2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

A Measurement Distribution Framework for Cooperative Navigation using Multiple AUVs

Maurice F. Fallon, Georgios Papadopoulos and John J. Leonard

Abstract— **In recent years underwater survey and surveillance missions with more than a single Autonomous Underwater Vehicle (AUV) have become more common thanks to more reliable and cheaper platforms, as well as the addition of remote command and control communications using, for example, the WHOI acoustic modem. However cooperative navigation of AUVs has thus far been limited to a single AUV supported by a dedicated surface vehicle with access to GPS. In this paper a scalable and modular framework is presented in which any number of vehicles can broadcast, forward and acknowledge range, dead-reckoning, feature and GPS measurements so that the full fleet of AUVs can navigate and cooperate in a consistent and accurate manner. The approach is independent of the resultant application — such as recursive state estimation or full pose optimization. Trade-offs between the number of vehicles, the condition of the communication channel and rate at which updates are available are also discussed. Finally performance is illustrated in a realistic experiment.**

I. INTRODUCTION

Multiple Autonomous Underwater Vehicle (AUV) deployments are becoming more common as the technologies upon which the individual vehicles rely become more stable and the acoustic communications technology that they use to share commands and information becomes standardized. However there is a need for a navigation framework which makes a consistent and accurate estimate of the positions of the full fleet of AUVs available to each of the vehicles online. This approach would allow multiple inexpensive AUVs to share the capabilities of a single accurately instrumented vehicle or to capitalize upon a single vehicle surfacing for a GPS fix, so as to avoid the need for the other vehicles to surface. Deeper cognizant capabilities such as distributed task assignment and decision informed by these positions estimates would then be possible.

Such a system should be both distributed and scalable — both to vehicles entering and leaving the fleet as well as to changes in the rate of data transmission between the vehicles. Finally the approach should be flexible enough to allow the resultant data be used by any multi-vehicle recursive state estimation (e.g. Particle Filter, Extended Kalman Filter) or pose optimization algorithms.

A major complication in any marine environment is the communication channel. Communication using acoustic modems (such as the WHOI Micromodem [7]) is at a very low rate (as low as 32 bytes per 10 seconds) with a range of several kilometers in open water. It is unreliable and typically unacknowledged and the singl channel is shared amoungst

The authors are with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA
mfallon,gpapado,jleonard@mit.edu

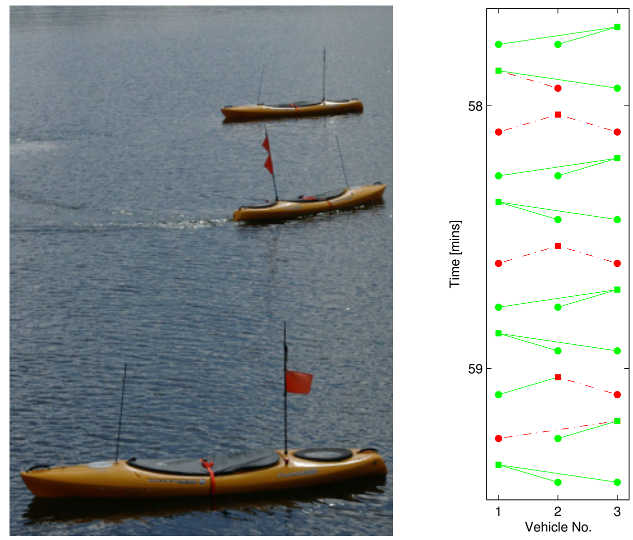

Fig. 1. Left: Three SCOUT kayaks which are used in our experiments as AUVs surrogates in testing the proposed measurement distribution framework. Right: An illustration of the typical multi-vehicle acoustic communication sequence with several failed tranmissions (illustrated in red) while green lines indicate successful tranmissions. Each vehicle transmits in turn during a 30 second TDMA (Time Division Multiple Access) cycle. Note that due to failed tranmissions between 57.5–59.5 mins, no data from Vehicle 2 reaches Vehicle 3. This would cause a filter reconstruction backlog on Vehicle 3.

obviously multiple vehicles. We propose a framework which fits these hugely demanding limitations. The framework is tested using experimental data collected using MIT's SCOUT kayaks (as illustrated in Figure 1) fitted WHOI acoustic modems but is intended to be utilized on an AUV fleet, such as the OceanServer Iver2s or Hydroid Remus 100s.

Section II will discuss previous cooperative navigation and localization research. Section III will discuss the specifics of the marine environment for this application. The core idea of the cooperative navigation algorithm — a measurement distribution framework — is outlined in Section IV. An experimental demonstration with 3 autonomous surface vehicles (serving as AUV surrogates) is outlined in Section VI. Finally conclusions and future work are presented in Sections VII and VIII.

II. COOPERATIVE NAVIGATION

Our previous work approached the problem with one or more surface vehicles employed in the role of a Communications and Navigation Aid (CNA) supporting an AUV. This vehicle had two roles, firstly functioning as a communication and control moderator for the underwater fleet, while sec-

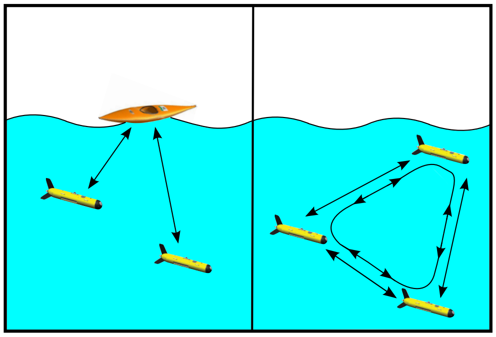

Fig. 2. Our previous work [5] has considered *top down* cooperative navigation with communication between a surface vehicle and one or more AUVs (left). The proposed framework is intended to allow multiple AUVs to communicate navigation information in a distributed and scalable manner — without either a surface craft or centralization (right).

ondly providing its own position information and an estimate of its range so that the AUV could use that information to better navigate.

Initial work used two such surface vehicles so as to estimate the full state vector of the AUV at once [13]. Subsequently the configuration was changed to a single CNA supporting an Iver2 AUV [1], [5], which can itself be extended in a straight forward fashion to supporting any number of AUVs within range of the transmission.

However this approach did not utilize information transmitted *between* the AUVs and required a surface vehicle providing uncorrelated measurements to the underwater vehicles to avoid overconfidence. An illustration of the difference between the two scenarios is illustrated in Figure 2.

A. Distributed Cooperative Navigation

The distributed localization problem has been studied in great depth by Roumeliotis and colleagues. Early work by Roumeliotis and Bekey [11] pioneered the concept of distributed localization based on recursive state estimation. Simulation results demonstrated that a group of vehicles measuring the distance to one another, but without access to global location estimate, can estimate their global position more accurately than any of the individual vehicles. Mourikis and Roumeliotis [9] performed a detailed performance analysis of cooperative localization. More recently, Nerurkar *et al.* [10] proposed a Distributed Conjugate Gradient maximum a posteriori algorithm for distributed localization, developing efficient methods to limit the communication cost and computational complexity for large multi-robot teams. Simulation results are presented for a team of 18 robots.

Our work targets the underwater environment, where severe communications constraints would make such an approach difficult to implement. Trawny *et al.* [12] investigated cooperative localization with limited communication with a quantized maximum *a posteriori* estimator. Their approach assumes that each robot transmits all its data to all other robots at each time step, with reliable bi-directional communication; it would be interesting to combine their quantized state estimator with the measurement distribution strategy described in this paper.

In the underwater domain, a number of authors have developed distributed localization approaches that are compatible with the capabilities of acoustic modems [1], [8], [4], [14]. In our earlier work, we proposed an approach for distributed localization in which the multi-vehicle navigaton filters are continually transmitted to and from each vehicle [2]. However such an approach, while reasonable for a few vehicles, does not scale well beyond 3–4 AUVs due to the large covariance matrices that need to be transmitted from vehicle to vehicle. Maczka et al. [8] approached the problem by transmitting only a scalar function of the main diagonal elements of the covariance matrix from one vehicle to the other. This simplification allows two vehicles to cooperatively navigate via a single 32–byte packet. However, the approach neglects inter-vehicle cross covariance terms, which as the authors acknowledge can lead to overconfidence and perhaps divergence. Eustice et al. [4] and Webster et al. [15] have investigated cooperative navigation between a surface ship and a single AUV, using the same WHOI micromodem that is utilized in our work.

Instead of trying to transmit all current covariance data to all vehicles at once, in this paper we propose a method which distributes the individual vehicle dead reckoning information to each vehicle before reconstituting the tracking filter on each vehicle. Before explaining the approach we will outline some of the details of the equipment that will be used.

III. Equipment

While the proposed solution is designed to be independent of the AUV platform and not to require any hardware beyond what is typically already installed, in this section we will outline the envisaged platform. While the experiments presented in Section VI utilize MIT's SCOUT kayaks, the proposed framework is intended for used on a low-cost platform such as the OceanServer Iver2. The Iver2's onboard board sensing will be limited to a compass, a depth sensor (allowing the problem to be reduced to two dimensions) and an acoustic modem and as such is a very low cost platform.

Our approach will utilize the WHOI Acoustic Modem [7], which uses low-rate Frequency-Shift Keying (FSK) or Phase-Shift Keying (PSK) to transmit small packets of information. Previously the basic 32 byte FSK packet was transmitted, however in future the Iver2 will use the 192 byte PSK packet which utilizes a newer co-processor extension board. In general for the same SNR, multipath environment and range this PSK message can send 4–6 times the data in the same packet length [6].

Transmission of a packet consists of two stages: first a *mini-packet* lasting 1.5 seconds is transmitted to initiate the communication sequence. The inter-vehicle range is a simple function of the time-of-flight of this mini-packet and the speed of sound in water (which is instrumented separately). The noise range standard deviation used in what follows will assumed to be $\sigma_r = 3$m. Following this, the information packet is transmitted in a process which lasts approximately 5-6 seconds. In all, it is prudent to reserve 10 seconds per transmission.

The time-of-flight is determined by initiating transmissions precisely at start of a second using a pulse-per-second (PPS) signal. The AUV contains a low-drift temperature-compensated timing board so as to maintain this synchronization for 10's of hours [4].

If one or more other vehicles are in the vicinity the message packet can be received by the transmitting vehicles and the inter-vehicle range can be estimated. However, there are two ways in which transmission can fail. Firstly, if the mini-packet is not properly decoded by the receiving modem, then neither the range nor packet will be received. Secondly, if the mini-packet is recognized properly but the longer main packet is not properly decoded then only the range will be successfully measured for that transmission.

One important point is that the range measurement will not be available to the transmitting vehicle. For the transmitting vehicle to learn of this range measurement the receiving vehicles must re-transmit the range as part of a subsequent packet.

Finally the control of the acoustic modem and the implementation of the distributed cooperative navigation algorithm takes place using the MOOS-IvP platform [3] on a 'backseat' computer at present, while the control missions described in Section VI are controlled separately on a 'frontseat' computer.

IV. Measurement Distribution Framework

A. Measurement Bookkeeping

In what follows we propose a measurement exchange system which allows for a fully consistent and distributed cooperative navigation solution. The approach remains within the strict limits of the underwater communication problem and remains flexible to varying communications conditions (changing inter-vehicle range, relative orientation and the presence of other noise sources in the vicinity) as well as accommodating scenarios such as vehicle surfacing.

Consider a dead-reckoning filter: at each time-step k a vehicle i will integrate sensor readings to form an independent cumulative state vector, dX_k^i, representing its dead-reckoned position change (and associated increase in uncertainty) during that time. Combining all of these vectors would allow the entire vehicle pose and uncertainty be reconstructed. Furthermore observations of objects with known or estimated location (in our case via range measurements) would allow this estimate to be improved and the uncertainty to be reduced.

Some failed message transmissions will occur when transmitting these values to other vehicles. As a result the receiving vehicles will not be able to reconstruct the pose of the transmitting vehicle. So as to circumvent this issue and to ensure that all dead-reckoning and range measurements are available to each vehicle's own multi-vehicle navigation filter, a 'bookkeeping' scheme will be used:

1) On each vehicle, all local and all received dead-reckoning measurements as well vehicle-to-vehicle ranges will be collated into a single database. This will also include range measurements between vehicles other than the local vehicle.
2) Where measurements have yet to be received by the local vehicle they will be requested as part of subsequently transmitted packets. The other vehicles will keep track of such requests and retransmit the requested messages until successful.
3) Only when the complete set of data up to a particular timestep is received by the local vehicle can any recursive state estimation or full pose optimization be carried out for that timestep.

Algorithm 1: Multi-Vehicle Measurement Distribution

```
for each vehicle do
    Integrate local cumulative dead-reckoning regularly
    if scheduled to surface then
        Log local dead reckoning to measurement table
        and reset cumulative DR filter
        Measure GPS location and insert into local
        measurement table
    else if scheduled to transmit data then
        Log local dead reckoning to measurement table
        and reset cumulative DR filter
        Choose useful dead reckoning data and
        outstanding request markets to form a packet
        Transmit packet using WHOI modem
    else if packet received successfully then
        Log local dead reckoning to measurement table
        and reset cumulative DR filter
        Measure vehicle-to-vehicle range using modem
        Insert range and data into local measurement
        table and update communications table
        Recompute best multi-vehicle pose if possible
    else if packet received unsuccessfully then
        Log local dead reckoning to measurement table
        and reset cumulative DR filter
        Measure vehicle-to-vehicle range using modem
        (if possible)
        Insert range into local measurement table
```

This will typically mean that the filter will operate in a partly delayed fashion as the required data simply cannot be gathered together instantaneously. Simple model based prediction can be used if necessary to best estimate the vehicle positions.

A general description of the method is presented in Algorithm 1 while a more specific illustration is presented in the following section.

B. Example Scenario

So as to illustrate the operation of proposed measurement distribution algorithm, consider the following example sequence of three iterations of the transmission algorithm. Further explanation is presented in Figures 3 and 4.

A fixed cycle of vehicle transmissions is decided upon (typically a simple repeating loop) and the vehicles begin

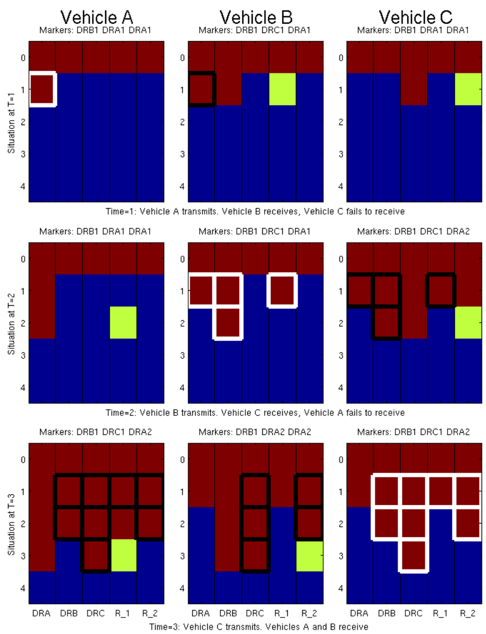

Fig. 3. Evolution of the multi-vehicle ledger system for three successive transmissions — one from each of the 3 vehicles. Vertical blocks represent the evolution for a particular vehicle. Individual blocks show the information available to a particular vehicle at a certain time. Red represents known measurements, blue unknown measurements (even if in the future) while green indicates a recent range measurement. A white box indicates the contents of a transmitted packet and a black box indicates the measurements from the packet received at another vehicle. As time proceeds, the individual ledgers are filled in with incoming data until all the data from a particular timestep becomes available to each vehicle. See Section IV-B for a fuller explanation.

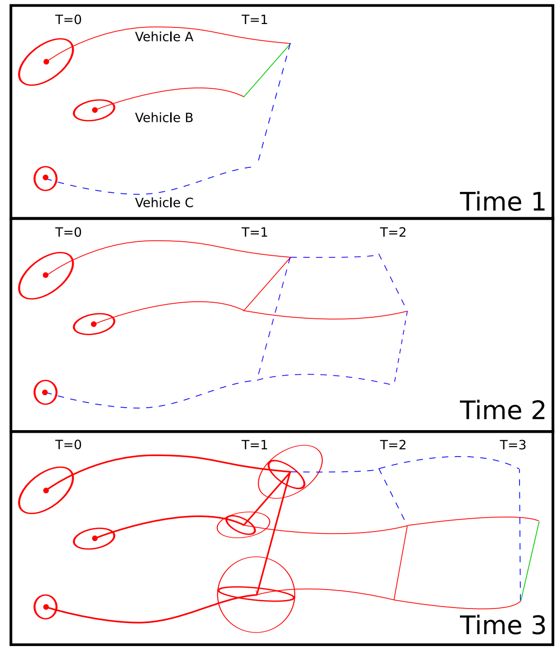

Fig. 4. Development of the system discussed in Section IV-B — from the point of view of **Vehicle B**. The required dead-reckoning and range measurements are gradually pieced together when data packets are received (as shown by solid red lines). Note how the complete multi-vehicle data for timestep 1 is available (at $t = 3$), at which time a filtering algorithm can be used to integrate all the measurements (as indicated in bold red). Portions of the ranging and dead-reckoning data that have yet to be received at $t = 3$ are indicated by dashed blue lines.

their respective missions and dive underwater. Obviously each vehicle will learn its own cumulative deadreckoning as time progresses, but will lack that of the other vehicles as well as some of the range measurements. This data will be disseminated using the proposed algorithm.

Illustrated in Figure 3 is the way that each of the three vehicles would build up their ledger of measurements and request data unknown to that vehicle. Each row represents the development of a particular vehicle's measurement ledger from transmission time $t = 1$ to $t = 3$. As measurements are learned the associated entries are filled in.

Initially at $t = 0$, the positions and associated uncertainty of each vehicle is known to all the others. At $t = 1$ Vehicle A will transmit its dead-reckoning from $t = 0$ to $t = 1$ (identified as as DRA1) as well as a markers indicating the earliest measurements unknown to each of the vehicles (which are DRB1, DRA1, DRA1 respectively). We suppose in this example that Vehicle B receives the packet and enters that data into its ledger. It will also note that Vehicle A requires the measurements from DRB1 and thereafter. As mentioned previously, by way of the transmission process the range between Vehicle A and B is measured — but only at Vehicle B. Simultaneously we suppose that Vehicle C fails to receive the transmission, so only the range measurement to Vehicle A (or the failure to measure the range) can be added to Vehicle C's ledger.

Continuing to $t = 2$, Vehicle B transmits all the data that it deems to be useful to the other vehicles, as well as its request markers (DRB1, DRC1 and DRA1). We suppose that Vehicle A fails to receive, but that Vehicle C does receive. Using this received data, Vehicle C now has all the dead-reckoning and ranging measurements up to $t = 1$ and can use any filtering or smoothing algorithm to estimate the multi-vehicle positions or poses for that short segment. Note also that its request markers have been updated to acknowledge that Vehicle B requires only the measurement DRC1 or later data.

Next at time $t = 3$, Vehicle C transmits data and we suppose that both Vehicle A and B receive the packet. Both vehicles can now compute the multi-vehicle trajectory up to $t = 1$, while Vehicle A can do so up to $t = 2$. Observe that the dead-reckoning of Vehicle B, DRB1 and DRB2, was first received by Vehicle C and then forwarded to Vehicle A indirectly. At this time-step DRA1 was not transmitted as Vehicle C was aware that it was already known to Vehicles B and A.

In summary using this algorithm the vehicles transmit, forward and acknowledge measurement data and in doing so gradually fill in the table of dead-reckoning and range measurements until all required data is available to each of the vehicles. This example makes a couple of simplifications for the sake of clarity (size of packets, variable indexing) but as presented illustrates the core algorithm concept.

C. Avoiding Data Buildup: Keyframes

An issue with the setup as proposed above is that should measurements be created at a rate which exceeds the rate at which they are distributed to the other vehicles a backlog will build.

To avoid such a backlog it is necessary to modify this simple system to be flexible to the rate of data transmission. We will instead focus on the reconstruction of a less detailed estimate of the multi-vehicle pose using what we will call 'keyframes'. Instead of attempting to reconstruct the fully detailed pose synchronized with our transmission cycle (a transmission every 10 seconds), we instead focus on building a less detailed version (in the case of a keyframe rate of 4, with poses separated by 40 seconds) and using the other transmission slots solely to distribute and marshal that data, as well as leaving aside bandwidth for the usual command and control functionality unrelated to this measurement system.

To do so requires the calculation of the rate at which measurements are created and the effective rate at which they are distributed. If each datapacket can contain N_p numbers and any packet transmitted by a communicating vehicle has a probability of P_p of being received. This will give an effective rate of data transmission of

$$N_{p,\text{eff}} = N_p P_p \tag{1}$$

Furthermore, if for each independent update of the tracking system N_k measurements are generated and must be distributed to the set of vehicles, to avoid a backlog of messages, the keyframes need to be spaced so that $P_{\text{keyframe}} N_{p,\text{eff}} \geq N_k$.

However this estimate assumes an up-to-date and global knowledge of which packets are required by which vehicles at any time. As illustrated in Section IV-B, until a measurement is acknowledged by the other vehicles, the transmitting vehicle is forced to continue retransmitting messages which perhaps have already been received, but this is unavoidable if we require the complete pose trajectory.

An alternative probabilistic approach would be for retransmission to occur for as long as it is likely that measurements have yet to be received (using the estimated rate of successful transmission). Thereafter transmission of that piece of data would pause until acknowledgment can be determined (See Section VIII).

For these reasons the above calculation instead represents a lower bound for the keyframe spacing

$$P_{\text{keyframe}} \geq \left\lceil \frac{N_k}{N_{p,\text{eff}}} \right\rceil \tag{2}$$

In the experiments that follows the keyframes spacing is determined experimentally and was found to be significantly higher than this value. Furthermore it has been experimentally observed that the probability of successful transmission, P_p, is a complex time-varying function of depth, range, conditions and orientation. As such the calculation of the optimal spacing is a theoretical rather than a practical bond.

D. Typical Packet Contents

Thus far, we have avoided precisely detailing the measurement data we wish to share so as to maintain generality. In this section we will outline the specific data which will be distributed for our envisaged application.

We assume a multi-vehicle state vector in three dimensions per vehicle. The incremental state vector of the ith AUV of N_i vehicles moving, at time k', will

$$d\mathbf{X}^i_{k'} = [dx^i_{k'}, dy^i_{k'}, d\theta^i_{k'}]^T \tag{3}$$

with an associated 3x3 block of the full covariance matrix, $d\mathbf{P}^{ii}_{k'}$ while an accurate global estimate of the vehicle depth, $x^i_{k'}$, will be known at all times and will allow the simplification of the problem to two dimensions in the horizontal plane. This filter will be reinitialized at the beginning of each keyframe.

The vehicle will integrate forward and starboard velocity estimates, $v^i_{k'}$ and $w^i_{k'}$, and a heading estimate from a compass, $\theta^i_{k'}$ so as to propagate this incremental estimate as follows

$$dx^i_{k'} = dx^i_{k'-1} + \triangle_{k'}(\hat{v}^i_{k'} \cos\theta^i_{k'} + \hat{w}^i_{k'} \sin\hat{\theta}^i_{k'}) \tag{4}$$

$$dy^i_{k'} = dy^i_{k'-1} + \triangle_{k'}(\hat{v}^i_{k'} \sin\theta^i_{k'} - \hat{w}^i_{k'} \cos\hat{\theta}^i_{k'}) \tag{5}$$

$$d\theta^i_{k'} = \hat{\theta}^i_{k'} \tag{6}$$

as well as updating the covariance block in a similar manner. The sensors which generate these estimates will have associated measurement uncertainties $(\sigma_{v_i}, \sigma_{w_i}, \sigma_{\theta_i})$, though the quality of these estimates may vary from vehicle to vehicle. The prediction step will be carried out at a relatively high frame rate ($\triangle_{k'} = 5$Hz) compared to the rate of the transmission system (one transmission per 10 seconds) and the keyframe rate (multiples of the transmission rate) and will be integrated in the period between two keyframes, $k-1$ and k, so as to form a cumulative dead-reckoning estimate for the local vehicle for that block of time.

As the 3x3 block of the covariance matrix is symmetric, only 9 numerical values will be required to represent the movement of a vehicle between two keyframes (with the vehicles operating at a known depth). To reconstitute the full multi-vehicle propagation step requires sharing these values across each of the vehicles. This means that the amount of information to be shared increases linearly with the number of vehicles — which allows for reasonable scaling of the proposed solution.

In addition, inter-vehicle ranges will be measured by receiving vehicles at each keyframe (or will note a null value if transmission is unsuccessful). This results in a $N_i - 1$ range values and hence the total number of measurement values generated for each keyframe will be

$$N_k = 9N_i + N_i - 1 = 10N_i - 1 \tag{7}$$

As well as the raw values of the measurements, each packet requires some overhead to explain what the measurements correspond to, as well as the data request markers mentioned in Section IV-B. A preliminary examination suggests that it is necessary to assign 4-6 bits per number to this task, depending on the method used, although a fuller examination will be carried out in future work.

In summary, using a suitable spacing of the keyframes, this approach allows for a completely scalable and distributed tracking filter — trading off the spacing of multi-vehicle poses against the number of vehicles in the system and the effective rate of data transmission.

V. Applying the Data: Multi-Vehicle EKF

Having received the required measurement data from each of the vehicles, it can then be used as the input of a sequential state estimator (such as a particle filter or an EKF) or used to optimize the entire multi-vehicle pose (using a technique such as conjugate gradient or matrix factorization). The proposed filter is recomputed in its entirety on each vehicle computer and as the measurement data is identical the resultant position estimates will be identical[1]

In this section we will discuss a specific implementation of a multi-vehicle Extended Kalman Filter (EKF) using distributed measurement data. For the N_i vehicle scenario, tracking in the X, Y and θ dimensions, the 3xN_i dimensional state vector will be as follows, at time k,

$$\mathbf{X}_k = [\mathbf{x}_k^1, \mathbf{y}_k^1, \theta_k^1, \ldots, \mathbf{x}_k^{N_i}, \mathbf{y}_k^{N_i}, \theta_k^{N_i}]^T \quad (8)$$

with an associated $3N_i$x$3N_i$ covariance matrix, P_k. This filter will be initialized using the known position estimates of the full set of vehicles — typically just before submerging. Each time a complete portion of the multi-vehicle cumulative deadreckoning is received at a vehicle, it is used to predict the filter.

Similarly when the corresponding range measures become available the multi-vehicle correction step is carried out. In the three vehicle case, where vehicle 2 and 3 receive a data packet from vehicle 1, the observation matrix would be given by the following Jacobian, H_k:

$$\begin{bmatrix} (x^1_{k|k-1} - x^2_{k|k-1})/d^{12} & (x^1_{k|k-1} - x^3_{k|k-1})/d^{13} \\ (y^1_{k|k-1} - y^2_{k|k-1})/d^{12} & (y^1_{k|k-1} - y^3_{k|k-1})/d^{13} \\ (x^2_{k|k-1} - x^1_{k|k-1})/d^{12} & 0 \\ (y^2_{k|k-1} - y^1_{k|k-1})/d^{12} & 0 \\ 0 & (x^3_{k|k-1} - x^1_{k|k-1})/d^{13} \\ 0 & (y^3_{k|k-1} - y^1_{k|k-1})/d^{13} \end{bmatrix}^T \quad (9)$$

where $d^{ij} = \|\mathbf{X}^i_{k|k-1} - \mathbf{X}^j_{k|k-1}\|$ is the Euclidean distance between two estimated vehicle positions i and j. If instead a range measurement was determined between vehicle 1 and 2 but not to vehicle 3, the observation matrix would will be modified accordingly.

A. Vehicle Surfacing

While the rate at which the multi-vehicle filter degrades will be slower than that of the single vehicle case, without access to a global landmark the multi-vehicle filter will become increasingly uncertain. At some (typically prearranged) point we propose that one of the AUVs will surface to access the GPS.

[1] Using this assumption of identical distributed position information, each individual vehicle's mission planner can then independently follow the same decision making process without the need to communicate.

This will obviously reduce the uncertainty of the surfacing vehicle: by sharing its very accurate position estimate with the other vehicles the fleet will gain a reduction in uncertainty via a simple EKF correction. The GPS measurement will simply be a direct observation of a portion of the state vector. These two numbers can be shared with the other vehicles in the same way as the range and dead reckoning.

More specifically we are interested in the scenario in which an inexpensive AUV with cheap on-board sensors surfaces to measure its GPS position before diving and sharing its position with other much more expensive AUVs with expensive dead-reckoning units — which can then themselves become confident of their positions. In this way the expensive vehicles need not halt their mission during a run or surface in a potentially hostile environment. See Section VI for an experimental simulation of this concept.

VI. Experiment

So as to demonstrate the proposed concept, a realistic experiment was carried out on the Charles River beside MIT. Three of the MIT SCOUT autonomous kayaks were used (see Figure 1). Each of the vehicles had a compass, GPS receiver and WHOI modem as well as access to common timeserver and a precise pulse-per-second trigger (via GPS) which allowed us to establish fully synchronized clocks on each vehicle and to carry out one-way-ranging between the vehicles.

Having synchronized the clocks, a (de-centralized) TDMA cycle was established. Each vehicle was assigned a one-packet ten-second transmit slot in a repeating 30 second cycle and listened for messages during the other two slots.

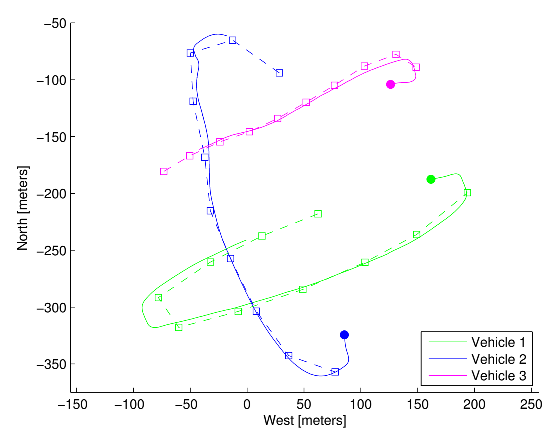

Fig. 5. Paths travelled by each vehicle (6 minutes of the 70 minute mission). The smooth line represents the GPS position of each vehicle, with the current location indicated by a circle. The dashed line segments illustrate the estimated path as reconstructed using a multi-vehicle EKF as described in Section V. The covariance of the estimated path is omitted for clarity. With a data transmission every 10 seconds and keyframes spaced every 4 transmissions, each multi-vehicle pose update is spaced 40 seconds apart. Note also that the reconstructed pose lags the current position by approximately 30-40 seconds in this figure. See Section VI-A.

During the experiment, modem messages were transmitted between each vehicle so as to determine package transmission statistics: (1) was a particular message received (2) at which vehicle (3) what was the measured inter-vehicle range. Furthermore when a modem transmitted, the precise transmission time was recorded by the transmitting vehicle. When a modem packet was received, the precise receiving time was recorded by the local vehicle. This times were used to establish the start and end points of the dead-reckoning integrations, as discussed in Section IV-C. In total over 420 messages were transmitted between the vehicles, at the lower FSK rate (due to a lack of availability of the higher rate modems).

For the 70 minute experiment 93% of transmissions resulted in successful range estimation while in 80% of cases the data packet was also successful transferred, averaging across all vehicles. This performance, better than in our previous experiments, was aided perhaps by calm weather conditions and little river traffic.

Each vehicle carried out a series of pre-planned overlapping loops of approximately 800-1000 meters in length (Figure 5). In total the vehicles carried out 7, 5 and 4 of their respective loops, thus travelling several kilometers each. Towards the end of the experiment the loops of Vehicles 1 and 2 were lengthened while Vehicle 3 floated in the center of the location. With increased range between the vehicles, a greater proportion of message transmissions failed. The effect this had on the performance of the message distribution algorithm is discussed in the following section.

While the experiment was obviously not carried out on underwater vehicles, we believe that the approach taken does not in any way modify the constraints of the communication and navigation systems. An advantage of using the kayaks is that GPS ground truth was available at all times.

A. Results

Having collected the experimental data, the algorithms discussed in this paper were applied to the data log. This allows us to experiment with different packet encoding techniques, priority systems, measurement quantization level and to compare the effect on the reconstructed vehicle poses.

For the scenario presented in Figures 6 and 7, vehicle transmissions were every 10 seconds while keyframes occurred every 4 transmissions, thus spacing multi-vehicle poses/keyframes 40 seconds apart. Each transmission cycle, the transmitting vehicle's message distribution system encoded 24 numbers — either cumulative dead-reckoning, range or GPS measurements. Assuming the PSK packet size of 192 bytes, this allows 6 bytes per number and 48 for message indexing and the request of unknown data.

While the intention of this experiment was to illustrate the measurement distribution concept, nonetheless Figure 6 illustrates that having reassembled the measurements at the remote vehicle, the data values successfully represented the multi-vehicle navigation. Again we emphasize that each vehicle has access to the identical multi-pose online, allowing distributed decision making.

Figure 7 illustrates the performance of some of the statistics of the experiment. The upper figure shows that distributing 24 numbers per transmission was sufficient to keep pace with the rate of measurement creation, except for the portion at the end of the experiment when a backlog formed. This issue will be studied as part of future work. Directly related to this backlog is the amount of lag time between the current time and most recently available update of the EKF. The system typically operated with a lag of 50-60 seconds while the data was distributed and marshalled.

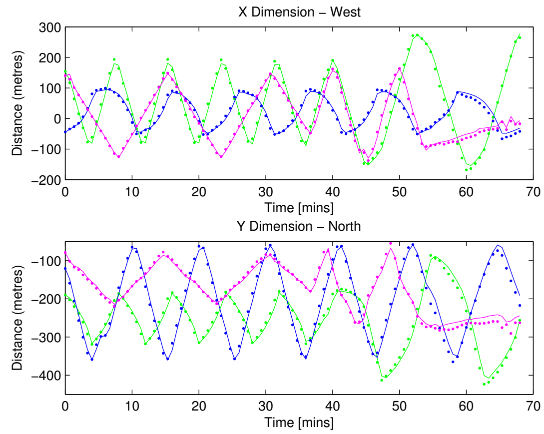

Fig. 6. Navigation of each vehicle in X and Y dimensions. The solid lines represent the GPS locations of each vehicle in the X (west) and Y (north) dimensions. Each dot represents the multi-vehicle pose, as reconstructed via the measurement distribution framework. It is emphasized that this information was available to each vehicle online. Note the 3 minute gap at the end of the mission is illustrative of the reconstruction lag that occurs while recent messages are distributed throughout the fleet.

The first 12 minutes of the lower figure illustrates the core benefit of cooperative navigation. The sensor noise of the overall cooperating vehicle fleet is reduced below what of the individual vehicles and the rate of uncertainty grows at a shorter rate as a result.

Secondly, the error and uncertainty of Vehicle 2 was reduced by a simulated GPS fix occurring with a 20 minute period. The fix position was shared to each vehicle, just as any another measurement, and causes the uncertainty of Vehicle 1 and 3 to be reduced — thus sharing the quality navigation filter of Vehicle 2 with all present vehicles.

Regular surfaces of only Vehicle 2 can allow the uncertainty of the entire vehicle fleet to be bounded. This allows Vehicles 1 and 3 to remain 'submerged' and to continue theirs missions *in situ* for as long as their battery life permits.

VII. CONCLUSIONS

In this paper we presented a framework for a distributed measurement communication system for an extremely low-data rate multi-vehicle system suitable for deployment on a fleet of Autonomous Underwater Vehicles. The proposed system is flexible to different fleet sizes, communication rates and communication environments, as well as harmonizing with the usual command and control communication currently used.

An application of the system was then outlined and experimentally tested in which the distributed measurements

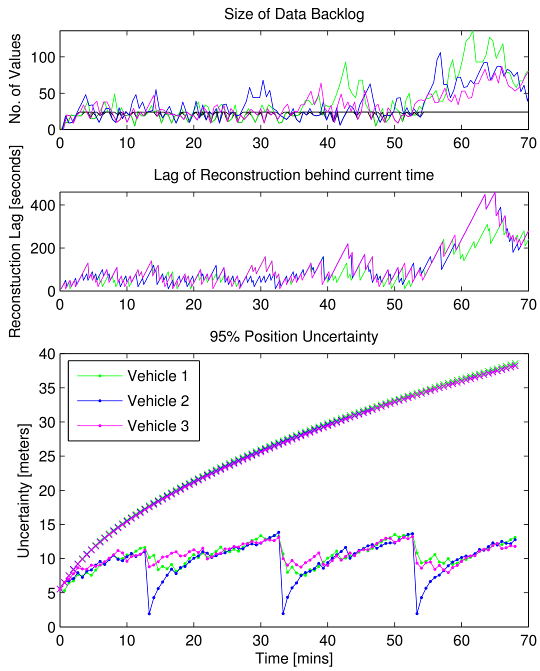

Fig. 7. Performance Statistics of the experiment discussed in Section VI. In the first portion of experiment, despite occasional transmission failures each vehicle maintain a fixed backlog of navigation messages (top figure). Each vehicle reconstructs the multi-vehicle pose a consistent 50-70 seconds behind the current time (center figure). Using an EKF, the reconstructed measurement set combines dead-reckoning and inter-vehicle ranges to reduce the rate of uncertainty growth (dotted lines) below that of the vehicles functioning individually (crosses). Finally Vehicle 2 'surfaces' every 20 mins, which allows the full vehicle fleet to bound their uncertainty to 10 meters (lower figure).

were integrated, as part of a multi-vehicle Extended Kalman Vehicle, to allow each vehicle to estimate the full multi-vehicle pose in real-time during the mission. It is emphasized again that any filtering or smoothing algorithm could be applied to the distributed data.

VIII. FUTURE WORK

The proposed framework has not considered flexibility to allow vehicles to enter and leave the network. We propose that a single cycle in the cycle be continually left open — either for emergency external commands, for a new vehicle to join the network or a vehicle to indicate that it will depart the network. Future work will examine this circumstance more closely.

In the static keyframe scenario proposed, it would be required to choose the keyframe spacing in advance which in turn would require an accurate and stable estimate of the transmission channel. If the quality of the transmission channel were to deteriorate during the mission, each vehicle would gradually build up a backlog of data. Future work will investigate how the spacing can be dynamically determined.

From an experimental view point, steps will continue towards implementing the framework within the MOOS-IvP platform, [3], and its testing on a number platforms and in a number of different concept scenarios.

IX. ACKNOWLEDGMENTS

This work was supported by ONR Grant N000140711102 and by the ONR ASAP MURI project led by Princeton University. Georgios Papadopoulos was funded by the Singapore National Research Foundation (NRF) through the Singapore-MIT Alliance for Research and Technology (SMART) Center for Environmental Sensing and Modeling (CENSAM).

The authors wish to thank Andrew Patrikalakis, Michael Benjamin, Joseph Curcio and Keja Rowe for their help with the collection of the experimental data.

REFERENCES

[1] Alexander Bahr, John J. Leonard, and Maurice F. Fallon. Cooperative localization for autonomous underwater vehicles. *The International Journal of Robotics Research*, 28(6):714–728, 2009.

[2] Alexander Bahr, Matt Walter, and John J. Leonard. Consistent cooperative localization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3415–3422, May 2009.

[3] Michael R. Benjamin, John J. Leonard, Henrik Schmidt, and Paul Newman. An overview of moos-ivp and a brief users guide to the ivp helm autonomy software. (CSAIL-2009-028), June 2009.

[4] Ryan M. Eustice, Louis L. Whitcomb, Hanumant Singh, and Matthew Grund. Experimental results in synchronous-clock one-way-travel-time acoustic navigation for autonomous underwater vehicles. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4257–4264, Rome, Italy, April 2007.

[5] Maurice F. Fallon, Georgios Papadopoulos, and John J. Leonard. Cooperative AUV navigation using a single surface craft. In *Proceedings of 7th Conference on Field and Service Robotics (FSR)*, July 2009.

[6] Lee Freitag. Personal Communication, 2009.

[7] Lee Freitag, Matt Grund, Sandipa Singh, James Partan, Peter Koski, and Keenan Ball. The WHOI micro-modem: An acoustic communications and navigation system for multiple platforms. In *Proccedings of IEEE/MTS Oceans Conference*, volume 1, pages 1086–1092, September 2005.

[8] Darren K. Maczka, Aditya S. Gadre, and Daniel J. Stilwell. Implementation of a cooperative navigation algorithm on a platoon of autonomous underwater vehicles. *Oceans 2007*, pages 1–6, 2007.

[9] Anastasios I. Mourikis and Stergios I. Roumeliotis. Performance analysis of multirobot cooperative localization. *IEEE Transactions on Robotics*, 22(4):666–681, 2006.

[10] Esha D. Nerurkar, Stergios I. Roumeliotis, and Agostino Martinelli. Distributed MAP estimation algorithm for cooperative localization. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA'09)*, pages 1402–1409, May 2009.

[11] Stergios I. Roumeliotis and George A. Bekey. Collective localization: A distributed kalman filter approach to localization of groups of mobile robots. *IEEE International Conference on Robotics and Automation*, pages 2958–2965, April 2000.

[12] Nikolas Trawny, Stergios I. Roumeliotis, and Georgios B. Giannakis. Cooperative multi-robot localization under communication constraints. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA'09)*, pages 4394–4400, May 2009.

[13] Jerome Vaganay, John J. Leonard, Joseph A. Curcio, and J. Scott Willcox. Experimental validation of the moving long base line navigation concept. In *Autonomous Underwater Vehicles, 2004 IEEE/OES*, pages 59–65, June 2004.

[14] Sarah E. Webster, Ryan M. Eustice, Christopher Murphy, Hanumant Singh, and Louis L. Whitcomb. Toward a platform-independent acoustic communications and navigation system for underwater vehicles. In *Proceedings of the IEEE/MTS OCEANS Conference and Exhibition*, Biloxi, MS, 2009. Accepted, To Appear.

[15] Sarah E. Webster, Ryan M. Eustice, Hanumant Singh, and Louis L. Whitcomb. Preliminary deep water results in single-beacon one-way-travel-time acoustic navigation for underwater vehicles. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009. Accepted, To Appear.