: https://api.semanticscholar.org/CorpusID:110581908

[13] H. Abbas, Y. Kim, J. B. Siegel, and D. M. Rizzo, "Synthesis of pontryagin's maximum principle analysis for speed profile optimization of all-electric vehicles," Journal of Dynamic Systems, Measurement, and Control, vol. 141, no. 7, p. 071004, 2019.

[14] J. Hu, Y. Shao, Z. Sun, M. Wang, J. Bared, and P. Huang, "Integrated optimal eco-driving on rolling terrain for hybrid electric vehicle with vehicle-infrastructure communication," Transportation Research Part C: Emerging Technologies, vol. 68, pp. 228–244, 2016.

[15] R. L. Hulstrom, Insolation Models and Algorithms, 2003, pp. 61–141.

[16] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers et al., "The era5 global reanalysis," Quarterly Journal of the Royal Meteorological Society, vol. 146, no. 730, pp. 1999–2049, 2020.

[17] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," SIAM review, vol. 59, no. 1, pp. 65–98, 2017. [Online]. Available: https://doi.org/10.1137/141000671

[18] I. Dunning, J. Huchette, and M. Lubin, "Jump: A modeling language for mathematical optimization," SIAM Review, vol. 59, no. 2, pp. 295–320, 2017.