Spatiotemporal Path Planning in Strong, Dynamic, Uncertain Currents

David R. Thompson[1,2], Steve Chien[2], Yi Chao[2], Peggy Li[2], Bronwyn Cahill[3], Julia Levin[3], Oscar Schofield[3], Arjuna Balasuriya[4], Stephanie Petillo[4], Matthew Arrott[5], Michael Meisinger[5]

***Abstract*— This work addresses mission planning for autonomous underwater gliders based on predictions of an uncertain, time-varying current field. Glider submersibles are highly sensitive to prevailing currents so mission planners must account for ocean tides and eddies. Previous work in variable-current path planning assumes that current predictions are perfect, but in practice these forecasts may be inaccurate. Here we evaluate plan fragility using empirical tests on historical ocean forecasts for which followup data is available. We present methods for glider path planning and control in a time-varying current field. A case study scenario in the Southern California Bight uses current predictions drawn from the Regional Ocean Monitoring System (ROMS).**

I. INTRODUCTION

We consider path planning for autonomous underwater gliders based on advance forecasts of an uncertain, time-varying current field. The system we describe will be tested as part of the Ocean Observatories Initiative (OOI), a science-driven oceanic sensor network. The OOI will integrate moorings, radar and Autonomous Underwater Vehicles (AUVs), assimilating real-time data into regional forecasting models. The initial tests of this system will use underwater gliders, autonomous submersibles designed for long duration missions [4]. Gliders travel efficiently by changing buoyancy and using winged surfaces to produce a sawtooth trajectory (Figure 1). Gliders can travel for months on a single battery charge, pausing at intervals on the surface to transmit science data and receive commands.

Gliders are highly sensitive to prevailing currents. Coastal current forces are often larger than the propulsion force of the glider, and can change on time scales of hours. Effective mission planning must exploit these forces by placing the glider strategically to exploit beneficial tides and eddies. Fortunately ocean models such as the Regional Ocean Monitoring System, or ROMS [1], [2], [11] can now produce current forecasts at time scales and resolutions sufficient to inform glider plans. While a growing body of work addresses path planning in a static current field [9], [5], [7], currents in the glider scenario vary on time scales much shorter than the mission duration. Dynamic currents of the Slocum glider domain demand true spatiotemporal path planning. Advances in this area include work by Zhang et al. [6] who optimize paths represented as continuous trajectories, and Soulignac et al. whose wavefront method propagates path cost as a function of departure time [8].

[1] david.r.thompson@jpl.nasa.gov. [2]Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109. [3] Rutgers University, 71 Dudley Road, New Brunswick, NJ 08901. [4] Center for Ocean Engineering, Massachusetts Institute of Technology, 77 Massachusetts Ave. Cambridge, MA 02139. [5] Calit2, University of California, San Diego, 9500 Gilman Drive #0436, La Jolla, CA 92093.
Copyright 2010 California Institute of Technology. All Rights Reserved.
Government support acknowledged.

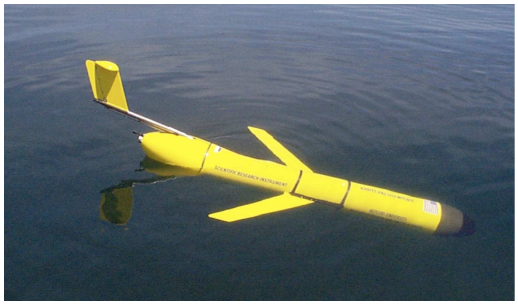

Fig. 1. Slocum Glider submersible, courtesy Rutgers University.

These approaches find optimal paths to a given spatial location. However, our deployments will also require that gliders be present at a specific time. Specifying *spatiotemporal* goals is important both for tracking time-varying ocean phenomena [10], [12] and for time-coordination of multiple assets within the OOI. Early arrival will not suffice if the glider cannot hold position against ocean currents. This requires a new approach to path planning in time-varying current fields. Recently we presented a wavefront algorithm for path planning in time varying currents based on an *earliest valid arrival* criterion. Our approach minimizes the time of travel to the goal position under the constraint that the glider is then capable of holding position against currents until the desired time is reached [13]. This permits goals specified as arbitrary spatiotemporal volumes.

Previous tests of variable-current path planning have assumed that planning occurs with perfect knowledge of future currents. In practice it is difficult to forecast currents accurately and actual conditions can quickly diverge. This work evaluates performance by simulating realistic discrepancies between planning and runtime currents. We use historical ocean current predictions for which followup data is available, providing an advance current prediction and the actual conditions. Our study considers a case scenario using ROMS predictions in the Southern California Bight. We present our path planning methodology and some simple control methods for tracking the path. Simulations compare current-agnostic and current-sensitive mission planning using both perfect and realistic forecasts. Finally, we conclude with some quantitative guidelines for ocean modeling accuracy and some qualitative recommendations for mission planning in the future OOI.

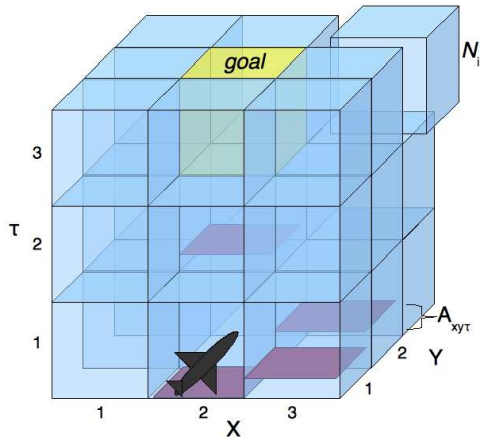

Fig. 2. The 3D spatiotemporal grid. The glider begins in the lower center and travels a path that visits 4 nodes before arriving at the goal location (yellow node at top). The red rectangles show $\mathcal{A}_{xyt}$, the earliest valid arrival time for each node's spatial location. Arrival times are computed for all nodes, but for clarity we hide arrival times for nodes not participating in the optimal path.

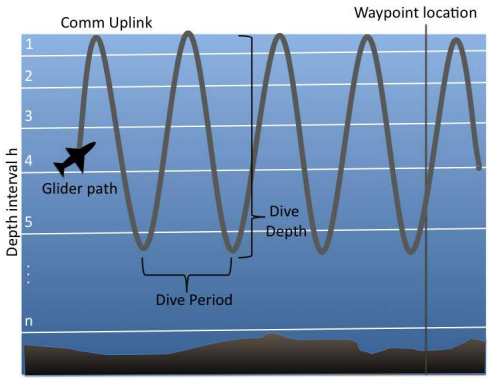

Fig. 3. Dive profile for our simplified glider motion model using current predictions at different depths. GPS updates are extended pauses at the surface that occur at regular time intervals.

II. PATH PLANNING APPROACH

We use ROMS forecast models provide to provide a depth-averaged current estimate $C_{xy\tau}$ for an ocean volume with discrete grid squares indexed by time intervals τ, latitude x, longitude y. A glider path is comprised of *segments* traveling between adjacent nodes in the 3D grid. Segment endpoints take specific continuous time values in their nodes' intervals. We treat the temporal and spatial dimensions differently: segment endpoints must lie exactly on the physical locations represented by grid points, but they can lie anywhere within the continuous time interval covered by the grid square. Figure 2 illustrates a simple four-segment path through the spatiotemporal grid. The glider begins in the lower center and travels a path to the goal location represented by the yellow grid square at the top. The final segment of this trip consists of a *station keeping* action; the glider waits at the same location until it enters the time interval of the destination node.

A. Travel time estimates

Path planning requires that we calculate the travel time between adjacent grid squares. In practice forward velocity is affected by dynamics of the dive trajectory and periodic pauses for GPS acquisition. Moreover, the glider may deviate from the straight-line path between grid squares since it can only correct its heading at each GPS waypoint. Over long distances these effects average out, so this high-level path planner assumes fixed flight parameters yielding a constant still-water velocity.

The path planner assumes that the glider will choose a control propulsion that combines with the prevailing current to yield a velocity in the desired direction of travel. We define 2-vectors corresponding to control propulsion $v_{control}$, a constant local current with velocity $C_{xy\tau}$, and a resulting net velocity v_{actual}. The control velocity must have a magnitude no greater than the nominal velocity of the vehicle in still water, so that $|v_{control}| \leq m$. The net velocity has magnitude λ and follows the desired direction of travel d:

$$v_{actual} = v_{control} + C_{xy\tau} = \lambda d \tag{1}$$

We choose the propulsion velocity $v_{control}$ to maximize travel speed, producing the largest possible motion in the desired direction of travel:

$$\lambda = \max_{v_{control}} \quad |v_{control} + C_{xy\tau}| \tag{2}$$

A path between adjacent nodes is considered impossible if the glider would require longer than a single time interval to travel between them. Otherwise, the transit time between adjacent nodes is determined by the net velocity λd.

Note that ROMS models actually supply depth-variable current predictions $C_{xy\tau h}$ with depth indices h. We must take these into account as well since the glider visits many depths during travel. The current predictions vary in vertical resolution from $5m$ near the ocean surface down to $500m$ resolutions at depth. The glider follows a sinusoidal path with maximum dive depth determined by the user-defined maximum, the vehicle limit, or the sea floor (Figure 3).

The dive period is determined by flight parameters such as pump displacement and the dive and climb angles. In this study the vertical flight profile is taken as fixed; our planner controls the glider's heading but not its depth. The fraction of time spent at each depth level approaches a known quantity over long distances, so we integrate currents over the depth dimension to estimate straight-line travel times between latitude/longitude positions. The result is that the glider can construct plans in the 3D volume comprised by spatial dimensions x and y, and the temporal dimension τ. The total current force experienced by the glider is taken to be constant within each grid square; it is given by $C_{xy\tau}$:

$$C_{xy\tau} = \frac{1}{\sum_h f(h)} \sum_h f(h) C_{xy\tau h} \tag{3}$$

where $f(h)$ is the time in each glide cycle spent in the depth interval h. Our planner uses a constant dive angle and a maximum depth of $200m$ or the sea floor depth, whichever is less.

B. Wavefront path planning

Our path planning uses the spatiotemporal wavefront algorithm of Thompson et al. [13]. The planner uses a wavefront expansion to compute the time required to reach any location and time in a 3D spatiotemporal volume. It optimizes an *earliest valid arrival* criterion that seeks a path to reach a given spatial location as early as possible, provided that the glider can then hold position against currents until the desired time is reached. This objective favors fast travel to the goal location to provide a margin of error for recovering from execution uncertainty. At the same time it permits goals to be specified with a destination *time*.

Input: Nodes $\mathcal{N}$, Current predictions C, Start location $\mathcal{N}_{\text{start}}$
Output: Path $\mathcal{P} = \{P_i\}_{i=1}^{n}$

1 Initialize $\mathcal{A}_{\text{start}} \leftarrow 0$, all other $\mathcal{A} \leftarrow \inf$
2 Initialize wavefront queue $Q = \{\mathcal{N}_{\text{start}}\}$
3 **while** Q *not empty* **do**
4 find "parent node" $\mathcal{N}_p \in Q$ at x_p, y_p, τ_p minimizing $\mathcal{A}_p$
5 $Q \leftarrow Q \setminus \mathcal{N}_p$
6 define "hold position destination" $\mathcal{N}_h$ at $(x_p, y_p, \tau_p + 1)$
7 **if** $\mathcal{N}_h$ *reachable from* $\mathcal{N}_p$ **then**
8 $\mathcal{A}_h \leftarrow \mathcal{A}_p$
9 $\mathcal{M}_h \leftarrow \mathcal{N}_p$
10 $Q \leftarrow Q \cup \mathcal{N}_h$
11 **foreach** (x_c, y_c) *neighboring* $\mathcal{N}_p$ **do**
12 **if** (x_c, y_c) *reachable from* $\mathcal{N}_p$ **then**
13 **if** $\mathcal{A}_p \notin \tau_p$ **then**
14 $t_{\text{depart}} \leftarrow \min t : t \in \tau_p$
15 **else**
16 $t_{\text{depart}} \leftarrow \mathcal{A}_p$
17 $t_{\text{arrival}} \leftarrow t_{\text{depart}} +$ travel time (from $\mathcal{C}$)
18 $\tau_c = \tau : t_{\text{arrival}} \in \tau$
19 "child" node $\mathcal{N}_c$ at (x_c, y_c, τ_c)
20 **if** $t_{\text{arrival}} < \mathcal{A}_c$ **then**
21 $\mathcal{A}_c \leftarrow t_{\text{arrival}}$
22 $\mathcal{M}_c \leftarrow \mathcal{N}_p$
23 $Q \leftarrow Q \cup \mathcal{N}_c$
24 $\mathcal{P} \leftarrow \{\mathcal{N}_{\text{end}}\}$ for $\mathcal{N}_{\text{end}}$ minimizing $\mathcal{A}_{\text{end}}$
25 $\mathcal{N}_{\text{prev}} \leftarrow \mathcal{M}_{\text{end}}$ (parent of $\mathcal{N}_{end}$)
26 **while** $\mathcal{N}_{\text{start}} \notin \mathcal{P}$ **do**
27 $\mathcal{N}_{\text{prev}} \leftarrow M_{\text{prev}}$ (parent of $\mathcal{N}_{prev}$)
28 $\mathcal{P} \leftarrow \mathcal{P} \cup \mathcal{N}_{\text{prev}}$
29 **return** $\mathcal{P}$

Algorithm 1: Computation of arrival times and parent nodes satisfying the "earliest valid arrival" criterion. Q is a priority queue of unexpanded nodes maintained in order of increasing arrival times. Subscripts p, c, h indicate data associated with the current parent, child, and station-keeping destination nodes. After finding the set of "best parents" $\mathcal{M}$ we compute the optimal path $\mathcal{P}$ by tracing parents from the goal back to the origin.

The path planning algorithm computes $\mathcal{A} \in \mathbb{R}$ for each node; this is defined as the earliest possible time of arrival to location (x, y) for the "hold position" action ending in time interval τ. Often a node directly following another at the same physical location will take its predecessor's value of $\mathcal{A}$ due to position-holding. For example, in Figure 2 the final segment consists of a hold position action. Physical locations have not changed between the final two nodes so their earliest arrival times are identical. We record the earliest possible arrival time yet discovered for each node in the grid. A wavefront algorithm recursively expands nodes until the goal location is reached.

Our path planning strategy begins by initializing a queue Q of unexpanded reachable nodes. Initially this queue contains the starting location. Each expansion operation takes the top node $\mathcal{N}_p$ of the queue and simulates travel in each of the 8 compass directions, as well as holding position at the current node into the next time step.

If the hold position action is successful it results in the parent's time of arrival $\mathcal{A}_p$ being preserved in the same physical location at the next time interval. For travel between different physical locations we compute a travel time t based on the earliest possible time of departure from the parent node. If the parent node was itself reached by a hold position action, then it is occupied from the start of its interval. In this case the earliest possible departure t happens at the start of the time interval. We can identify this case whenever a parent is occupied prior to its own time interval, i.e. $\mathcal{A}_p \notin \tau_p$. This is a subtle point: a node's associated arrival time need not be within the time interval represented by the node. Instead it refers to the earliest valid arrival, i.e. the earliest possible arrival which then permits position holding until entering the node.

We compute travel times t from equation 2 assuming a constant local current velocity. Depending on the time required, the destination child node could turn out to be in either the parent's time interval or the following one. The resulting arrival time is the best arrival yet realized in the child node $\mathcal{N}_c$. We retain paths corresponding to the earliest possible arrival time for each child, which requires storing a single best parent $\mathcal{M}_c$ for each reachable node in the grid.

The algorithm continues recursively expanding child nodes until it has accounted for all grid locations (Algorithm 1). We expand new nodes in order of increasing arrival times which requires maintaining a sorted queue tantamount to a wavefront expansion [7] or Dijkstra method. This guarantees that discounting discretization effects, the earliest possible arrival is known for every reachable node; the wavefront solution is optimal up to discretization accuracy. After completely expanding all the nodes in the grid one can trace a valid path from any goal node back to the start using the recorded parent nodes (lines 24-29). The result of the planning procedure is a sequence of segments through adjacent 3D grid squares that is theoretically within the glider's physical capability and optimizes the earliest valid arrival criterion.

The arrival times $\mathcal{A}$ also define a *reachability function*: the locations for a particular time step that are reachable by a glider beginning at the original start location. The isocontours of this function constitute a useful data product for mission planners; they show legal locations for waypoint placement during glider planning. Figure 4 shows a glider path together with reachability isocontours.

We discounted any trial for which the currents at the

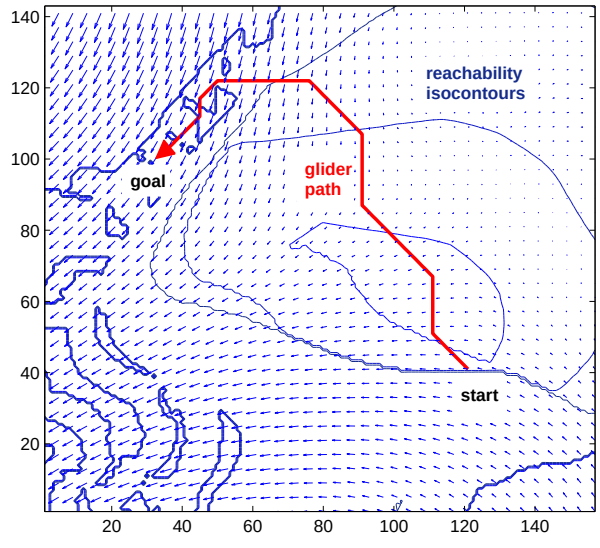

Fig. 4. Typical solution for earliest arrival showing the final simulation time step. The red path represents the optimal glider path. Isocontours of the reachability function are shown in yellow and green, with blue arrows showing the final timeslice of the time-varying current field.

endpoint were greater than the glider's own propulsion. These scenarios resulted in attempts to pass through the target node at speed, by first lingering in calmer waters and then dashing through the target at the appropriate time. Naturally these plans were quite fragile, but requiring at least one time interval-worth of station-holding remedied this problem. If the mission planners don't care about targeting a particular time of arrival but simply want to reach a given spatial location as quickly as possible, the algorithm can easily be modified to accommodate this scenario. One simply stops the planning when any node in the target spatial position is surrounded by expanded neighbors, after which one can be certain that there is no earlier-arriving path.

III. EXPERIMENTAL METHOD

Our test scenario uses historical ROMS forecasts for the Southern California Bight in August and September 2009. ROMS is a 4D-VAR based oceanographic simulation that models ocean convection and current processes to predict currents, sea height, and salinity. It uses boundary conditions provided by real-time data streams such as CODAR (radar measurements of surface current speed), sea surface temperature, and moorings. The numerical model provides two main current-related data products. *Nowcasts* are estimates of current ocean conditions based on the latest available real-time data. The nowcasts are updated every 6 hours. Alternatively, current *forecasts* use a numerical simulation to predict future ocean conditions up to 48-hours in advance.

Figure 5 shows typical examples of 48-hour forecast and nowcast estimates for currents at two locations in the simulation area. Note that current velocities often reach magnitudes of $0.8m/s$; this exceeds the glider's own propulsion which would carry it at just $0.2 - 0.3m/s$ in still water. There is also a strong tidal influence apparent with 6-hour periodicity.

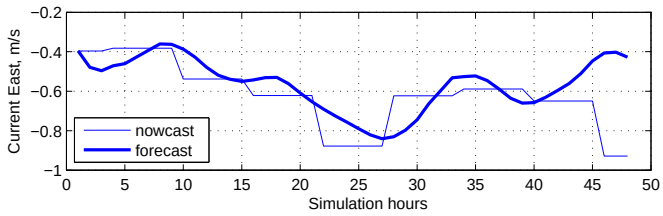

Fig. 5. Estimates of Eastward current velocity for a location near the Los Angeles coastline (33.5 N, 241.5 E). Both the hourly forecast and 6-hour nowcast data products are shown.

Finally, the forecast and nowcast products tend to diverge in the later hours of the simulation; this is typical of the simulation domain and reflects the accumulation of errors in the forecast predictions.

Glider deployment operators conventionally refine glider plans during most communications periods; this allows them to refine the entire mission plan at intervals of $6 - 8$ hours. However, in principle it is still advantageous to plan paths based on the full predictive window of ROMS current forecasts. Not only does this optimize glider travel by placing the vehicle in more favorable long-term positions but it could alleviate the onerous operational requirements of manually monitoring glider progress throughout the day. Thus, our experiment evaluates path planning for the full 48-hour time scale of ROMS predictions.

We model a vehicle similar to the Webb Research Slocum glider [14] that navigates by dead-reckoning between periodic GPS acquisitions. We translate the path planner's optimal trajectory into a timed sequence of segments, with each segment active during a specific time duration. The physical vehicle cannot follow this path perfectly because it receives only periodic position updates. Instead we use a simple pursuit method to track the trajectory; we define intermediate goal waypoints that are revised at the start of each dive cycle. At each surfacing the simulated glider updates its position using GPS and adjusts the glider heading based on the current active waypoint. Our simulation presumes surfacing activities every 60 minutes. Each planning trial requires that the glider travel from a predefined start location to some end location approximately $30km$ distant. Ground-truth currents are drawn from ROMS nowcasts. We selected 25 test paths to cover the area of interest with a variety of headings.

The tests compare three different path planning strategies:

- An *omniscient wavefront* approach computes a path based on the wavefront planning and perfect knowledge of the run-time currents. Naturally this is unrealistic for real ocean conditions.
- A *realistic wavefront* method uses wavefront planning based on the error-prone current *forecast* data products.
- A simple alternative does not use path planning at all. Instead it always travels directly towards the final goal.

Each segment of the trajectory is active during a specific time period; we compute a heading to reach the active segment's endpoint at each surfacing. We also consider several methods of deriving this direction (Figure 6). Current-sensitive control

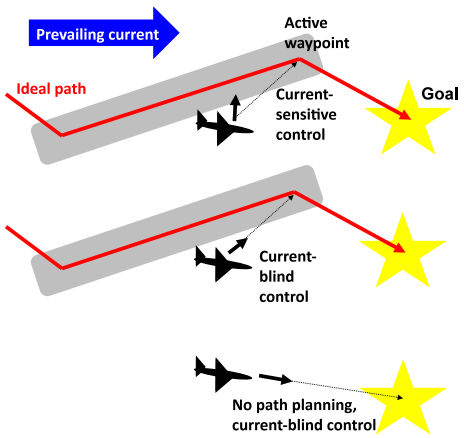

Fig. 6. We use several control strategies to track the planned trajectory. These include current-sensitive control that accounts explicitly for predictive forces, current-blind control that aims directly toward the current active waypoint, and a third option that forgoes path planning altogether.

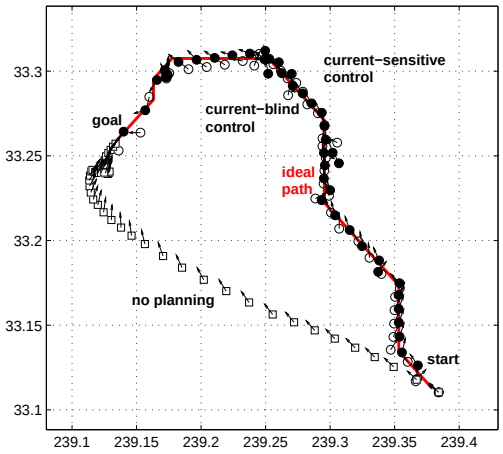

Fig. 7. Emperical glider positions and control propulsions for our three control strategies with hourly surfacings. Current-blind control (white circles) generally strays slightly down stream of the ideal path.

computes the control propulsion that combines with predicted currents to carry the vehicle toward the intermediate goal. Current-blind control applies a propulsion force in the direction of the next waypoint, so that the path itself uses current predictions but the local waypoint-following does not. This is more realistic for deployment scenarios where bandwidth limits preclude uploading current forecasts to the glider. Finally, a greedy strategy with no path planning does not have access to any current forecast information so it always propels the vehicle in the direction of the final goal.

Our experiments use wavefront planning to find interpolating waypoints for each path, in both directions, for two 48-hour periods spanning Sept 1 to 5, 2009. The goal for each trial is to position the glider at the opposite waypoint 48-hour planning window. We set gliders' still-water horizontal velocity to $0.3m/s$. Gliders plan within a travel zone consisting of the rectangular region enclosing the start and end waypoints, with an additional latitude/longitude margin of 0.2 degrees. Any valid glider path must lie entirely within this spatiotemporal volume. We simulate two different spatial resolutions for the current predictions. The first, a full-resolution simulation with $1km$ grid squares, is the highest resolution typically used for ROMS forecasts. We also simulate the common $6km$ grid square resolution by subsampling the current forecast data.

After the planning is complete we simulated the glider using the simple motion model described above. The plan execution environment uses the nowcast data, which differs somewhat from the realistic current forecasts. This simulates a realistic discrepancy between predicted and actual currents. Another source of execution uncertainty for both realistic and omniscient planners is the period of the dive cycle. We update the simulated glider's position every 60 minutes to simulate GPS acquisition and course corrections at periodic surfacing events.

IV. RESULTS

Each test has one of three outcomes. The first possibility, which can happen for either the realistic or omniscient planner, occurs whenever there is no valid path to reach the goal in the 48-hour window. This generally happens in the case of strong countervailing currents that overwhelm the glider's propulsion. In this case no plan is generated and we do not run the followup simulation. This case comprises approximately 50% of the planning trials. Alternatively when a trial does not use a planner, or the path planning system identifies a valid plan, in which cases the simulated glider either succeeds to fails to reach the goal.

We compute an error score for each trial based on the distance between the final glider position and the desired physical location. Figure 8 shows the distributions of these errors over all simulation trials. This comparison excludes runs where the path planner finds no valid path. Note that our simulation updates heading at 60 minute intervals, so it is unreasonable for a glider to hold position exactly at the target. Therefore distances less than $2km$ are tantamount to perfect successes, and distance errors within $5km$ are often sufficient for mission planning. Our tests did not show a significant performance difference between the different current-sensitive planning methods, or between control strategies utilizing the same current-sensitive plan. Anecdotal evidence from trials such as Figure 7 suggest that current-sensitive control could provide a benefit, but this is not significant for the majority of the simulated trials we considered.

On the other hand there is a clear difference between methods utilizing current predictions and greedy methods that do not involve a path planner. Simulations without a path planner can fail spectacularly; the worst quartile of greedy trials ends with the glider at least $13km$ off target, but when planning is involved only a handful of outlier trials misses the target. The benefit is largely due to the planner's ability to reject invalid paths. However, even in trials where valid paths

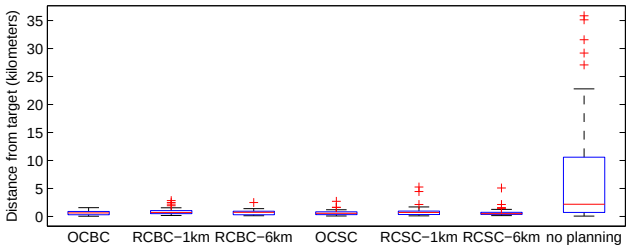

Fig. 8. Performance over all trials, measured in terms of final distance from the target location at the desired 48-hour time. Rectangles show median and quartiles, whiskers the extent of the data (excluding outliers, which are plotted individually). OCBC: Omniscient, current-blind control. OCSC: Omniscient, current-sensitive control. RCBS: Realistic, current-blind control. RCSC: Realistic, current-sensitive control. The 1km and 6km suffixes refer to the spatial size of the current prediction's grid squares for planning and execution. All planning scores exclude trials for which no valid path exists. The "no planning" case always propels the glider directly toward the goal.

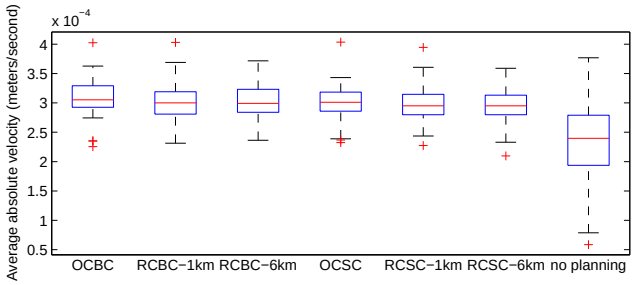

Fig. 9. Mean absolute velocities over all trials.

are found there is a significant median difference of errors, suggesting that planning-based approaches provide improved accuracy over planner-free methods (significant to 95%).

These results are corroborated by mean absolute velocities (Figure 9). Current-sensitive path plans make faster progress toward the goal, offering a higher rate of progress per time step by exploiting time-varying currents.

V. CONCLUSIONS

Glider mission planning is a challenging problem domain because the phenomena under observation and vehicle motion constraints can change dramatically on mission time scales. This work has sought to quantify the effect of forecast inaccuracy on the validity of mission plans. In our trials, glider performance using path planning from 48-hour forecasts was statistically indistinguishable from performance using perfect current predictions.

The main advantage of the path planner in these tests is to *recognize and exclude dangerous destinations*. The path planner is quite effective at rejecting infeasible paths where the glider would be pushed off course, which provides a valuable means of reducing risk in automated mission planning. At the same time, the reachability estimates computed by the path planner offer an intuitive way to visualize the envelope of planning possibilities and can assist gliders to identify safe mission goals that don't carry the glider through dangerous tides or eddies.

The path planner described in this work will shortly be deployed as part of a broader mission planning system within the Ocean Observatories Initiative Observing System Simulation Experiment (OSSE). Later it will join a mission planning system for physical gliders in the Mid-Atlantic Bight. We hope to improve the system with more sophisticated motion models and the ability to modify glider flight parameters mid-mission.

VI. ACKNOWLEDGMENTS

A Portion of the research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work also received support from the Ocean Observatories Initiative (OOI) and the National Science Foundation. We thank the Rutgers Glider Team and John Kerfoot for their support and counsel. Copyright 2010 California Institute of Technology. All Rights Reserved. US Government Support Acknowledged.

REFERENCES

[1] Y. Chao, L. Zhijin, J. D. Farrara, M. A. Moline, O. M. E. Schofield, and Sharanya J. Majumdar. Synergistic applications of autonomous underwater vehicles and the regional ocean modeling system in coastal ocean forecasting. *Limnol. Oceanogr*. 53:2251-2263, 2008.

[2] Y. Chao, Z. Li, J. Farrara, J. C. McWilliams, J. Bellingham, X. Capet, F. Chavez, J.-K. Choi, R. Davis, J. Doyle, D. Frantaoni, P. P. Li, P. Marchesiello, M. A. Moline, J. Paduan, S. Ramp. Development, implementation and evaluation of a data-assimilative ocean forecasting system off the central California coast. *Deep-Sea Research II*, doi:10.1016/j.dsr2.2008.08.011, 2009.

[3] J. K. Kerfoot, S. M. Glenn, O. M. Schofield, H. J. Roarty, and R. Chant, J. T. Kohut, "The View from the Cool Room: The Rutgers University Coastal Ocean Observatory." *AGU*, 2008.

[4] O. Schofield, J. Kohut, D. Aragon, L. Creed, J. Graver, C. Haldeman, J. Kerfoot, H. Roarty, C. Jones, D. Webb, and S. Glenn. "Slocum Gliders: Robust and Ready," *Journal of Field Robotics*, Vol. 24:6, pp.473–485.

[5] C. Pêtrès, Y. Pailhas, P. Patrón, Y. Petillot, J. Evans and D. Lane. "Path Planning for Autonomous Underwater Vehicles," *IEEE Trans. Robotics*, 2007.

[6] W. Zhang, T. Inanc, S. Ober-Blöbaum, and J. E. Mardesn, "Optimal Trajectory Generation for a Glider in Time-Varying 2D Ocean Flows B-spline Model," *Proceedings of the IEEE International Conference on Robotics and Automation*, 2008.

[7] M. Soulignac and P. Taillibert and M. Rueher. "Adapting the wavefront expansion in presence of strong currents," *IEEE International Conference on Robotics and Automation*, 2008, pp. 1352–1358.

[8] M. Soulignac, P. Talilibert, and M. Rueher. "Time-minimal Path Planning in Dynamic Current Fields," *Proceedings of the IEEE International Conference on Robotics and Automation*, 2009.

[9] D. Kruger, R. Stolkin, A. Blum, J. Briganti, "Optimal AUV Path Planning for Extended Missions in Complex, Fast-flowing Estuarine Environments," *ICRA*, 2007.

[10] N. E. Leonard, D. A. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. Davis. "Collective Motion, Sensor Networks, and Ocean Sampling." *Proceedings of the IEEE*, Vol. 95, No. 1, January 2007.

[11] P. Li, Y. Chao, Q. Vu, Z. Li, J. Farrara, H. Zhang and X. Wang. "OurOcean: An Integrated Solution to Ocean Monitoring and Forecasting," *OCEANS* 2006.

[12] R. N. Smith, Y. Chao, B. H. Jones, D. A. Caron, P. P. Li, and G. S. Sukhatme, "Trajectory design for autonomous underwater vehicles based on ocean model predictions for feature tracking," *Proceedings of Field and Service Robotics*, 2009.

[13] D. Thompson, S. Chien, M. Arrott, A. Balasuriya, Y. Chao, P. Li, M. Meisinger, S. Petillo, O. Schofield. Mission Planning in a Dynamic Ocean Sensorweb, *International Conference on Planning and Scheduling (ICAPS) SPARK Applications Workshop*, 2009.

[14] Webb Research Corporation, *gSlocum Shallow Battery Glider Operations Manual*. Ver 1.6, 11 Jan. 2005.