A Multiple-Hypothesis Approach to Concurrent Mapping and Localization for Autonomous Underwater Vehicles*

Christopher M. Smith and John J. Leonard
(cmsmith@mit.edu and jleonard@mit.edu)
Department of Ocean Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139, U.S.A.

Abstract

This paper describes a multiple hypothesis approach to concurrent mapping and localization (CML) for autonomous underwater vehicles (AUVs). The objective of CML is to enable a mobile robot to build a map of an unknown environment, while simultaneously using that map to navigate with bounded position error. Multiple hypothesis concurrent mapping and localization (MHCML) has potential to provide a theoretically consistent framework that incorporates navigation error, sensor noise, data association uncertainty, and physically-based sensor models. MHCML is fundamentally different from conventional multiple hypothesis tracking because multiple hypotheses are considered for both the location of the vehicle and the locations of features. New techniques for evaluation of decision dependencies and calculation of likelihoods for vehicle and feature tracks are introduced. Simulation results are presented to illustrate the viability of the approach for an AUV equipped with a forward-look sonar.

1 Introduction

This paper describes a multiple hypothesis approach to concurrent mapping and localization (CML) for autonomous underwater vehicles (AUVs). The objective of CML is to enable a mobile robot to build a map of an unknown environment, while simultaneously using that map to navigate with bounded position error. This problem has been a popular topic in the research community, due to its theoretical challenges and critical importance for many mobile robot applications.

A seminal technique for concurrent mapping and localization, called the stochastic map, was published by Smith, Self, and Cheeseman [1990]. The stochastic map consists of a single state vector that represents the estimates of the vehicle and feature locations and an associated covariance matrix. As the vehicle moves around its environment, taking measurements of environmental features, the stochastic map is updated using an extended Kalman filter. Moutarlier and Chatila extended upon this approach and provided the first experimental implementation using laser range data [Moutarlier and Chatila, 1989]. Rencken implemented CML using Polaroid sonar data on a mobile robot [Rencken, 1993]. More recently, Chong and Kleeman have investigated CML using a novel imaging sonar that can classify and accurately localize features such as planes and corners in a typical office environment [Chong and Kleeman, 1997].

The major problems encountered by the stochastic map are the failure of the extended Kalman filter to properly track the highly nonlinear transformations involved in geometric estimation and the fact that the technique scales (at best) quadratically with the number of features present [Leonard and Durrant-Whyte, 1991]. Also, Smith, Self, and Cheeseman did not address errors which may arise from ambiguity in the source of sensor data. Uhlmann recently provided a theoretical investigation of some of these problems, suggesting a new alternative to the extended Kalman filter that may have a broad impact on the field of estimation [Uhlmann, 1995]. However, many important issues remain for future research.

Our research contributes to existing work in this area by presenting the first multiple hypothesis approach to the problem and by considering the first underwater implementation of feature-based CML. (Stewart [1988] has considered concurrent mapping and positioning of underwater vehicles using a grid-based representation similar to the certainty grid of Moravec [1989] and Elfes [1987].) Implementations of CML performed to-date have tended to employ ad-hoc methods to determine the correspondence between sensor measurements and map features, such as such as manual matching or nearest-neighbor association [Bar-Shalom and Fortmann, 1988]. In contrast, multiple hypothesis concurrent mapping and local-

*To appear in International Conference on Field and Service Robotics, Canberra, Australia, December 1997.

ization (MHCML) offers the potential to provide a theoretically consistent framework that incorporates navigation error, sensor noise, data association uncertainty, and physically-based sensor models.

Our approach constitutes a significant extension to multiple hypothesis tracking (MHT), a technique for multitarget tracking pioneered by Reid [1979] and subsequently extended by Mori et al. [1986]. Cox and Leonard [1994] provided the first implementation of MHT for a mobile robot application, demonstrating accurate map-building from Polaroid sonar data, under the assumption that precise vehicle position information was available. The objective of multiple hypothesis tracking is to group measurements that originate from the same geometric feature, while rejecting spurious measurements. Each new observation initializes a tree of possible interpretation hypotheses that classify the measurements according to different target models. The tree is grown as new observations are validated with hypothesized targets, and is subsequently pruned to choose the single best interpretation of all past measurements [Cox and Leonard, 1994]. An important feature of the approach is that ambiguous assignment decisions can be deferred until more data is acquired and a better decision can be made.

As pointed out in [Cox and Leonard, 1994], conventional MHT techniques are inadequate once vehicle position uncertainty is included in the problem formulation. A brute force MHT approach to implement Smith, Self, and Cheesman's stochastic map would require that *exponentially many* global covariance matrices be constructed and maintained, presenting an inordinate computational burden. MHCML is fundamentally different from conventional MHT because multiple hypotheses are considered for both the location of the vehicle and the locations of features. Because cluster partitioning [Kurien, 1990] cannot be performed, new techniques are introduced for evaluation of decision dependencies and calculation of likelihoods for vehicle and feature tracks.

Our work is directed to support the navigation and mapping requirements of the US Navy's LDUUV and 21UUV underwater vehicles, currently in development at the Naval Undersea Warfare Center in Newport, RI, U.S.A. [Levine *et al.*, 1995]. The problem of CML is especially challenging for AUVs. The ocean environment is dynamic and inherently three-dimensional and underwater sensor data interpretation can be very complex. Our approach assumes that the AUV is equipped with a forward-look sonar system, such as the Navy's high resolution array (HRA) imaging sonar [Nussbaum *et al.*, 1996]. We assume that sonar image processing techniques can be employed to extract features from the sonar data, in the form of compact features such as mines, lobster-traps, rock outcroppings, or distinctive bathymetric features [Medeiros and Carpenter, 1996]. The CML objective for an AUV is to process this information to build a consistent map of the locations of environmental features, while simultaneously using that map for accurate position determination during extended missions of many hours or days. Figure 1 illustrates this process via an implementation of the stochastic map in a simulation run *without correspondence ambiguity* that combines a dynamic model for an AUV and the geometry of a high resolution forward-look sonar array [Medeiros and Carpenter, 1996]. The objective of the MHCML algorithm is to achieve similar performance for situations in which data association ambiguity is present.

Good navigation information is essential for safe operation and recovery of an autonomous underwater vehicle (AUV), especially for under-ice deployments [Bellingham *et al.*, 1993] or in regions of high currents [Schmidt *et al.*, 1996]. For the data gathered by an AUV to be of scientific, commercial, or military value, the location from which the data has been acquired must be accurately known. A popular method for underwater vehicle navigation is long baseline (LBL) navigation, in which the vehicle operates within a pre-calibrated array of acoustic beacons [Vaganay *et al.*, 1996]. For many applications of AUVs, however, such as coastal operations over large areas, the use of acoustic beacons is undesirable or impractical. If an accurate *a priori* map of the environment is available, one approach to globally-referenced position estimation is to use measurements of geophysical parameters, such as bathymetry, magnetic field, or gravitational anomaly [Tuohy *et al.*, 1996; Geyer *et al.*, 1987]. These approaches are based on matching sensor data with an *a priori* environment map, under the assumption that there is sufficient spatial variation in the parameter(s) being measured to permit accurate localization. In practice, because an up-to-date, high-quality map will usually be unavailable for the operating area of interest, the capability for an AUV to build a map of an unknown environment and use that map to navigate is highly attractive.

2 Multiple Hypothesis Concurrent Mapping and Localization Algorithm

The concurrent mapping and localization algorithm tracks the vehicle and a set of proposed features through time [Smith *et al.*, 1997a]. Each of these objects has a set of one or more estimates at any given time. A state is an estimate of the position and error covariance of the vehicle or feature (assuming some model). A track is a tree of possible states linked by causality. There are four major processes at each time cycle in CML. First, current estimates are projected using dynamic models to generate predicted possible states. Second, as measurements are processed, hypotheses are formed about

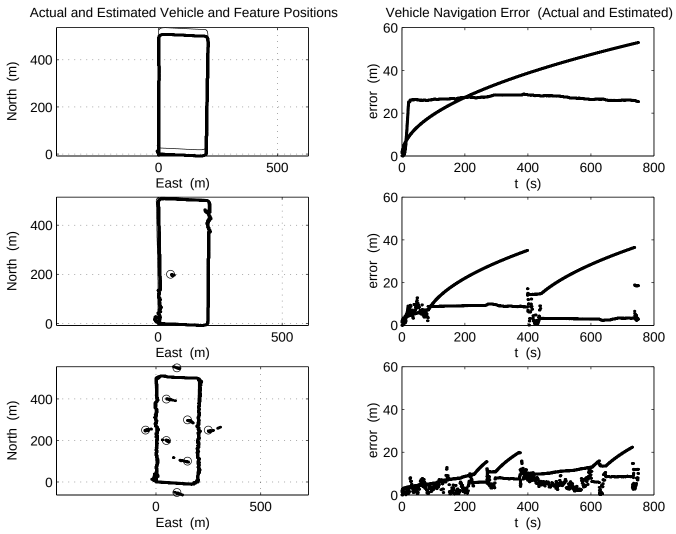

Figure 1: Simulation sample run of a stochastic map approach *without data association error* showing the benefit of using environmental cues to reduce navigational error. In each of the three runs, the vehicle travels in a rectangle; the number of observable features changes in each run. The left graphs show the estimated (thick) and actual (thin) vehicle tracks and the estimated (points) and actual (circles) feature positions. The right graphs show estimated (upper) and actual vehicle navigation error. In the first run, only dead-reckoning is available for navigation information. In the second run, one feature is present. In the third run, eight features are present. The presence of features not only reduces the navigation error, but also improves the uncertainty estimate.

the possible origins of the measurements and the disposition of the proposed features. Projected possible feature states are updated using the measurements. Third, assignments, or sets of hypotheses, are formed to explain, from a global perspective, what has transpired. These assignments are used to update the vehicle projected possible states using the measurements and the feature estimates. Finally, the track trees are pruned to remove from consideration unlikely possibilities and to improve the computational efficiency of the overall algorithm. Pruning does not, in general, take place based on the assignments for the current time step. Instead, decisions are delayed to allow accumulation of corroborative evidence.

As shown in Figure 2, three processes, measurement, navigation, and feature mapping, are combined to improve estimates and to provide global explanations for what has happened. In the stochastic map of Smith *et al.* [1990], vehicle and feature track models are combined in a single (large) state vector. In MHCML, vehicle and feature tracks are separated to enable a multiple hypothesis approach. In each time cycle, predicted vehicle states $X_{k|k-1}$ and feature states $\Xi_{k|k-1}$ are projected from the previous cycle. These predicted states are compared with sensor measurements Z using a gating function , to check for possible matches. Hypotheses Θ are formed to cover the following possibilities: a measurement has come from a feature, a measurement is spurious, a feature is not detected, or a measurement is the result of an unknown feature. Feature track states are updated $\Xi_{k|k}$ based on these hypotheses Θ, the predicted vehicle states $X_{k|k-1}$, and the measurements Z. The hypotheses Θ are combined to form global assignments Ω. Each assignment is a set of compatible hypotheses which explains the source of each measurement (a particular known feature, a new feature, or no feature) and the disposition of each feature (a particular measurement or a miss). Vehicle track states are updated $X_{k|k}$ based upon the possible global assignments Ω, the predicted feature track states $\Xi_{k|k-1}$, and the measurements Z. (At present, correlations between vehicle and feature states are ignored, see Uhlmann [1995] for a discussion of this issue.) A Bayesian framework is used to evaluate the likelihood of each of these states, hypotheses, and assignments [Reid, 1979]. Determination of what has, in fact, happened, is made at the global level by rejecting unlikely assignments. The result of this determination is a pruning of assignments. This produces, teleologically, a removal of unsupportable hypotheses and states.

2.1 Projection

At any given time, each track is represented by multiple possible state estimates. During projection, each (current) state produces a set of possible subsequent states. The process of projecting a state using a given dynamic model is encapsulated by an observer, which is the (implementation of the) model used to project the state. Our prior understanding of the feature ontology is captured in a Markov network of observers which capture the possible behavior(s) of the feature estimate using multiple state models [Smith *et al.*, 1997b]. Each observer in a network is linked to subsequent observers with *a priori* transition probabilities. This implementation allows delayed decisions about the structure of a feature's behavior based on a set of possible prior models. Each state marks its place within the observer network by knowing what observer produced it. Projection is then just propagation of the current state estimates through the observer network. In the simplest case, a feature observer network may be a single model linked to itself with probability one, in which case projection proceeds as in a Kalman filter with a fixed plant model. Likelihood of projected states is the product of the parent state likelihood and the transition probability between the observers used to project the parent and the projected state.

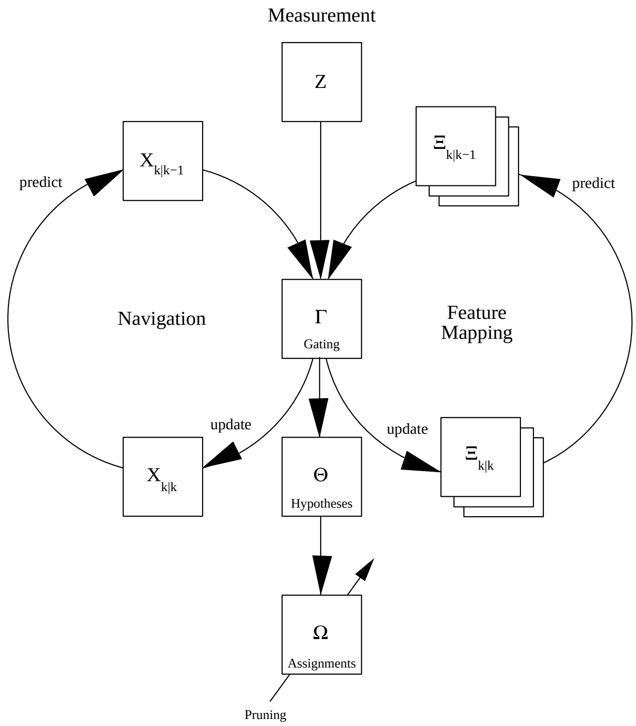

Figure 2: Process flow for MHCML. Information from relative measurements is used to update vehicle navigation and feature models simultaneously. Hypotheses and assignments are formed to allow *a posteriori* situational assessment at a global level.

2.2 Hypothesis formation and feature updating

When measurements are taken, hypotheses need to be formed to explain the underlying situation. A hypothesis, in general, is the association of a vehicle track, a feature track, and a measurement, asserting that a given measurement taken by a particular vehicle arises from a specified proposed feature. This is somewhat in contrast to the concept of target-to-track assignments in multitarget tracking [Mori *et al.*, 1986]. In that case, vehicle navigation uncertainty is not taken into account and multiple feature and/or measurement models are not permitted, leading to the conjecture that all target-to-track assignments are equi-probable. These assumptions are not made by MHCML, making target-to-track assignment inseparable from data-to-data assignment, and, in fact, central to the project of track estimation. The hypotheses formed represent the logical possibilities of the sensor physics based on out current understanding of the vehicle context. We follow (for the present) a somewhat typical approach within the multiple hypothesis tracking literature [Cox and Leonard, 1994]. For each measurement, we entertain the hypothesis that it may be spurious. For each feature, we entertain that it may have been missed. For each combination of projected possible vehicle state and measurement, we hypothesize a new proposed feature track. Finally, for each combination of projected possible vehicle state, projected possible feature state, and measurement, we consider a match hypothesis, which, as above, associates the vehicle, the feature, and the measurement. Match hypothesis formation is subject to two restrictions: state compatibility and gating.

In the development of multiple hypothesis tracking, it was observed that some data association decisions did not affect each other at all [Reid, 1979]. If the full set of possibilities were considered, many duplicate estimates would result from these orthogonal decisions. To address this, methods for clustering dependent decisions together (and more importantly separating independent decisions) were developed [Kurien, 1990]. When vehicle navigational uncertainty is introduced, uncertainty about what has happened to a feature is spread to the vehicle track during matching. Subsequently, this uncertainty is passed to other feature tracks as well. Thus, even features in separate 'clusters' quickly come to depend on each other's decisions about feature disposition. Clustering is no longer useful in maintaining sets of independent decisions. However, a full enumeration of the possibilities still contains many duplicates. Each branching in a feature track tree can be thought of as a decision to be made at a later time. While multiple branches still exist, the decision is active; once only one possibility remains, the decision is resolved. We want to make sure that (1) there will only be one state left for each track tree once all decisions are resolved, and (2) all states that are tracked are *possible* given some decision calculus. Each state maintains a set of the active decisions it depends on and the necessary resolution required for that state to hold. States are compatible if there are no active decisions upon which they both depend for which they require different hypotheses (decision resolutions) [Smith *et al.*, 1997b]. Checking state compatibilities serves to minimize the number of state combinations considered and, to the degree possible, separate orthogonal decisions.

In addition to restricting hypothesis formation based on state compatibility, a gating criterion is enforced. Although the probability density functions used (explicitly or implicitly) to model state estimates are everywhere nonvanishing, we recognize that unlikely matches will not provide improved state estimates, even if such matches are representative of the process being modeled. We only consider a match hypothesis when some combination of states and a measurement pass a gating test. Currently, the typical test is used [Bar-Shalom and Fortmann, 1988; Cox and Leonard, 1994]. The probability density function of the measurement innovation is estimated. The actual innovation is compared to this.

If the estimated probability of the innovation being no larger than the realized amount exceeds a threshold, the gate is successful and a match hypothesis is supported. Note that this is identical to restricting the Mahalanobis distance of the innovation from the origin in innovation space to be less than a specified value (calculated from the chi-squared shape of the innovation probability density function).

Feature tracks are updated by comparing the measurements, the set of projected possible feature states, and the set of projected possible vehicle states. Each feature state may be updated with a miss hypothesis. Each positively gating, compatible combination of feature state, vehicle state, and measurement provides an additional updated state. Note that this set of updated states is an exhaustive, mutually exclusive set of possible events, conditioned on the occurrence of the particular projected possible feature state, the set of measurements received, and the set of projected possible vehicle states. The gating probabilities (as given above) and the *a priori* probability of the feature track being missed provide the relative likelihoods of the updated states. The actual likelihood is the relative likelihood normalized and multiplied by the projected possible feature state likelihood. Updated states inherit the decision dependencies of their parent feature state, the matched vehicle state, if any, and the hypothesis dependency of the current gating decision.

2.3 Assignment formation and updating vehicle states

While each feature is restricted to one of a set of hypotheses about its disposition, the vehicle takes part in all of the measurements in a given time step. Because of this, vehicle track updating cannot take place on the basis of individual hypotheses, but rather must be based on sets of hypotheses spanning the decision space of the tracks and measurements for that time step. An assignment is an exhaustive, consistent set of hypotheses spanning this decision space. Each assignment will explain the disposition of each proposed feature and the origin of each measurement in a consistent manner. The first step in updating the vehicle track tree is the formation of all such possible assignments from the feature gating and measurement origin decisions. All assignments must be considered (in contrast to traditional multiple hypothesis tracking [Cox and Miller, 1993]), because (1) target-to-track associations are not equi-probable for MHCML and (2) the assignment likelihood calculation process requires consideration of all possible assignments. The efficient generation of possible assignments from a set of decisions has been addressed in some depth [Cox and Leonard, 1994].

Vehicle track tree updates originate from either the most recent set of projected possible states or the current set of updated possible states (just in case some other form of measurement, such as INS, has already been processed during the current time cycle). Each of these originating states is given a child state corresponding to each of the possible assignments. For each match hypothesis in an assignment, child states are produced by considering possible matches with projected possible states from the specified feature track. These updates are accomplished in the usual Kalman sense. The current implementation (suboptimally) updates the states sequentially; however, some algebra can be used to reconstruct the Kalman optimal update from these updated estimates and the projected possible feature states.

During vehicle update tree growth, state likelihood is calculated conditioned on the decision to accept the root likelihood, but not on the truth of that likelihood. Forming the assignment root state results in some base likelihood derived from the number of features detected and missed and the number of measurements which are false alarms and real targets. Match hypothesis updates are multiplied by their gating probability (as in feature track updating). The leaves of the vehicle track then form an exhaustive, mutually exclusive set of possible events. The *a posteriori* probability of each assignment is then the normalized sum of the leaf states to which it gives rise. The second stage of vehicle track update tree likelihood calculation begins by setting the assignment root state likelihoods to the product of the parent state likelihood and the *a posteriori* assignment likelihood. The update portion of the tree is then renormalized so that the leaves of any branch sum to the probability of the root state for that branch.

2.4 Pruning

The consideration of multiple hypotheses necessarily leads to rapid growth in the number of states to be considered. The basic multiple hypothesis tracking problem (when vehicle position is assumed to be known) grows as the exponential of the number of measurements received [Chang *et al.*, 1990]. The MHCML algorithm would, unchecked, grow quicker still. The viability of this algorithm is predicated on the timely resolution of delayed decisions. The rapid growth in algorithmic complexity means a fine line between increased explanatory power in decision-making and computational efficiency. Thus the efficient pruning of track trees to eliminate unlikely or unsupportable state estimates is quite important. Fortunately, a number of effective pruning strategies have been developed for multiple hypothesis tracking, including n-backstep, k-best, and threshold pruning [Kurien, 1990; Cox and Leonard, 1994]. MHCML can, in large part, adopt these pruning methods directly.

The essential fact is that pruning is carried out on the basis of assignments. Pruning is restricted to assignments to maintain the global framework of context explanation which is one of the chief theoretical advantages of MHCML. Assignment likelihood can be calculated by summing the likelihoods of the vehicle track leaf states which are compatible with the decision dependencies represented by the assignment. Once assignments are ranked by current support, a number of them may be rejected. This results in a subset of the hypotheses relating to the set of assignments for the time cycle in question becoming unsupported by any accepted assignments. Specific track states are then pruned if they depend on any of these rejected hypotheses. Subsequent, recursive consideration of additional hypotheses to reject is necessary to remove hypotheses which are wholly supported by states which have been pruned. Finally, the track trees are renormalized and the cycle can begin again with state projection.

3 Results

The MHCML algorithm described above has been completely implemented in C++ on SGI workstations and simulation testing is in progress. Figures 3 to 6 illustrate the algorithm's performances for a simple scenario in which a moving vehicle makes noisy range and bearing observations of a single feature. Simulation parameters are given in Table 1.

The vehicle starts out at (0,0) and moves north at 1 meter per second. The feature is stationary at the location (100,0). Dead reckoning measurements (depth, speed, pitch, and yaw) are taken every second. A single sonar measurement is received each second. The sonar measurement has the noise characteristics given in Table 1 95% of the time. Outlier measurements with 5 times the variance of these values are generated 5% of the time. Decisions are resolved to a single assignment at the previous time step (i.e., n-backstep pruning is being performed with n=1). This example shows how the MHCML algorithm converges to a single track in the presence of considerable noise.

4 Conclusions and Future Research

The three key elements of our multiple hypothesis approach to concurrent mapping and localization are:

1. explicit handling of vehicle position uncertainty,
2. delayed decisions based on multiple, competing hypotheses, and
3. a feature-based representation of the environment.

To mitigate the impact of vehicle position uncertainty, MHCML differs from conventional multiple hypothesis

Table 1: Simulation parameters.

parameter	value
range measurement variance	0.05 m^2
angle measurement variance	0.05 m^2
feature position process noise variance	2.0 m^2
probability of detection	0.90
probability of false alarm	0.05
vehicle depth measurement variance	0.05 m^2
vehicle speed measurement variance	0.25 m^2/s^2
vehicle pitch measurement variance	0.05 rad^2
vehicle yaw measurement variance	0.05 rad^2
vehicle position process noise variance	0.5 m^2
vehicle pitch process noise variance	0.05 rad^2
vehicle yaw process noise variance	0.05 rad^2
vehicle depth process noise variance	0.05 m^2
vehicle speed process noise variance	0.25 m^2/s^2

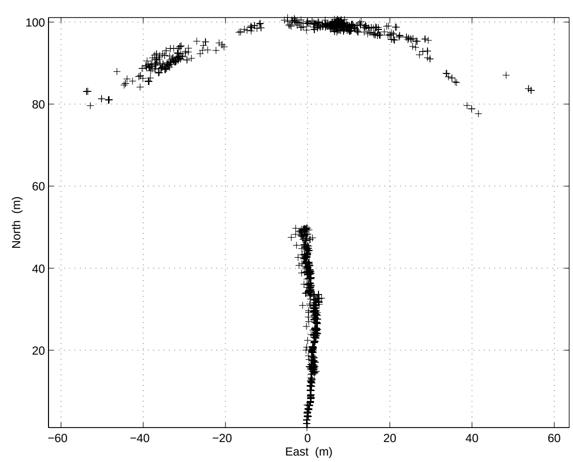

Figure 3: All vehicle and feature states for a simulated run with one feature. The vehicle starts out at (0,0) and moves north at 1 m/s. The feature is stationary at the location (100,0). 1954 total states are hypothesized, 747 are feature states and 1207 are vehicle states.

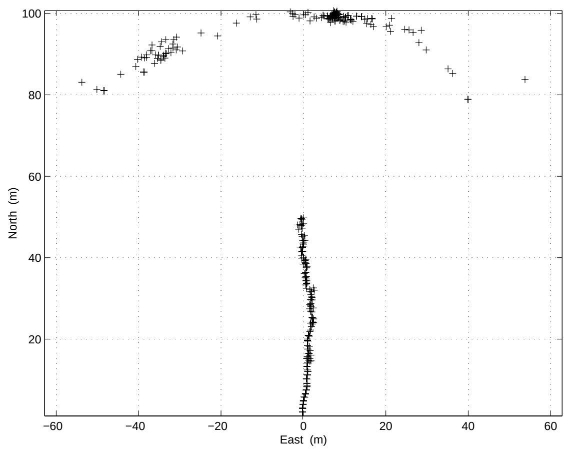

Figure 4: All vehicle and feature states that survive at least one pruning pass. 399 total states passed at least one cycle of pruning, 217 are feature states and 182 are vehicle states.

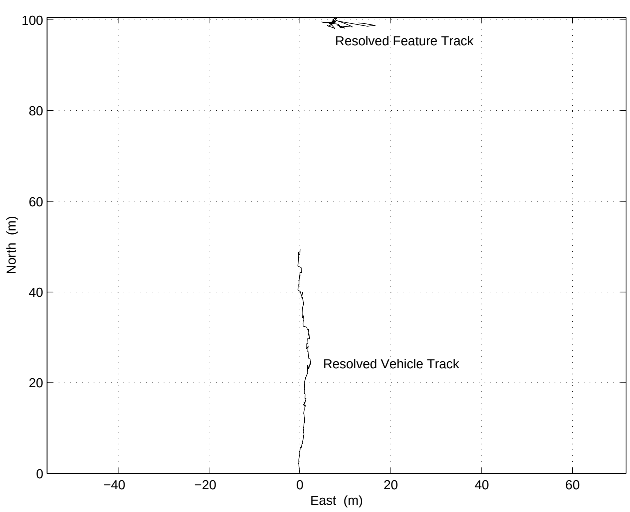

Figure 5: Resolved vehicle and feature tracks.

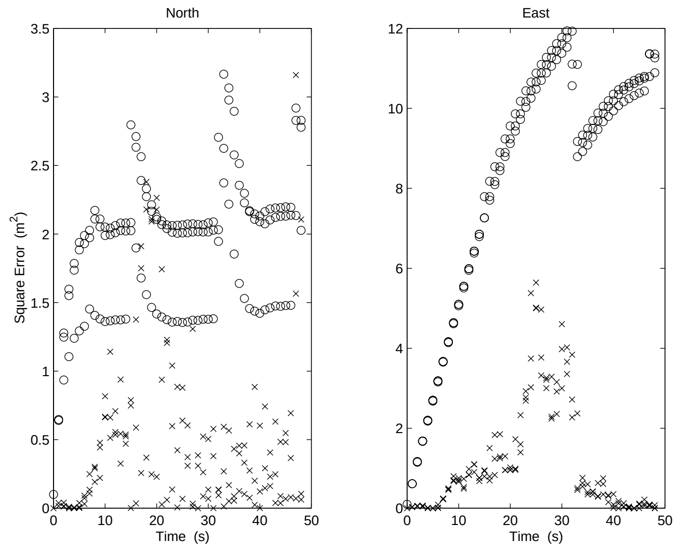

Figure 6: Estimated (circles) and actual (crosses) vehicle position squared error in the north and east directions. (This includes both projected and updated states.)

tracking in several significant ways. Feature ontology is generalized to a Markov network of observers. In this way, *a priori* knowledge about what features will be encountered and how information about these features will be gathered is fully and explicitly represented as a noumenal framework. Although the combination of a multiple-hypothesis approach and vehicle navigational uncertainty prevents the clustering or partitioning common to multitarget tracking, state estimate separability is maximized using decision dependency sets. In this way, orthogonal decisions do not result in increased track tree size. A new technique for assignment formation and likelihood calculation integrates target-to-track and data-to-data associations. New track updating procedures account for the inseparability of target-to-track and data-to-data associations. Finally, all pruning is accomplished at the level of global assignments. Global pruning is necessary both to maintain theoretical consistency at the state estimate level (as well as the track level) and to ensure proper *a priori* normalization when comparing estimate likelihoods. These particular algorithmic innovations (feature ontology, estimate separation, and global pruning) combine to make CML practicable.

Current research aims to validate the new MHCML implementation for more complex simulation scenarios and to perform post-processing of real-data from the US Navy HRA sonar sensor [Medeiros and Carpenter, 1996]. Topics for future research include investigation of methods for improved computational efficiency, experimentation in a controlled tank setting using a biomimetic underwater sonar system, and at-sea implementation onboard a NUWC unmanned underwater vehicle.

Acknowledgement

This research has been funded in part by the Naval Undersea Warfare Center, Newport, RI, U.S.A.

References

[Bar-Shalom and Fortmann, 1988] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association.* Academic Press, 1988.

[Bellingham *et al.*, 1993] J. G. Bellingham, C. Chryssostomidis, M. Deffenbaugh, J. J. Leonard, and H. Schmidt. Arctic under-ice survey operations. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 50–59, 1993.

[Chang *et al.*, 1990] K. Chang, S. Mori, and C. Chong. Performance evaluation of a multiple hypothesis multitarget tracking algorithm. In *IEEE Int. Conference on Decision and Control (CDC)*, pages 2258–2263, 1990.

[Chong and Kleeman, 1997] K. S. Chong and L. Kleeman. Sonar feature map building for a mobile robot. In *Proc. IEEE Int. Conf. Robotics and Automation*, 1997.

[Cox and Leonard, 1994] I. J. Cox and J. J. Leonard. Modeling a dynamic environment using a Bayesian multiple hypothesis approach. *Artificial Intelligence*, 66(2):311–344, April 1994.

[Cox and Miller, 1993] I. J. Cox and M. L. Miller. On finding ranked assignments with application to multi-target tracking and motion correspondence. Technical report, NEC Research Institute, July 1993.

[Elfes, 1987] A. Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, RA-3(3):249–265, June 1987.

[Geyer *et al.*, 1987] E. Geyer, P. Creamer, J. D'Appolito, and R. Gains. Characteristics and capabilities of navigation systems for unmanned untethered submersibles. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 320–347, 1987.

[Kurien, 1990] T. Kurien. Issues in the design of practical multitarget tracking algorithms. In Y. Bar-Shalom, editor, *Multitarget-Multisensor Tracking: Advanced Applications*, pages 43–83. Boston: Artech House, 1990.

[Leonard and Durrant-Whyte, 1991] J. J. Leonard and H. F. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, Osaka, Japan, 1991.

[Levine *et al.*, 1995] E. Levine, D. Connors, R. Shell, T. Gagliardi, and R. Hanson. Oceanographic mapping with the Navy's large diameter UUV. *Sea Technology*, pages 49–57, 1995.

[Medeiros and Carpenter, 1996] M. Medeiros and R. Carpenter. High resolution array signal processing for AUVs. In *AUV 96*, pages 10–15, 1996.

[Moravec, 1989] H. Moravec. Sensor fusion in certainty grids for mobile robots. In *Sensor Devices and Systems for Robotics*, pages 253–276. Springer-Verlag, 1989. Nato ASI Series.

[Mori *et al.*, 1986] S. Mori, C. Chong, E. Tse, and R. Wishner. Tracking and classifying multiple targets without a priori identification. *IEEE Transactions on Automatic Control*, AC-31(5), May 1986.

[Moutarlier and Chatila, 1989] P. Moutarlier and R. Chatila. Stochastic multisensory data fusion for mobile robot location and environment modeling. In *5th Int. Symposium on Robotics Research*, Tokyo, 1989.

[Nussbaum *et al.*, 1996] F. Nussbaum, G. Stevens, and J. Kelly. Sensors for a forward-looking high resolution auv sonar. In *AUV 96*, pages 141–145, 1996.

[Reid, 1979] D. B. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automatic Control*, AC-24(6), December 1979.

[Rencken, 1993] W. D. Rencken. Concurrent localisation and map building for mobile robots using ultrasonic sensors. In *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, pages 2192–2197, Yokohama, Japan, 1993.

[Schmidt *et al.*, 1996] H. Schmidt, J. Bellingham, M. Johnson, D. Herold, D. Farmer, and R. Pawlowcisz. Real-time frontal mapping with AUVs in a coastal environment. In *IEEE Oceans*, pages 1094–1098, 1996.

[Smith *et al.*, 1990] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. Cox and G. Wilfong, editors, *Autonomous Robot Vehicles*. Springer-Verlag, 1990.

[Smith *et al.*, 1997a] C. S. Smith, J. J. Leonard, A. A. Bennett, and C. Shaw. Concurrent mapping and localization for autonomous underwater vehicles. In *Undersea Defence Technology*, pages 338–342, 1997.

[Smith *et al.*, 1997b] C. S. Smith, J. J. Leonard, A. A. Bennett, and C. Shaw. Feature-based concurrent mapping and localization for autonomous underwater vehicles. In *IEEE Oceans*, 1997.

[Stewart, 1988] W. K. Stewart. *Multisensor Modeling Underwater with Uncertain Information.* PhD thesis, Massachusetts Institute of Technology, 1988.

[Tuohy *et al.*, 1996] S. T. Tuohy, J. J. Leonard, J. G. Bellingham, N. M. Patrikalakis, and C. Chryssostomidis. Map based navigation for autonomous underwater vehicles. *International Journal of Offshore and Polar Engineering*, 6(1):9–18, March 1996.

[Uhlmann, 1995] J. Uhlmann. *Dynamic Map Building and Localization: New Theoretical Foundations.* PhD thesis, University of Oxford, 1995.

[Vaganay *et al.*, 1996] J. Vaganay, J. G. Bellingham, and J. J. Leonard. Outlier rejection for autonomous acoustic navigation. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 2174–2181, April 1996.