

Navigation of Unmanned Marine Vehicles in Accordance with the Rules of the Road

Michael R. Benjamin
NAVSEA Division Newport
Newport RI 02841
Center of Ocean Engineering, MIT
Cambridge MA 02139
Email: mikerb@csail.mit.edu

Joseph A. Curcio
John J. Leonard
Center of Ocean Engineering
Dept. of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge MA 02139
Email: jcurcio@mit.edu, jleonard@mit.edu

Paul M. Newman
Dept. of Engineering Science
Oxford University
Oxford, OX1 3PJ, UK
Email: pneuman@robots.ox.ac.uk

Abstract—This paper is concerned with the in-field autonomous operation of unmanned marine vehicles in accordance with convention for safe and proper collision avoidance as prescribed by the Coast Guard Collision Regulations (COLREGS). These rules are written to train and guide safe human operation of marine vehicles and are heavily dependent on human common sense in determining rule applicability as well as rule execution, especially when multiple rules apply simultaneously. To capture the flexibility exploited by humans, this work applies a novel method of multi-objective optimization, interval programming, in a behavior-based control framework for representing the navigation rules, as well as task behaviors, in a way that achieves simultaneous optimal satisfaction. We present experimental validation of this approach using multiple autonomous surface craft. This work represents the first in-field demonstration of multi-objective optimization applied to autonomous COLREGS-based marine vehicle navigation.

I. INTRODUCTION

A. Motivation

Mobile robotic platforms deployed in the marine environment offer substantial benefits to society while bringing a multitude of policy and legal challenges. Introducing mobile robotic vessels into navigable waterways presents the risk of collision with other vessels (both manned and unmanned), personal injury and property damage. Until policy, law and specifications evolve to address these issues, one can only speculate on the requirements imposed on developers, owners and operators of mobile robotic marine vehicles. However, an inspection of the relevant legal standards concerning safe operation of vessels in navigable waters reveals the likely need of owners, operators and programmers to abide by the current “rules of the road” given by the “International Regulations for Prevention of Collision at Sea”, or the “COLREGS” [1]. It is likely that as the use of mobile robotics continues to proliferate within the marine environment a new legal framework will evolve to address the ramifications of ownership and operation of these assets. A prudent operator might take the stance that, until the law catches up with the operation of these vehicles, the smart move is to make the vehicles compliant with the existing standards applicable to safe navigation [2], [3].

B. Solution Framework

Although the COLREGS is a document suitable for guiding human behavior, it is not suitable for direct input into a vehicle control system. In practice, there are often multiple rules simultaneously in effect, and to varying degrees. This is particularly true in congested waters. In many situations there are also multiple distinct vehicle maneuvers that would satisfy a given rule. Humans are fairly good at dealing with conflicting rules and capitalizing on flexibility within rules, but these situations present the harder challenges for autonomous vehicle control.

To address this problem, we have used a novel method of multi-objective optimization, interval programming (IvP), [4], within a behavior-based architecture for capturing COLREGS rules. Each rule corresponds to a behavior that produces an objective function over the vehicle’s decision, i.e., actuator, space. The objective functions capture the behavior prescribed by the COLREGS rule (in the peak areas of the function), but also capture its flexibility (in the non-peak areas). Each iteration of the vehicle control loop then involves the creation and solution of a multi-objective optimization problem, where each module contributes one function. This approach is suitable for building additional mission modules, on top of a COLREGS foundation where the mission modules also produce additional functions alongside the COLREGS modules.

Results from simulation and results from in-field experiments with multiple autonomous surface craft are reported to validate these algorithms and architecture.

II. BACKGROUND

A. Behavior-Based Control

In behavior-based systems, robot or vehicle control is the result of set of independent, specialized modules working together to choose appropriate vehicle actions. It can be viewed as an alternative to the traditional sense-plan-act control loop as shown in Fig. 1, where decision-making and planning are performed on a single world model that is built up and maintained over time.

Commonly cited virtues of behavior-based systems include: the ease of development of the independent modules, the lack




Fig. 1. Behavior-based control differs from conventional control by composing overall vehicle behavior into distinct modules that are developed and operate largely in isolation, and coordinated through an action selection mechanism. In this case, action selection is in the form of a new multi-objective optimization technique to overcome known difficulties associated with behavior-based control.

of a single complex world model, and the potential for a highly reactive vehicle with certain behaviors triggered by the appropriate events in a dynamic environment.

The origin of such systems is commonly attributed to Brooks’ “subsumption architecture” in [5]. Since then, it has been used in a large variety of applications including: indoor robots, e.g., [6], [7], [8], [9], [10], [11], [12], [13], [14], land vehicles, e.g., [15], planetary rovers, e.g., [16], [17], [18], and marine vehicles, e.g., [19], [20], [21], [22], [23], [24], [25].

Action selection, as indicated in Fig. 1, is the process of choosing a single action for execution, given the outputs of the behaviors. The “action space” is the set of all possible distinct actions. For example, all combinations of rotational and angular velocity for a robot, or all speed, heading and depth combinations for a marine vehicle.

B. Known Difficulties in Behavior-Based Control

The primary difficulty often associated with behavior-based control concerns action selection - namely how to ensure the chosen action really is in the best overall interest of the robot or vehicle. An action generally is a vector of values, one for each actuator being controlled. For example, the rotational and angular velocity for a land robot, or heading, speed and depth for a marine robot.

Generally there are two techniques used in practice. The simplest method is to pick (at every iteration of the control loop) a single behavior to have exclusive control of the vehicle. Some approaches, like [21], [5], [26] assign a set of fixed priorities to behaviors, and conditions for their activation. The priorities do not change dynamically. In other implementations, like [23], priorities may be determined dynamically. Although appealing due to its simplicity, it is problematic in applications where the outright ignoring of the “secondary” behaviors leads to gross vehicle inefficiency, as is the situation with task described in this work.

The other common form of action selection, known variably as “action averaging”, “vector summation” etc., takes the output of each behavior in the form of a vector and uses the average numerical value as the action sent to the vehicle’s actuators. Summation is typically weighted to reflect behavior priority. This method has been used effectively in a number of applications, [6], [27], [28], [22], [29].

When the preferred actions of two distinct behaviors disagree, this approach rests on the idea that the alternative actions degrade in effectiveness in a manner depicted in Fig.2. In




Fig. 2. In action-averaging, each behavior outputs a *single* best action. The best action presumably is the most effective among alternative actions for that particular behavior. The effectiveness levels of alternative actions are rendered here only for illustration and do not participate in the action averaging process. When two behaviors are non-mutually exclusive and share common action choices with high levels of effectiveness, as shown here, then action averaging typically reflects an appropriate compromise between behaviors.

such a case, the action, or actuator setting, in between the two individually preferred actions may indeed be the most effective action overall. However, action averaging is problematic in cases when alternative actions degrade in effectiveness in a manner depicted in Fig. 3, where the numerical average does not represent an effective compromise between two behaviors that are, in effect, mutually exclusive.




Fig. 3. The average of the best action produced by two behaviors may have poor value for *both* behaviors. The chooser of the action is oblivious to the error since the behaviors output a single preferred action and do not communicate the underlying effectiveness of their alternatives, rendered here only for illustration. In this case, interests being pursued by the two behaviors are mutually exclusive, and the “compromise” is detrimental to both.

III. THE “IVP” ARCHITECTURE

A. Behavior-Based Control with Interval Programming

By using multi-objective optimization in action selection, behaviors produce an *objective function* rather than a single preferred action ([10], [15], [30]). In the examples in Figs. 2 and 3, the objective functions are what distinguish opportunities for compromise. Note the pair of preferred actions in Figs. 2 and 3 are virtually the same despite the differences in utility of secondary alternatives.

There are two practical challenges in handling objective functions: (1) the method must be fast enough to accommodate the vehicle control loop, typically 1-20Hz. On each iteration new functions are created and a new problem solved (speed is a primary advantage of action averaging). (2) the method

cannot be overly restrictive on objective function form. Non-convex, multi-modal functions are quite common.

The interval programming model specifies (1) a scheme for representing functions of unlimited form and (2) a set of algorithms for finding the globally optimal solution. All functions are piecewise linearly defined, thus they are typically an *approximation* of a behavior’s true underlying utility function. Search is over the weighted sum of individual functions and uses branch and bound to search through the combination space of pieces rather than the decision space of actions. The only error introduced is in the discrepancy between a behavior’s true underlying utility function and the piecewise approximation produced to the solver. This error is preferable compared with the error of restricting all behaviors to a quadratic function for example. Furthermore, the search is much faster than brute force evaluation of the decision space, as done in [15]. The decision regarding function accuracy is a local decision to the behavior designer, who typically has insight into what is sufficient. The solver guarantees a globally optimal solution and this work validates that such search is feasible in a vehicle control loop of 4Hz on a 600MHz computer.

To enhance search speed, the initial decision provided to the branch and bound algorithm is the output of the previous cycle, since typically what was a good thing to do a fraction of a second prior, is not a bad thing at the present. (In fact, when something *does* change dramatically in the world, such as hitting a waypoint, the solve time has been observed to be roughly 50% longer, but still comfortably under practical constraints). See [4] for more on IvP and search algorithms.

B. The Decision Space and Vehicle Helm

Action selection here consists of deciding the variables, heading (θ), speed (v), and time-on-leg (t). The latter is the “intended” duration of the chosen action. The *helm* is the module consisting of the behaviors and the optimization (action selection) engine. The helm produces a tuple $\langle \theta, v, t \rangle$ on every iteration of the control loop, and the values of heading and speed are fed into PID control to produce rudder and thruster commands.

The helm, through GPS, has access to its own position (x, y) , and through wireless communication has access to the position, heading, and speed of a given vehicle $(x_b, y_b, \theta_b, v_b)$. Each helm behavior has access to these variables, and they comprise all the necessary input to the behaviors described below for this work.

C. Closest Point of Approach

For COLREGS behaviors, an important quality of a candidate action $\langle \theta, v, t \rangle$, is the *closest point of approach (CPA)* between two vehicles during a candidate leg. Behaviors producing objective functions over candidate actions (legs), need to calculate this value quite often, and need to be efficient. Thus, the algorithm and efficiency measures are given here.

For a given $\langle \theta, v, t \rangle$, we compute the time t' when the minimum distance between two vehicles occurs given by:

$$dist^2(\theta, v, t) = k_2 t^2 + k_1 t + k_0, \quad (1)$$

where

$$k_2 = v^2 - 2 \cos(\theta) v \cos(\theta_b) v_b + v_b^2 - 2 \sin(\theta) v \sin(\theta_b) v_b$$

$$k_1 = 2 \cos(\theta) v y - 2 \cos(\theta) v y_b - 2 y \cos(\theta_b) v_b + 2 \cos(\theta_b) v_b y_b + 2 \sin(\theta) v x - 2 \sin(\theta) v x_b - 2 x \sin(\theta_b) v_b + 2 \sin(\theta_b) v_b x_b$$

$$k_0 = y^2 - 2 y y_b + y_b^2 + x^2 - 2 x x_b + x_b^2$$

The stationary point is obtained by taking the first derivative with respect to t :

$$dist^2(\theta, v, t)' = 2k_2 t + k_1.$$

Since there is no “maximum” distance, this stationary point always represents the closest point of approach, and therefore:

$$t' = \frac{-k_1}{2k_2}.$$

The value of t' is clipped by $[0, t]$, and t' is zero when the two vehicles have the same heading and speed (the only condition where k_2 is zero). The CPA value is obtained by plugging t' back into (1). For fast behavior implementation, when a behavior begins the job of creating an objective function, prior to many CPA calculations, all terms in (1) comprised exclusively of $x, y, x_b, y_b, \theta_b, v_b$ are calculated once and stored for later calculations.

IV. THE “BREAD AND BUTTER” BEHAVIORS

A primary motivation for applying multi-objective optimization to the COLREGS navigation problem is that such behaviors could augment existing behaviors without any mutual design consideration. We present two such “bread and butter” behaviors sufficient for illustrating the subsequent description of the COLREGS behaviors.

A. A Waypoint Behavior

The waypoint behavior is populated with a set of (x_i, y_i) waypoints, and has access to the vehicle’s current position (x, y) via GPS. It ranks candidate legs $\langle \theta, v, t \rangle$ based on the proximity of the resulting position to the next waypoint.




Fig. 4. The objective function produced for the waypoint behavior rates decisions higher that bring it closer to the next waypoint. The utility drops linearly. Darker shades represent higher utility. Typically about 600 linear pieces are used to represent this function.

B. A Collision Avoidance Behavior

The collision avoidance behavior differs from the COLREGS behaviors only in that it doesn't care *how* collisions are avoided. It is based primarily on the CPA distance for a candidate $\langle \theta, v, t \rangle$ decision. The CPA distance is applied to a metric function that assigns a utility based on parameterizable a "inner" distance and an "outer" distance. CPA distances lower than the inner-distance are treated as collisions, and values greater than the outer-distance have a plateau utility nominally set to 100. (Functions are normalized prior to the application of the priority weight, so utility ranges are insignificant). CPA distance in between the outer-distance and inner-distance degrade linearly, illustrated by the example in Fig. 5.




Fig. 5. The objective functions produced by the AvoidCollision behavior for two situations. In both cases, the controlled vehicle has a top speed of 4 meters/second with the contact moving on the indicated heading. Darker colors represent more favorable actions, and larger radii on the plot indicate higher candidate speeds. The vehicles are 200 meters apart. CPA distances less than 10 meters are considered collisions (in white) and those greater than 75 meters are neutral (in black). Distances in between degrade linearly. In (a) the contact is moving at 3 m/s and in (b) the contact is moving at 5 m/s.

The priority of the behavior is determined by the CPA distance of a hypothetical continuation of the current heading and speed out another n seconds. A simulation track is shown in Fig. 6.




Fig. 6. In simulation, the lefthand vehicle is guided by a waypoint and collision avoidance behavior to the point on the right. (Note this vehicle passes to the opposite side as would be prescribed by the COLREGS. Compare this trajectory with Fig. 12.) The righthand vehicle is executing a waypoint behavior with no collision avoidance to the waypoint on the left. The function rendered represents the addition of the two objective functions at that point in time.

V. THE COLREGS BEHAVIORS

There are nearly 40 rules that comprise the "COLREGS", nearly half of which concern lighting and sounds. We focus our attention on the four most challenging rules, from an autonomous navigation perspective, that cover "head-on" situations and "crossing" situation, rules 14-16. It is also

worth noting rules 8(b), (d) which address collision avoidance generally (all excerpts are from [1]):

Rule 8 : "Action to Avoid Collision"

(b) Any alteration of course and/or speed to avoid collision shall, if the circumstances of the case admit, be large enough to be readily apparent to another vessel observing visually or by radar; a succession of small alterations of course and/or speed should be avoided.

(d) Action taken to avoid collision with another vessel shall be such as to result in passing at a safe distance. The effectiveness of the action shall be carefully checked until the other vessel is finally past and clear.

This rule reveals a measure of the flexibility common in the rules, suitable for humans, but tricky for robots, such as "large enough to be readily apparent", and "small alterations of course". Generally the flexibility is found in both the *condition* of the rule and the *application* of the rule. Exploiting the latter is of paramount importance, since the rules need to at times co-exist with other rules as well as the efforts of the vehicle to complete its task.

A. The "Head-on" Behavior

The rule regarding two vessels approaching head-on is Rule 14 in [1]:

Rule 14 : "Head-on Situation"

(a) Unless otherwise agreed, when two power-driven vessels are meeting on reciprocal or nearly reciprocal courses so as to involve risk of collision each shall alter her course to starboard so that each shall pass on the port side of the other.

(b) Such a situation shall be deemed to exist when a vessel sees the other ahead or nearly ahead and by night she could see the mast-head lights of the other in a line or nearly in a line or both sidelights and by day she observes the corresponding aspect of the other vessel.

(c) When a vessel is in any doubt as to whether such a situation exists she shall assume that it does exist and act accordingly.

The objective function produced by this behavior is also based on the closest point of approach for a given candidate maneuver leg $\langle \theta, v, t \rangle$. The "head-on" condition referred to in the rule is interpreted to be in effect when the relative bearing between the two vehicles is within 15 degrees of the heading of the contact. To achieve the desired effect, the candidate heading is compared against the current relative bearing and starboard maneuvers are rated higher, and likewise lower for port maneuvers, as shown in Fig. 7.

In addition, the behavior given a range outside of which the priority of the behavior is zero and is inactive (see Fig. 12(a)).

B. The "Crossing" Behaviors

COLREGS Rule 15 and 16 serve to define a "crossing" situation.

Rule 15: "Crossing Situation"




Fig. 7. The “head-on” behavior produces objective functions based in part on the closest point of approach for a candidate maneuver and in part on a preference for starboard maneuvers passing the contact on the port side. Darker colors represent more favorable actions, and larger radii on the plot indicate higher candidate speeds. Compare against Fig. 5(b) where maneuvers to either side of the contact are nearly equal in preference.




Fig. 9. The “crossing” behavior produces objective functions based in part on the closest point of approach for a candidate maneuver and in part on a preference for maneuvers that do not cross ahead of the other vessel. Darker colors represent more favorable actions, and larger radii on the plot indicate higher candidate speeds. (Compare with Fig. 5(b)).

(a) When two power-driven vessels are crossing so as to involve risk of collision, the vessel which has the other on her starboard side shall keep out of the way and shall, if the circumstances of the case admit, avoid crossing ahead of the other vessel.




Fig. 8. The Give-way vessel yields to the Stand-on vessel.

Rule 16: “Action by Give-way Vessel”

Every vessel which is directed to keep out of the way of another vessel shall, so far as possible, take early and substantial action to keep well clear.

The objective function produced by this behavior also utilizes closest point of approach for a given candidate maneuver leg $\langle \theta, v, t \rangle$ in its objective function formulation.

The “crossing” condition referred to in the rule is interpreted to be in effect when the relative bearing between the two vehicles is greater than 15 degrees of the heading of the contact, but less than 90 degrees. According to Rule 15, crossing ahead of the other vessel is to be avoided. To represent this preference in the objective function, a candidate leg, $\langle \theta, v, t \rangle$, is further evaluated to determine if it crosses ahead or behind the other vessel. The ranking of utility of an action is penalized further if it crosses ahead, as shown in Fig. 9.

VI. EXPERIMENTS

Testing is done both in simulation and on two kayak-based autonomous surface crafts depicted in Figs. 10 and 11. Each vehicle had access to a compass and Garmin 18 GPS, the latter with updates of 1Hz. The GPS also provided the vehicle speed information, and at sufficiently high enough speed ($> 0.5\text{m/s}$), the GPS was preferred over the compass for heading measurements. Each vehicle communicated its position, heading and speed to the other vehicle at a rate of



Fig. 10. Two kayak-based autonomous surface craft for used for in-field experiments. Each had access to GPS and shared their current position and trajectory with the other.

4Hz, via a 802.11b wireless link.

Fig. 12 shows a representative experimental result that we have achieved using the algorithm described above. This experiment was designed to test “Rule 14 (Head-on Collision)”. The caption in the figure provides a detailed step-by-step account of how the correct behavior emerges based on the IvP optimized action selection strategy described above in Section III.

Each vehicle was running a suite of processes each communicating through a central database process. This collection of software is known as MOOS and is further described in [26]. It has been used effectively in a number of autonomous marine vehicle projects, e.g., [31]. Each process subscribes to the database process for information it needs and in turn publishes information to the database that it would like to make available to other processes. MOOS is an open source project that includes many utility processes besides the core communication and scheduling functionality. The IvP Helm is a single process in a community of MOOS processes that subscribes to vehicle position information and publishes actuator commands (which are in turn subscribed to by the processes interfacing with the actuators). Many essential processes used to conduct these experiments are part of the MOOS distribution.

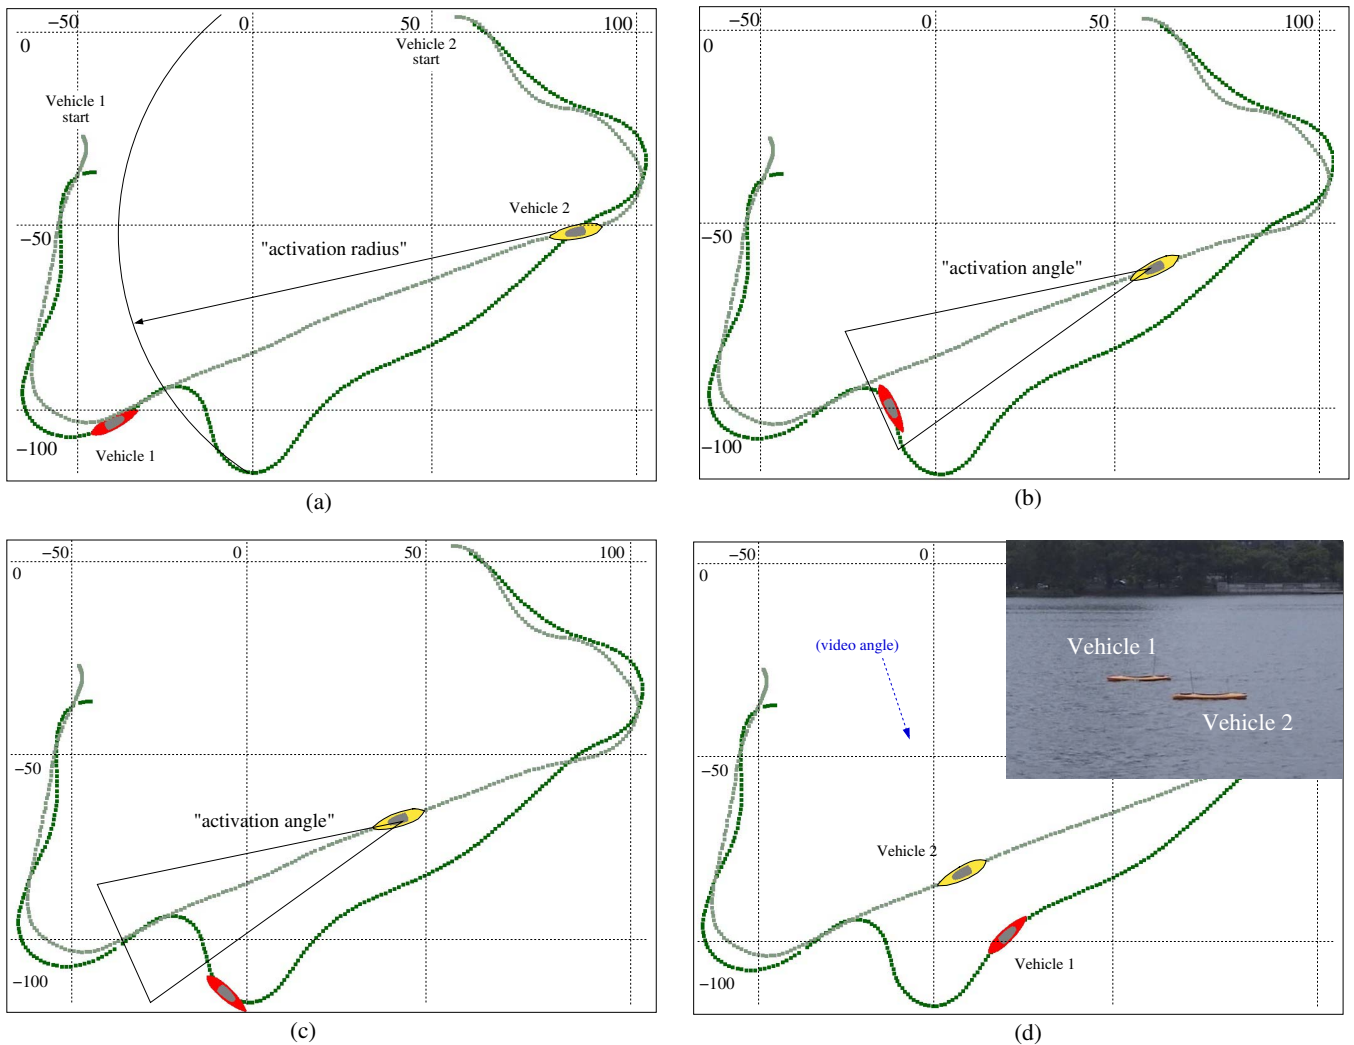


Fig. 12. In-field experiments with two autonomous kayaks verifying the COLREGS “Head-on” Rule 14 behavior. Vehicle 1 and 2 are put on a head-on collision course through a series of waypoints. Vehicle 1 is utilizing a waypoint behavior and a “rule-14” behavior. Vehicle 2 is only using a waypoint behavior and does not make any attempt at collision avoidance with vehicle 1. In (a) the two vehicles are on a head-on collision course with vehicle 1 heading to waypoint (105, -35), and vehicle 2 heading to waypoint (-50, -110). In (a) only the waypoint behavior is active in vehicle 1 because vehicle 2 is still outside the activation range. In (b) vehicle 1 is within the activation range and within the activation angle specified to the rule-14 behavior and is thus making a starboard maneuver to avoid collision. In (c) vehicle 1 has just moved outside the activation angle and thus the rule-14 behavior becomes inactive, and the influence of the waypoint behavior begins to dominate again. In (d) vehicle 1 is proceeding uninhibited toward its destination. The image is from video shot during the experiment that produced the data shown here.

VII. CONCLUSION

This paper has investigated the problem of autonomous collision avoidance and navigation for autonomous surface craft. We have presented a novel method using IvP-based multi-objective optimization to coordinate distinct vehicle behaviors representing both task execution and established human protocol for safe navigation. This paper also provides, to our knowledge, the first ever demonstration of such a system on a physical marine platform. Further research is necessary to explore the robustness of this method in more complex navigation scenarios, both in simulation and on the water.

ACKNOWLEDGEMENTS

This work was funded in part by the NUWC Division Newport ILIR program managed by Dick Philips, as well as by the National Oceanic and Atmospheric Administration (NOAA) in a program managed by Justin Manley. Development of the interval programming model is funded by Dr. Don Wagner and Adam Nucci at ONR, and formerly the NUWC ILIR program. This work is also funded in part by the ONR-Global NICOP program, Dr. John Fay.

REFERENCES

- [1] U. C. G. Commandant, *International Regulations for Prevention of Collisions at Sea, 1972 (72 COLREGS)*, US Department of Transportation, US Coast Guard, August 1999, COMMANDANT INSTRUCTION M16672.2D.




Fig. 11. A kayak-based autonomous surface craft.

- [2] E. D. Brown and N. J. Gaskell, "Report on the Law Relating to Autonomous Underwater Vehicles," Society for Underwater Technology, Tech. Rep., 2000.
- [3] S. Showalter, "The Legal Status of Autonomous Underwater Vehicles," *The Marine Technology Society Journal*, vol. 38, no. 1, pp. 80–83, 2004.
- [4] M. R. Benjamin, "The Interval Programming Model for Multi-Objective Decision Making," Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, Tech. Rep. AIM-2004-021, September 2004.
- [5] R. A. Brooks, "A Robust Layered Control System for a Mobile Robot," *IEEE Journal of Robotics and Automation*, vol. RA-2, no. 1, pp. 14–23, April 1986.
- [6] R. C. Arkin, "Motor Schema Based Navigation for a Mobile Robot: An Approach to Programming by Behavior," in *Proceedings of the IEEE Conference on Robotics and Automation*, Raleigh, NC, 1987, pp. 264–271.
- [7] R. C. Arkin, W. M. Carter, and D. C. Mackenzie, "Active Avoidance: Escape and Dodging Behaviors for Reactive Control," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 5, no. 1, pp. 175–192, 1993.
- [8] J. Hoff and G. Bekey, "An Architecture for Behavior Coordination Learning," in *Proceedings of the 1995 IEEE International Conference on Neural Networks*, Perth, Australia, November 1995, pp. 2375–2380.
- [9] S. Lenser, J. Bruce, and M. Veloso, "A Modular Hierarchical Behavior-Based Architecture," in *RoboCup-2001: The Fifth RoboCup Competitions and Conferences*, A. Birk, S. Coradeschi, and S. Takokoro, Eds. Springer Verlag, 2002.
- [10] P. Pirjanian, "Multiple Objective Action Selection and Behavior Fusion," Ph.D. dissertation, Aalborg University, 1998.
- [11] J. Riecki, "Reactive Task Execution of a Mobile Robot," Ph.D. dissertation, Oulu University, 1999.
- [12] A. Saffiotti, E. H. Ruspini, and K. Konolige, "Using Fuzzy Logic for Mobile Robot Control," in *Practical Applications of Fuzzy Technologies*, H. J. Zimmerman, Ed. Kluwer Academic Publishers, 1999, ch. 5, pp. 185–206.
- [13] E. Tunstel, "Coordination of Distributed Fuzzy Behaviors in Mobile Robot Control," in *IEEE International Conference on Systems, Man, and Cybernetics*, Vancouver, BC, Canada, October 1995, pp. 4009–4014.
- [14] M. Veloso, E. Winner, S. Lenser, J. Bruce, and T. Balch, "Vision-servoed Localization and Behavior-Based Planning for an Autonomous Quadruped Legged Robot," in *Proceedings of AIPS-2000*, Breckenridge, April 2000.
- [15] J. K. Rosenblatt, "DAMN: A Distributed Architecture for Mobile Navigation," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, 1997.
- [16] H.-H. Ju, P.-Y. Cui, and H.-T. Cui, "Autonomous Behavior Path Planning for Lunar Rover," *ACTA AUTOMATICA SINICA*, vol. 29, no. 2, pp. 324–329, 2002.
- [17] P. Pirjanian, T. L. Huntsberger, and P. S. Schenker, "Development of CAMPOUT and its further applications to planetary rover operations: a multi-robot control architecture," in *Proceedings of SPIE Conference on Sensor Fusion and Decentralized Control in Robotic Systems IV*, Newton, MA, October 2001.
- [18] S. Singh, R. Simmons, T. Smith, A. Stentz, V. Verma, A. Yahja, and K. Schwehr, "Recent Progress in Local and Global Traversability for Planetary Rovers," in *IEEE Conference on Robotics and Automation*, San Francisco, CA, April 2000.
- [19] S.-M. Lee, K.-Y. Kwon, and J. Joh, "A Fuzzy Logic for Autonomous Navigation of Marine Vehicle Satisfying COLREG Guidelines," *International Journal of Control, Automation, and Systems*, vol. 2, no. 2, June 2004.
- [20] M. R. Benjamin, "Multi-objective Autonomous Vehicle Navigation in the Presence of Cooperative and Adversarial Moving Contacts," in *OCEANS 2002*, Biloxi Mississippi, October 2002.
- [21] A. A. Bennet and J. J. Leonard, "A Behavior-Based Approach to Adaptive Feature Detection and Following with Autonomous Underwater Vehicles," *IEEE Journal of Oceanic Engineering*, vol. 25, no. 2, pp. 213–226, April 2000.
- [22] M. Carreras, J. Battle, and P. Ridao, "Reactive Control of an AUV Using Motor Schemas," in *International Conference on Quality Control, Automation and Robotics*, Cluj Napoca, Rumania, May 2000.
- [23] R. Kumar and J. A. Stover, "A Behavior-Based Intelligent Control Architecture with Application to Coordination of Multiple Underwater Vehicles," *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Cybernetics*, vol. 30, no. 6, pp. 767–784, November 2001.
- [24] J. K. Rosenblatt, S. B. Williams, and H. Durrant-Whyte, "Behavior-Based Control for Autonomous Underwater Exploration," *International Journal of Information Sciences*, vol. 145, no. 1-2, pp. 69–87, 2002.
- [25] S. B. Williams, P. Newman, G. Dissanayake, J. K. Rosenblatt, and H. Durrant-Whyte, "A decoupled, distributed AUV control architecture," in *Proceedings of 31st International Symposium on Robotics*, Montreal, 2000.
- [26] P. M. Newman, "MOOS - A Mission Oriented Operating Suite," MIT Department of Ocean Engineering, Tech. Rep. OE2003-07, 2003.
- [27] R. C. Arkin and T. Balch, "AuRA: Principles and Practice In Review," *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 9, pp. 175–189, 1997.
- [28] T. Balch and R. C. Arkin, "Behavior-Based Formation Control for Multi-agent Robot Teams," *IEEE Transactions on Robotics and Automation*, December 1998.
- [29] O. Khatib, "Real-Time Obstacle Avoidance for Manipulators and Mobile Robots," in *Proceedings of the IEEE International Conference on Robotics and Automation*, St. Louis, MO, 1985, pp. 500–505.
- [30] M. R. Benjamin, "Interval Programming: A Multi-Objective Optimization Model for Autonomous Vehicle Control," Ph.D. dissertation, Brown University, Providence, RI, May 2002.
- [31] M. Benjamin, M. Grund, and P. Newman, "Multi-objective Optimization of Sensor Quality with Efficient Marine Vehicle Task Execution," in *International Conference on Robotics and Automation (ICRA)*, Orlando, Florida, May 2006.