An Architecture for Cooperative Localization in Underwater Acoustic Networks[*]

Jeffrey M. Walls
University of Michigan
Ann Arbor, Michigan
jmwalls@umich.edu

Ryan M. Eustice
University of Michigan
Ann Arbor, Michigan
eustice@umich.edu

ABSTRACT

This paper outlines an architecture for underwater acoustic cooperative localization. Our system leverages communication within an acoustic network to both share navigation information and measure the relative range between vehicles. We employ a factor graph framework for vehicle position estimation. The underlying structure of the factor graph formulation provides a low-bandwidth estimation framework that is tolerant to communication dropout. We detail the hardware and software used to implement a three vehicle cooperative network and provide a performance summary over several field trials.

Categories and Subject Descriptors

I.2.9 [**Robotics**]: Autonomous vehicles

1. INTRODUCTION

Underwater acoustic communication enables subsea vehicles to broadcast and receive command, control, and health information from surface operators. Additionally, time-of-flight (TOF) measurements of acoustic broadcasts between vehicles can be used to augment an underwater navigation framework. In this paper, we outline the design, implementation, and deployment of a system architecture for multiple vehicle operations that exploits acoustic communication for online navigation.

Underwater vehicles are unable to directly observe XY position; seawater is opaque to electromagnetic signals such that standard terrestrial systems like the global positioning system (GPS) are unavailable. Underwater vehicles integrate body-frame Doppler velocity, attitude, and depth to compute a dead-reckoned navigation estimate. XY position errors from dead-reckoned navigation, however, grow unbounded in time. Bounded-error navigation can be achieved using
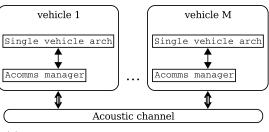
(a) Multiple vehicles exchange information over the acoustic communication channel and observe their TOF-based relative range.

(b) Independent navigation and control system onboard Iver AUVs. The cooperative localization process is highlighted in red.

Figure 1: Block diagram of a multi-vehicle system architecture. Single vehicle inter-process communication is handled via LCM. Communication between multiple vehicles is handled over the acoustic channel.

static acoustic beacon networks such as narrowband long-baseline (LBL) [1]. As vehicle network size increases, however, the rate at which each vehicle can query the LBL network decreases. Moreover, the range of vehicle operations is limited by the range of the LBL network.

With the addition of synchronous clock hardware, vehicles in an acoustic network can passively measure the one-way-travel-time (OWTT) of broadcast messages. Assuming a known sound velocity profile, receiving vehicles observe the relative range at their time-of-arrival (TOA) to the time-of-launch (TOL) position of the broadcasting vehicle. Unlike static beacon networks, OWTT-based underwater localization requires vehicles to share state information across the network. Solutions detailing information exchange and the resulting estimation frameworks have appeared throughout the literature including [2–8].

Below, we outline a recent distributed estimation framework for multi-vehicle OWTT cooperative localization [8]. The primary contribution of this paper includes the details
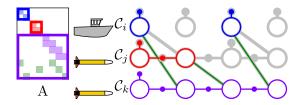
Figure 2: Example factor graph estimation framework and corresponding measurement Jacobian, A (left). Each row of pose nodes (large circles) represents a single vehicle (right). Small circles represent factors (observations). The full (centralized) graph is highlighted in gray, while the reconstruction on board the third (purple) vehicle is fully colored.

of the hardware and software processes employed within our system and their integration within a multi-vehicle architecture (Fig. 1). We also provide the cooperative localization functionality as an open-source software library. While the architecture described is applicable to a wide range of acoustic networks, we summarize an implementation on a three vehicle network. Finally, we present a performance summary over several field trials.

2. COOPERATIVE LOCALIZATION

Cooperative localization refers to the estimation of a vehicle state given relative position observations to other vehicles. Vehicles fuse measurements from onboard sensing (e.g., velocity and attitude) with external information such as relative vehicle observations and shared data. Herein relative position observations specifically refer to OWTT ranges.

The limitations of the acoustic communication channel [9] challenge the design of a cooperative localization system. Chiefly, low-bandwidth and unacknowledged data packets constrain each vehicle's ability to share information across the network. We show below that jointly estimating the full trajectory of each vehicle leads to a solution that can be efficiently distributed across an acoustic network.

2.1 Factor graph trajectory estimation

We apply a maximum-likelihood approach to jointly estimate the trajectory of each vehicle, that is, we compute the set of vehicle trajectories that best explains the set of observations. We use a factor graph to represent the joint probability distribution over vehicle trajectories. A factor graph is a bipartite graphical model composed of variable nodes (poses along a trajectory) and factor nodes (measurements). Factor graphs have become a standard within the robotics community for estimation, in particular, for the simultaneous localization and mapping (SLAM) problem [10]. For a more thorough derivation of the below framework, the interested reader is referred to [8].

Factor graphs represent a smoothing approach as opposed to standard state estimation, which only estimates the current vehicle state. The complete trajectory of the ith vehicle is represented by a set of rigid-body poses, $\mathbf{X}_i = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$. To minimize bandwidth in our distributed system, each rigid-body pose is represented as its XY position. The distribution over all poses can be factored given the set of all local observations, $\mathbf{Z}_i$, as

$$p(\mathbf{X}_i|\mathbf{Z}_i) \propto p(\mathbf{x}_1) \prod_i p(\mathbf{z}_{\mathrm{odo}_i}|\mathbf{x}_{i-1}, \mathbf{x}_i) \prod_j p(\mathbf{z}_{\mathrm{prior}_j}|\mathbf{x}_j). \quad (1)$$

We only consider two factor types: unary 'prior' factors (such as GPS) and binary 'odometry' factors (integrated velocity). The structure of the single vehicle factor graph is a chain. Herein, the ith vehicle's local factor graph will be referred to as factor chain $\mathcal{C}_i$. Poses along the chain with their associated factors are referred to as links. In Fig. 2, each row of pose nodes represents a separate vehicle chain.

When considering the factor graph over the full vehicle network, OWTT relative range measurements introduce constraints between the broadcasting vehicle TOL pose and the receiving vehicle TOA pose. Within the factor graph formulation, the joint distribution over the set of poses for each of M vehicles can be factored over the set of vehicle chains and relative ranges

$$p(\mathbf{X}_1, \ldots, \mathbf{X}_M|\mathbf{Z}_1, \ldots, \mathbf{Z}_M, \mathbf{Z}_r) \propto$$
$$\prod_{i=1}^M \underbrace{p(\mathbf{X}_i|\mathbf{Z}_i)}_{\mathcal{C}_i} \prod_k \underbrace{p(\mathbf{z}_k|\mathbf{x}_{i_k}, \mathbf{x}_{j_k})}_{\text{relative ranges}}, \quad (2)$$

where $\mathbf{Z}_r$ is the set of relative range factors. Fig. 2 illustrates the factor graph over a three vehicle network.

Each vehicle estimates its trajectory given the full set of factors. For zero-mean additive Gaussian noise models, the maximum likelihood estimate results in a nonlinear least-squares problem with linear subproblem

$$\min_{\mathbf{X}} \left\| \mathbf{AX} - \mathbf{b} \right\|^2, \quad (3)$$

where A is the sparse measurement Jacobian weighted by the square root information [10]. A is the matrix analogue of the factor graph (Fig. 2). Sparsity here implies efficient solutions.

We can exploit the factorization of the joint distribution over all vehicle poses, to arrive at a distributed solution. Within our estimation framework, each vehicle shares portions of its local chain with the network[1]. Properties of the constituent factors allow for low-bandwidth and dropout tolerant communication. In turn, each vehicle can then reconstruct a portion of the global factor graph using its local chain, the set of received chains, and the set of locally observed relative range factors (Fig. 2).

3. SYSTEM ARCHITECTURE

Our multiple vehicle system architecture spans two subsystems (Fig. 1): a single vehicle module containing sensor drivers, navigation, and control processes and a multiple vehicle module encompassing interaction between vehicles over the acoustic channel.

We discuss the system architecture as implemented on a three vehicle network consisting of two Ocean Server, Inc. Iver2 AUVs (Fig. 3) and a topside ship. While we make specific notes about this three vehicle network, the architecture is vehicle independent.

3.1 Single vehicle subsystem

Each vehicle executes several processes including sensor drivers, a pose estimator (Section 2), and, in the case of the AUVs, a vehicle controller. The vehicle software architecture is built around several open source projects, most notably

---

[1] The described approach only shares *local* vehicle information. Avenues for future work include forwarding information from the wider vehicle network.

Figure 3: Iver2 AUVs used within the described cooperative acoustic network.

lightweight communication and marshalling (LCM) [11] for inter-process communication. The single vehicle architecture is outlined in Fig. 1b and described below.

We rely on LCM to promote a modular software architecture in which sensor drivers, estimators, and controllers may be designed as separate processes. LCM allows processes to communicate over user datagram protocol (UDP) multicast within a publish/subscribe paradigm. Processes can publish user-defined message types to subscribers. LCM provides additional benefits including logging and playback utilities, a real-time messaging spy, and wrappers for several popular languages including C/C++, Java, and Python. Logging and playback are particularly well suited for robotics to aid with development and post processing.

We field two modified Iver2 AUVs detailed in [12]. Each AUV is driven by the Ocean Server controller, Underwater Vehicle Console (UVC). UVC provides an interface to a 'backseat' CPU [13] over a serial communication link. Our backseat Linux host is responsible for all sensor processing and acoustic communication. Additionally, the backseat executes a state machine and health monitor to automate data logging and handle vehicle behavior. All communication to the UVC is piped through the `Backseat driver` process (Fig. 1b), which translates between UVC and the LCM network.

Our topside vehicle executes sensor and acoustic communication software on a Linux host. The software stack onboard the topside vehicle is nearly identical to each AUV with the exception of the vehicle controller. We also run a viewer and acoustic communication user interface to visualize AUV state information provided by periodic acoustic broadcasts and submit control commands.

3.2 Multiple vehicle subsystem

The multiple vehicle subsystem includes all hardware and software processes that interface directly with the acoustic communication system. While LCM handles communication between processes on a single vehicle, the acoustic hardware and software handle communication between vehicles. Multiple vehicles interact via periodic acoustic broadcasts. Our architecture is similar to the one described by [14]. Each vehicle is equipped with a Woods Hole Oceanographic Institution (WHOI) Micro-modem and co-processor board [15], 25 kHz BTech Acoustics 2RCL transducer, and a synchronous-clock reference. The Micro-modem is configured to operate in a synchronous clock mode, using the user supplied pulse per second (PPS) signal to discipline its internal clock.

Stable synchronized clock references serve as the cornerstone for the acoustic network onboard each vehicle. The topside ship clock is disciplined by a Meinberg GPS timeserver. Each Iver2 AUV is outfitted with a PPSBoard [16], which provides a free-running low-drift PPS signal (aligned to a GPS time signal while at the surface) while the vehicles are subsea. The PPSBoard is built around a SeaScan, Inc., temperature compensated crystal oscillator which produces less than 1 ms drift over 14 h (corresponding to a ∼1.5 m bias in range observations). Newer commercially available free-running clocks provide improved performance. For example, Symmetricom provides less than 1 ms drift over 5000 h (∼208 d).

Our acoustic communication software, `Acomms manager` in Fig. 1a, is built around the C++ Goby acomms library [17] for medium access (MAC) scheduling and data quantization and serialization. During a mission, each vehicle executes a preconfigured fixed time division multiple access (TDMA) schedule. Each vehicle is assigned a slot, or several slots, within the TDMA period to broadcast messages. Slots are preconfigured for specific message types, for example, LBL pings or varying size data packets. We frequently use a simple TDMA in which vehicles broadcast a state data packet and a navigation data packet roughly once per minute. The Goby MAC library is responsible for signaling transmissions time so that messages are broadcast at the top of the second for synchronous communication.
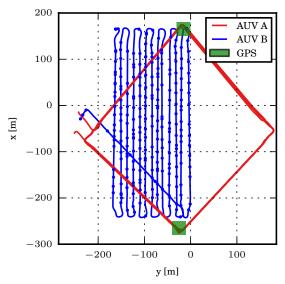
The Micro-modem with the co-processor board is capable of broadcasting various length frequency-shift keying (FSK) or phase-shift keying (PSK) data packets[15]. We typically use FSK rate 0 Micro-modem (one 32 Byte data frame) data packets to broadcast vehicle state and health to the topside ship. The topside ship uses a variety of 13 bit mini and 32 Byte data packets for control. We use rate 1 and rate 2 PSK data packets (three 64 Byte data frames) to relay navigation links for factor graph-based cooperative localization.

The Goby dynamic compact control language (DCCL) library [18] handles quantizing and marshalling data messages for acoustic communication. Goby relies on Google protobufs to specify the range and resolution of each field within a specific message type. Our acoustic communication process translates between LCM message types on the local vehicle and a DCCL type for broadcast. We employ message types for navigation data, state (e.g., estimated position, depth, battery health), and command (e.g., abort, abort-to-surface, jump to waypoint, manual override).
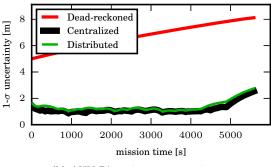
3.2.1 Estimator implementation

Our implementation of the distributed factor graph approach (Section 2) is built around the incremental smoothing and mapping (iSAM) [10] library and is released as an open-source C++ library[2]. iSAM updates the weighted measurement Jacobian to efficiently recover a navigation solution. The factor graph approach can easily include additional factor types in each vehicle's local graph. For example, we can incorporate ranges to LBL beacons by including LBL nodes along with the estimated vehicle trajectories. Running on relatively modest hardware, an Intel Core 2 Duo CPU, the `Estimator` process required approximately 3% CPU load during our field trials.

---

[2]http://robots.engin.umich.edu/SoftwareData/OWTT

(a) Relative vehicle trajectories (topside not shown).

(b) AUV-B's estimate uncertainty.

Figure 4: Summary of field trial and performance comparison. (a) An XY view of the vehicle trajectories. Blue dots indicate where AUV-B received range observations. (b) The smoothed uncertainty in each AUV-B pose as the fourth root of the determinant of the pose marginal covariance.

The cooperative localization process, `Estimator` in Fig. 1b, interfaces with the acoustic communication `Acomms manager` to determine when to add pose nodes to the graph and to share navigation links over the network. The `Estimator` packs odometry and GPS prior factors by subscribing to the sensor driver LCM stream.

At the TOL, the `Acomms manager` requests a data packet from the `Estimator`. The `Estimator` adds a TOL pose node to the graph and computes the navigation information to be broadcast. At the TOA of a data packet, the `Acomms manager` publishes the decoded navigation data and the OWTT range observation to the `Estimator`, which, in turn, adds the new factors and pose nodes to its graph. Following OWTT measurement updates, the `Estimator` publishes the current vehicle state estimate to the `Backseat driver` so that UVC can correct for navigation error. A similar interaction occurs when an LBL range observation is obtained. We leverage robust cost functions, specifically dynamic covariance scaling [19], within the factor graph to reduce the influence of outlier range observations.
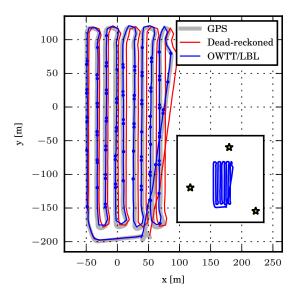

Figure 5: AUV trajectory (1 h) with topside OWTT and LBL support. Blue dots indicate where range observations were received. The inset plots LBL beacon locations.

4. FIELD TRIALS

The described multi-vehicle architecture has been successfully deployed in several shallow water field trials carried out at the University of Michigan Biological Station. We summarize two field trials below highlighting typical configurations for multiple vehicle operations. For convenience, we refer to each Iver2 AUV as AUV-A and AUV-B.

4.1 Two AUV deployment

The first trial represents a typical three vehicle deployment. The topside vehicle supported AUV-A and AUV-B during a 1.5 h mission. AUV-A followed a large diamond at a fixed depth of 8 m over AUV-B's lawnmower survey (fixed depth 6 m) while the topside vehicle drifted above the survey area (Fig. 4a). AUV-A had periodic access to GPS during brief surface intervals. The factor graph produced here is illustrated in Fig. 2. All vehicles remained within 500 m during the trial. Acoustic reception rates were asymmetric between the vehicles, ranging between 37–87%.

Fig. 4b shows the resulting uncertainty over trajectory poses for dead-reckoning, the post-process centralized estimator, and the distributed factor graph framework. Although AUV-B used only the local subset of range factors, its estimate was able to benefit from relative range observations and the difference compared to the centralized estimator is small.

4.2 Single AUV with LBL support

The second trial highlights the ability of our navigation architecture to include additional navigation constraints. In this trial, a single AUV, AUV-A, executed a lawn-mower survey at a constant 5 m depth. The topside ship provided OWTT support. AUV-A also interrogated the three-beacon LBL network roughly once per minute. 44% of acoustic broadcasts were successfully received between the vehicles. The estimated AUV position fuses only acoustic ranges and dead-reckoned navigation. During the trial, the AUV periodically received GPS for post-process comparison.

The estimated trajectory is shown in Fig. 5 along with the dead-reckoned only navigation result. We post-processed optimized the AUV trajectory with GPS observations to compare to the accuracy of the acoustic navigation solution. The on-line solution using acoustic ranges is able to closely follow the GPS optimized result and shows clear improvement of the dead-reckoned solution, which drifts approximately 10 m by the end of the mission.

5. CONCLUSIONS

As AUV network sizes grow, cooperative architectures will serve an important role in monitoring underwater vehicles. The cooperative localization system architecture described is extensible and vehicle independent. We have successfully deployed this system during several field trials with a three vehicle network and demonstrated the ability to augment dead-reckoned navigation with OWTT range constraints between vehicles. Moreover, the acoustic communication system allows topside operators to monitor AUV state and health as well as broadcast commands.

References

[1] L. L. Whitcomb, D. R. Yoerger, and H. Singh. Combined doppler/LBL based navigation of underwater vehicles. In *Proc. Int. Symp. Unmanned Untethered Subm. Tech.*, pages 1–7, Durham, NH, 1999.

[2] J. Vaganay, J. J. Leonard, J. A. Curcio, and J. S. Wilcox. Experimental validation of the moving long base-line navigation concept. In *Proc. IEEE/OES Autonomous Underwater Vehicles Conf.*, pages 59–65, Sebasco, ME, June 2004.

[3] M. F. Fallon, G. Papadopoulos, J. J. Leonard, and Nicholas M. Patrikalakis. Cooperative AUV navigation using a single maneuvering surface craft. *Int. J. Robot. Res.*, 29(12):1461–1474, 2010.

[4] R. M. Eustice, H. Singh, and L. L. Whitcomb. Synchronous-clock one-way-travel-time acoustic navigation for underwater vehicles. *J. Field Robot.*, 28(1): 121–136, 2011.

[5] S. E. Webster, J. M. Walls, R. M. Eustice, and L. L. Whitcomb. Decentralized extended information filter for single-beacon cooperative acoustic navigation: Theory and experiments. *IEEE Trans. Robot.*, 29(4):957–974, 2013.

[6] N. A. Cruz, B. M. Ferreira, O. Kebkal, A. C. Matos, C. Petrioli, R. Petroccia, and D. Spaccini. Investigation of underwater acoustic networking enabling the cooperative operation of multiple heterogeneous vehicles. *Marine Tech. Soc. J.*, 47(2):43–58, 2013.

[7] L. Paull, G. Huang, M. Seto, and J. J. Leonard. Communication-constrained multi-AUV cooperative SLAM. In *Proc. IEEE Int. Conf. Robot. and Automation*, pages 509–516, Seattle, WA, May 2015.

[8] J. M. Walls, A. G. Cunningham, and R. M. Eustice. Cooperative localization by factor composition over a faulty low-bandwidth communication channel. In *Proc. IEEE Int. Conf. Robot. and Automation*, pages 401–408, Seattle, WA, May 2015.

[9] J. Partan, J. Kurose, and B. N. Levine. A survey of practical issues in underwater networks. *ACM SIGMO-BILE Mobile Computing and Communications Review*, 11(4):23–33, 2007.

[10] M. Kaess, A. Ranganathan, and F. Dellaert. iSAM: Incremental smoothing and mapping. *IEEE Trans. Robot.*, 24(6):1365–1378, 2008.

[11] A. S. Huang, E. Olson, and D. C. Moore. LCM: Lightweight communications and marshalling. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, pages 4057–4062, October 2010.

[12] H. C. Brown, A. Kim, and R. M. Eustice. An overview of autonomous underwater vehicle research and testbed at PeRL. *Marine Tech. Soc. J.*, 43(2):33–47, 2009.

[13] J. DeArruda. OceanServer Iver2 autonomous underwater vehicle remote helm functionality. In *Proc. IEEE/MTS OCEANS Conf. Exhib.*, pages 1–5, Seattle, Washington, September 2010.

[14] S. E. Webster, R. M. Eustice, C. Murphy, H. Singh, and L. L. Whitcomb. Toward a platform-independent acoustic communications and navigation system for underwater vehicles. In *Proc. IEEE/MTS OCEANS Conf. Exhib.*, pages 1–7, Biloxi, MS, October 2009.

[15] L. Freitag, M. Grund, S. Singh, J. Partan, P. Koski, and K. Ball. The WHOI micro-modem: An acoustic communications and navigation system for multiple platforms. In *Proc. IEEE/MTS OCEANS Conf. Exhib.*, pages 1086–1092, Washington, D.C., September 2005.

[16] R. M. Eustice, L. L. Whitcomb, H. Singh, and M. Grund. Recent advances in synchronous-clock one-way-travel-time acoustic navigation. In *Proc. IEEE/MTS OCEANS Conf. Exhib.*, pages 1–6, Boston, MA, September 2006.

[17] T. Schneider and H. Schmidt. Goby-acomms version 2: Extensible marshalling, queueing, and link-layer interfacing for acoustic telemetry. In *IFAC Conf. Manoeuvring and Cont. Marine Craft*, pages 331–335, Arenzano, Italy, September 2012.

[18] T. Schneider and H. Schmidt. The dynamic compact control language: A compact marshalling scheme for acoustic communications. In *Proc. IEEE/MTS OCEANS Conf. Exhib.*, pages 1–10, Sydney, Australia, May 2010.

[19] P. Agarwal, G. D. Tipaldi, L. Spinello, Cyrill Stachniss, and Wolfram Burgard. Robust map optimization using dynamic covariance scaling. In *Proc. IEEE Int. Conf. Robot. and Automation*, pages 62–69, Karlsruhe, Germany, May 2013.