Marine Mammal Marker (MAMMARK) Dead Reckoning Sensor for In-Situ Environmental Monitoring

Gabriel Hugh Elkaim*, Eric B. Decker*, Guy Oliver†, and Brent Wright‡
*Autonomous Systems Lab, Computer Engineering, University of California, Santa Cruz, Santa Cruz, CA 95064 USA
†Institute for Marine Science, University of California, Santa Cruz, Santa Cruz, CA 95064 USA
‡Boulder Creek, CA 95006 USA

***Abstract*— Understanding the behavior of marine mammals is quite limited by the observation technology used. Surface tracking using either geolocation or Argos satellite tags have shown that these mammals range much farther than previously thought. Relatively simple time/depth recorders (TDR's) have shown dives to depths of over 1000 meters for durations of over one hour. To further the understanding of these aquatic creatures, a smaller and more capable tag is being developed that can be deployed for longer durations and with more sensing capabilities. This tag utilizes a sensor suite consisting of temperature, depth, speed, salinity, three axes of magnetic field, three axes of acceleration, and GPS.**

The three-axis magnetometers and three-axis accelerometers are used to reconstruct the full attitude quaternion of the creature. Fusing this attitude measurement with water speed, and both initial and final position estimates from GPS, a full three dimensional underwater trajectory can be reconstructed (distributing the error from the return surface as an estimate of the ocean currents). This paper looks at three types of dead reckoning filters used to process this data via simulation: (1) pure open loop integration, (2) "scaled" integration that feeds back measured depth, and (3) a Kalman filter used to estimate the position of the creature, with a correction term from the measured depth (pressure). Comparison of these three navigation filters on simulated data for a 20 minute dive shows that the Kalman filter works best, in terms of position drift, though at the cost of high frequency noise in the trajectory. At the end of a 2 Km path, it shows a total offset error of approximately two meters (or 0.1% drift). The scaled filter is the worst, demonstrating instability at certain points, and the open loop integration yields a position estimate that is off approximately 5 meters at the end of the dive.

I. Introduction

This paper descibes in detail the development of a low-cost, small, and capable dead reckoning sensor used to acquire data on diving sea mammals (pictured in Fig. 1). This work is a direct extension of the previous work found in [9], and as such, will have largely the same background, hardware, and software descriptions. These sections are included in this paper for completeness, and the work diverges in the section on Navigation Filtering (Section IV) where the work of this paper is formally treated and discussed.

Marine mammals are inherently difficult to study. The cetaceans (whales, dolphins, and porpoises) are totally aquatic and even the amphibious pinnipeds (seals and sea lions) spend most of their lives at sea. Biologists can only catch a glimpse of them as they surface and so have turned to technological solutions to study these animals at sea. The most extensively studied is the northern elephant seal, *Mirounga angustirostris*, with much of this work performed at the elephant seal rookery at Año Nuevo State Reserve, 65 kilometers north of Monterey.

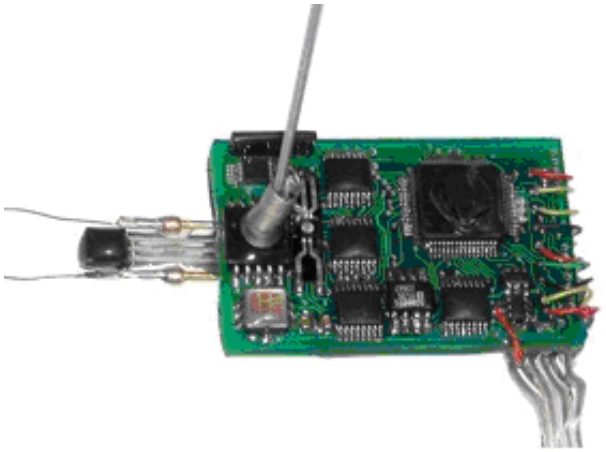

Fig. 1. The Prototype Pinniped Sensor Tag. This tag is designed to be attached to the skin of marine mammals and record their velocity, orientation, and depth, along with the environmental salinity and temperature.

In the early 1980s time-depth-recorders (TDRs), which record changes in water pressure over time, were first attached to elephant seals. Instruments were deployed on seals by gluing them to the seal's pelage just prior to their departure on a foraging trip and were recovered 2.5 - 8 months later when the seals returned to the rookery. The initial results revealed dives that were incredibly long, phenomenally deep and continuous 24 hours a day, day after day, week after week [2]. Mean dive duration of adult females was 22.1 minutes followed by a surface interval of 2.3 minutes [3]. One female in a 10 hour period made 10 dives, 7 of which exceeded an hour, with the longest lasting 97 minutes, and each of these dives was followed by a surface interval of 3 minutes or less. Modal dive depths ranged from 350 to 600 m with a maximum depth exceeding 1600 m.

By adding a photocell to the TDRs, locations could be calculated by determining the day length, which revealed latitude, and the offset of the times of sunrise and sunset

from the place where they were originally tagged, which revealed longitude [15]. This method is accurate to within approximately 100 km.

For elephant seals, this system of geolocation was adequate to describe their long-range movements throughout the north-eastern Pacific. It showed that they undertake two complete foraging migrations each year [7] and that the sexes segregate on their foraging migrations and employ different foraging strategies [4]. Adult males forage off the continental shelf, especially along the Aleutian Islands, and pursue benthic fish, rays, skates, and cephalopods. Females move well offshore and into the pelagic zone where they forage in the upper 1000 m of the water column. Their daily pattern of diving - deep in the day and shallower at night - tracks the diurnal vertical migration of the community of organisms known as the deep scattering layer upon which they feed.

Improvements in the ability to track the seal's movements occurred in the mid 1990s with the development of transmitters, which could be detected by the polar orbiting Service Argos/NOAA satellites. When a satellite was above the horizon and a seal equipped with an Argos transmitter surfaced, an uplink occurred and location could be calculated. The more uplinks which occurred in a single surfacing, the greater the accuracy of the location, and since the number of uplinks per fix is known, a location quality (LQ) could be determined. Because elephant seals are underwater about 90% of the time they are at sea, most seals only had one to four "good" locations a day and over 90% of these were Argos LQ 0, A, or B. These range from an accuracy of 9 km $\pm$ 16 km for LQ 0 hits to 48 km $\pm$ 71 km for LQ B hits [4], which is a considerable improvement over geolocation. Additional advantages of the Argos tags is that they can be tracked in near-real time and approximate locations of mortality can be determined if the transmitter stops transmitting and the seal is never seen again or if the transmitter appears to be moving as if on a ship and the seal is never seen again.

Additional sensors added to the TDRs in the 1990s included thermisters, velocity meters, hydrophones, video cameras and heart rate monitors. Suddenly biologists were data rich as the number of instrumented elephant seals soared past 200 (a sub-adult bull with one of these modern tags attached is pictured in Fig. 2). The range of insights into the biology of these seals was fascinating. The data revealed that the seals, while diving, were employing a variety of behavioral [1] and physiological [22] 'tricks' enabling them to have a lower metabolic rate than when sleeping on the beach!

But biologists are constantly impatient for technological advances to occur, and they can construct the most Rube Goldburg contraptions in their attempts to learn more about their animals. Two areas where improvements were sought were in the accuracy and frequency of surface locations and the ability to record the 3-dimensional movements of the seal between its surfacing locations. Several MAP tags that married a GPS receiver to a TDR, a velocity meter, and a 3-dimensional digital compass were constructed. It was conceptually successful [17], but was a 14-pound behemoth

Fig. 2. A sub-adult bull Elephant Seal with a modern data recorder affixed to its head. This is a modern time depth recorder (TDR) which transmits data back through the Argos satellite. This tag is approximately 4 times as large as the prototype MAMMARK tag.

requiring 12 d-cell batteries to power it and was only deployed on one translocated seal [18], which never returned to Año Nuevo.

Tags were getting increasingly larger and more expensive. Only the best-funded researchers could afford to deploy the newer tags and even they were limited in how many they could afford to deploy because of costs ranging from several thousands to tens of thousands of dollars. Clearly, there was a need for a newly designed tag with the capabilities of the MAP tag while shrinking its size and cost. It will be mounted on top of the seal's head so that when the seal surfaces, the GPS antenna will rapidly shed water and have maximum exposure to the sky. The tag needs to contain at least two external environment sensors, temperature and salinity, which will allow identification of water masses. The electronics will need to be potted to protect them from salt water and to allow them to withstand up to 3000 psi of pressure. To minimize disturbance and work for the seal, the cross-sectional area of the new tag should be less than 5% of the cross-sectional area of a seal's head. To be affordable and deployable in large numbers it needs to be marketed for $500-700. The new MAMMARK tag fulfills these requirements.

The new work in this paper is the development and testing of a dead reckoning filter to reconstruct the trajectory of

the diving pinniped. We simulate two distinct trajectories (a spiraling dive down and a long traversing dive) and use these trajectories to derive the outputs of the sensors on the MAMMARK tag. These sensor outputs are then corrupted with noise and used as the inputs to the dead reckoning filter which reconstructs both the attitude and position of the animal. Certain assumptions are made about the nature of the diving animal: it moves only in the direction that it is pointed, and that the currents are ignored. The current drift will be mitigated by the GPS measurements at both ends of the dive, and based on the very large drag associated with moving through the water, it is unlikely that the pinnepeds spend much time swimming at large angles of attack or sideslip.

The paper is organized as follows: Section I presents the background and motivations for developing the MAMMARK tag, Section II details the physical design of the prototype, Section III explains the software structure used to run the tag, with special emphasis on power conservation. The section on Navigation Filtering, Section IV, is itself separated into five subsections. These are, Subsection IV-A, which details the attitude estimation algorithm (based on two vector observations and a quaternion based solution to Wahba's problem [21]). Subsection IV-B, which addresses calibration issues for the sensors that comprise the MAMMARK tag. Subsection IV-C, which goes into the mathematical construction of the dead reckoning filter. Subsection IV-D, which elaborates on the Kalman filter implementation in order to feed back the measured (and noisy) depth back into the position domain in the optimal way. And finally, Subsection IV-F presents the simulation data and trajectory reconstruction. Lastly, we present conclusions in Section V and outline the future work in Section VI.

II. MAMMARK HARDWARE

Fig. 1 shows the physical prototype MAMMARK tag, and Fig. 3 shows a block diagram of the major hardware subsystems that make up the MAMMARK marine tag.

The core of the hardware is a TI MSP430 ultra-low power microcontroller [20]. The microcontroller includes several on-chip peripherals: SPI controllers, serial communication, clock control, watchdog mechanisms, DMA controllers, timers, and a small amount of program flash memory and several KB of RAM space.

As all of the sensors are analog in nature, the main interface between the microcontroller and the sensors is the analog-to-digital converter (ADC) subsystem. While the MSP430 has an onboard ADC capable of converting analog signals with 12 bits of precision, it was felt that this was not sufficiently precise to achieve the desired performance. Thus, an external 16 bit ADC is attached via one of the SPI channels. In order to sample all of the various sensors, this single ADC is multiplexed between each of the sensors; that is, one sensor at a time is converted, the value stored in RAM and then the next sensor converted.

There are two different kinds of sensors attached to the central microcontroller: differential and single-ended. Additionally, each of the sensors requires a different amplification to maximize the sensitivity of the sensor. Thus, two different operational amplifiers (op-amps) are used, one differential and one single-ended for each of the corresponding type of sensors. The gains for each of these op-amps are under microprocessor control, and thus each sensor can be optimized separately for maximum dynamic range in order to get the best reading from the signal of interest.

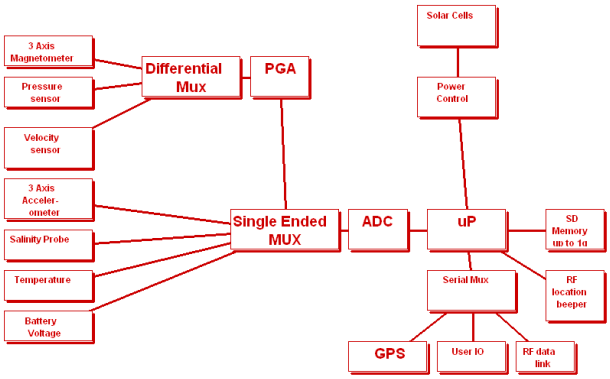

Fig. 3. System Block Diagram for Tag Hardware. The system block diagram shows the subsystem structure of the prototype MAMMARK hardware. The system has been designed to be small and very low power in order to record data in-situ for long duration on *very* limited battery power.

The differential sensors include the magnetometers (three axes), depth (pressure) transducer, and water velocity (measured using a two-axis strain gauge). The single-ended sensors include the accelerometers (three axes), salinity, temperature (both of the water and of the microcontroller), and battery voltage. These sensors are used to reconstruct the three dimensional trajectory (both in position and velocity domains), as well as salinity and temperature profiles.

In order to collect enough data for useful analysis, these sensors must be periodically sampled, filtered, corrected for calibration parameters, and stored for post-processing. Depending on the rate at which we are sampling each of the sensors, the amount of data collected can become very large (currently, the prototype limits the maximum sampling rate to 20Hz for all sensors). This is, however, an arbitrary limit imposed by the software. If, after experimentation, it is found that higher data rates are required, this can be changed without modifying the hardware.

Additionally, the most important aspect of the MAMMARK tag is the low-power nature of the system. In order to be of use in the field, the MAMMARK must collect and store the sensor data until the animal in question returns to a place where the tag can be recovered (provisions are in place for remote data retrieval, however, the limited bandwidth of the RF link would make this a very slow process). Thus, the main function of the software is to manage the power consumption of the hardware such that data can be collected for very long duration. It is expected that data collection for up to a year will be accomplished with a single lithium-ion cell. This will be discussed in more depth in the software section.

Clearly, the MAMMARK cannot store all of the data in RAM, as the very limited storage capability available on the microcontroller would not be sufficient for more than a few minutes of data. Furthermore, volatile memory storage runs the risk of data loss on failure. Instead, the main storage is provided by flash memory (either a Secure Digital or Multi-Media Card) connected to the SPI bus. This sub-system provides secure long term storage of up to 4 gigabytes. As the sensors provide data, it is aggregated in RAM until a certain block size is attained, at which point it is written to the flash memory subsystem. It is estimated that this storage will be sufficient for a significant period of time even at high sensing rates (even sampling continuously at 20Hz on all sensors the MAMMARK has over 20 days of storage capacity).

As previously stated, one of the key performance criteria for the device is long life and given its battery powered nature, power conservation is critical. Most sub-systems are kept in a low or powered-off state whenever possible. The power draw of each subsystem is balanced against the required time for power-up and stabilization for high quality sensor readings. The hardware includes power circuits that allow each individual sensor to be powered or de-powered as overall system requirements necessitate. In terms of the power budget, the sensor components are some of the most power hungry, and thus great care is taken when sequencing the power up in order to minimize overall power consumption.

Other subsystems include external communications, a Global Positioning System (GPS) module, RF Beeper, and the previously described mass storage module. Each module has provisions to be individually powered to minimize overall consumption. The hardware also includes provisions for recharging the battery via an external solar cell, which would of course only be active when the animal is at the surface.

Locating the tag after the animal returns to the beach and possibly molts is the function of the RF beeper. The RF beeper sends out a signal that allows the location to be determined using directional antennas on the receivers, and possibly will include a simple encoding of the latitude and longitude from GPS. Note that this will only be powered when the tag detects that it is on surface (this is relatively easy using both the depth gauge and salinity sensors). Once the tag is located, communications with the tag can occur either via a directly connected cradle or via radio on the 900 MHz ISM Band. The RF communication is done using an off-the-shelf RF communications solution (the Radiotronix Wi.232DTS) [19]. Within the tag and base communication module, provisions are made in the design to enable communications with numerous tags at speeds up to 115Kbps on multiple channels.

Given the large amount of data that will need to be transferred, and the power required to transfer that data over the RF link, it is definitely considered a backup option to physically recovering the tag itself. The foreseeable scenarios for RF only communication are, for instance, when the tag is attached to a large male on the beach who is too aggressive to approach (tranquilizers for the large bulls are sometimes lethal, and all efforts will be made to leave the animals unharmed). Smaller animals can be restrained using physical means that allow the tag to be removed.

III. MAMMARK SOFTWARE

The software running on the TI MSP430 is responsible for managing the communication, sensor, storage, and power subsystems. In order to make the MAMMARK useful, the software must carefully manage the sequencing and power consumption. Fig. 4 shows a simplified block diagram of the MAMMARK software system.

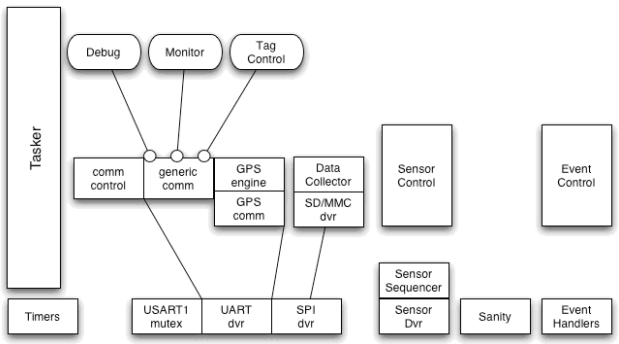

Fig. 4. The Block Diagram of System Software. The software to run the MAMMARK tag is divided into several subsystems and implemented as a hierarchical state machine in order to produce robust performance over long duration.

The system is broken into two main parts, Events and Tasks. Events are handled by event handlers; these are essentially messages passed between the handlers and can result in a task being awoken. Tasks themselves are implemented as simple one-shot, run-to-completion threads. This results in a structure where events cause event handlers to execute; the event handlers are assumed to be atomic (that is, they cannot be interrupted). They can execute in either an interrupt or non-interrupt context, depending on the specific event. An example of an event is the "data available event" generated by the sensor system. This event causes the "Data Collector" task to execute and the new data collected. When enough data is collected (this status itself is an event), this will trigger another task which will write the accumulated data to the flash data storage card via the SD/MMC driver.

The core of the MAMMARK tag (and its reason for existing) is the sensor system. The software that manages the sensors performs several tasks: scheduling, sequencing, reading, etc. Lists of sensors are maintained which determine the rate that each sensor should be sampled. These lists are analyzed by the "Sensor Control" task and a current operating sensor sequence is determined. The "Sensor Sequencer" in combination with the "Sensor Driver" implements data collection. This is done in a general way such that new sensors can be added as they become available.

Sensor drivers are responsible for the actual interaction with each physical sensor, including the application of power, any pre-conditioning required, and the associated timing. The sensor sequencer is driven primarily using time events provided

by the Timer subsystem. This system also provides events as needed to the "Tasker," itself responsible for overall thread sequencing.

Event Control provides for rendezvous between event producers and consumers. Entities interested in receiving events inform Event Control. When an event is signaled, Event Control determines what entities are interested and delivers the event. Event Control, in cooperation and coordination with the Tasker, is responsible for the mechanism where tasks waiting for an event are put to sleep and run when the corresponding event triggers them. The Tasker provides an implementation of a simple one-shot tasking system, which is consistent with the run-to-completion paradigm used within the software for the MAMMARK tag.

As previously noted in the hardware section, mass storage is provided by a SD/MMC card providing up to 4 Gigabytes of storage. The size of memory available is increasing and cost of these cards is decreasing, making them a "future-proof" technology. It also contributes to the low cost of the overall system. Data is generated via the sensor subsystem, collected by the "Data Collector" into blocks as required by the mass storage system and then written via the SD driver.

The lowest level of hardware implementing the data path to the SD card is shared with the serial communications hardware. This complicates the software as several hardware subsystems communicate through the same bus. Access to this hardware is controlled via the USART1 mutex module; mechanisms implemented via this module exist for a driver to request the hardware, wait until it is made available (if busy), and then proceed with its assigned task. As such access to both the serial (UART) hardware and mass storage (SPI) is quantized and this is reflected in the design of other subsystems interfacing through this hardware.

The GPS subsystem is one such serial user. When the GPS subsystem is active, the GPS communication module collects data packets from the GPS and hands these to the GPS engine. The GPS engine implements the actual state machine that generates appropriate GPS events for use by the rest of the system. This includes an "ignore GPS" state while the tag is submerged.

The other major user of the serial hardware is "Generic Communications." This module, coupled with "Communication Control," provides a generic packetized multi-port serial interface to the external world. Provisions are made for both local (via a physical cradle) and remote communications (via the Wi.232 RF module).

The Debug task is used for monitoring and controlling internal state of the tag, especially while testing and developing the tag. The Monitor task is used for monitoring normal operation of the tag, and collecting data to determine any malfunction for later analysis and repair.

The Tag Control task is used for controlling exactly what kinds of data the tag is collecting as well as uploading collected data to a host. It interacts with Sensor Control to establish sensor sequences. Note that some of these power-up sequences are not obvious, and require complicated staging in order to minimize the power and maximize the performance of the sensors.

The last major piece of the system is the "Sanity" module; this module is responsible for monitoring the health of the system and forcing a restart in case of problems. As the system may operate unattended for up to a year at a time, the intent is to increase the likelihood of valid data being written to storage, even in the case of some component failure. The Sanity module includes but is not limited to the hardware watchdog and oscillator monitors, which monitor low level tag hardware function. If a failure restart is needed, the Sanity module is also responsible for marking any data structures to allow for detection of the event and subsequent resynchronization of the data stream.

IV. Navigation Filtering

The MAMMARK tag differs from previous Time Depth Recorders (TDRs) both in the low-power and low-cost aspect, but more importantly in its ability to reconstruct the three dimensional trajectory of the tagged animal. This section explores the mathematical methods used for the navigation filtering on the MAMMARK tag. Note that these algorithms need not run in real time, as the raw sensor measurements are recorded, and thus the attitude and position reconstruction occurs once the tag has been recovered.

The MAMMARK tag is able to reconstruct the pinniped attitude (three dimensional orientation) by using the body-fixed measurements of Earth's magnetic field, along with the body-fixed measurements of gravity. By combining these two measurements, and with their known values in the navigation frame, sufficient information is available to determine the tag attitude precisely (within the limits of sensor noise). This is known as Wahba's problem, and the particular method used in this work to solve it is elaborated next, in Subsection IV-A.

Once the attitude is determined, the body measured velocity (with respect to the water) is rotated into the navigation frame and then integrated once to give position. This is a classical dead reckoning filter. The details of which are given in Subsection IV-C. Of course, integrating a noisy sensor measurement results in a linear error growth, and further the errors in attitude result in an error in projecting the velocity from the body frame to the navigation frame.

In order to get the best estimate of attitude possible, the magnetometers and accelerometers are calibrated using a two-step non-linear calibration algorithm that does not require any external truth reference (other than knowledge of the total magnitude of the magnetic and gravitation fields). The calibration techniques are detailed in Subsection IV-B.

Lastly, a direct measurement of depth (via the pressure sensor) is available. In this work, we use this measurement as a feedback signal in order to improve the estimate of position. This is done using a conventional Kalman filter. The formulation of the Kalman filter, as well as the justification for both the process and measurement noise matrices are detailed in Subsection IV-D.

Lastly, two trajectories for a diving seal are simulated in order to validate the various filtering models as well as the attitude algorithms. True errors are available as in these cases the true trajectory is known. The details of the simulations are detailed in Subsection IV-F.

A. Attitude Estimation

Attitude estimation is the process of determining the three dimensional orientation of an object based on some set of *noisy* measurements such as to minimize the estimation error. Within attitude estimation are two related issues: the first is the parametrization of attitude, and the second is the actual estimation filter. There are several different parameterizations of attitude, each with its own advantages and disadvantages. Briefly, the most common are the Direction Cosine Matrix (DCM), Euler Angles, Rodrigues Parameters, and Quaternions. Note that an excellent treatment of each of these appears in [16].

The Direction Cosine Matrix is the rotation matrix that rotates a vector (or transforms it) from one coordinate frame to the other. It is a 3×3 matrix that is orthonormal (that is, its transpose and its inverse are the same). It is called the direction cosine matrix because $R_{ij} = cos(\alpha_{ij})$. While there are 9 numbers making up the DCM, only three are independent, the other 6 are constrained by the orthonomality of the matrix.

Euler angles are the classic aircraft angles, usually described as yaw, pitch, and roll (the $[3-1-2]$ Euler angle set), though these are certainly not the only set of possible Euler angles. The advantage of Euler angles is that only three parameters are required, however, the construction of the rotation matrix requires the evaluation of transcendental functions, which can be costly in computational terms depending on the platform. Another advantage of the Euler angles is that they are immediately intuitive to a human observer, whereas none of the other parameterizations are. However, a disadvantage of Euler angles is the singularity that occurs when pitch, θ, reaches 90° and the remaining angles become poorly defined. This is the so-called "gimbal lock" problem, which is much less of an issue with strapdown attitude systems than it was with mechanical gimbals.

The Rodrigues parameters (also known as the Gibbs vector), is defined as $p = atan(\phi/2)$, where a is the unit vector about which the rotation occurs, and ϕ is the angle of rotation. Note that the Rodrigues parameters also have a mathematical singularity (which is, in fact, common to all three-parameter attitude parameterizations, see [5]).

Lastly, Quaternions are a four parameter representation of attitude, and thus do not suffer from the mathematical singularity. Quaternions represent rotation as a hyper complex number, that is the three dimensional analog of a complex number. The Quaternion is defined as:

$$\mathbf{q} = \left[\begin{array}{c|c} q_0 & \vec{q} \end{array} \right] \tag{1}$$

where

$$q_0 = cos(\phi/2) \tag{2}$$

and

$$\vec{q} = asin(\phi/2) \tag{3}$$

Note that all of the transformations and conversions with Quaternions are accomplished with simple multiplies and additions.

In this work, we specifically solve for attitude in the Quaternion domain, by solving Wahba's problem for the measurements of Earth's magnetic and gravitational fields. The basic formulation of attitude determination from two or more non-collinear, non-zero vectors measured in the body frame and known in the navigation frame is referred to as Wahba's problem, based on her 1966 formulation [21]. While there exist a variety of solutions for Wahba's probelm (see for instance, [6]), our solution is based on a linearization of the attitude estimation problem in the navigation frame (as opposed to the body frame). The great advantage of this formulation is that it results in a linear time invariant measurement equation, easily solvable by either an iterated least squares or Kalman filter solution.

While the complete treatment of this attitude estimation algorithm is beyond the scope of this work, see [8], [11], [12]. It is important to note that in this application, the two vector quantities being measured are Earth's magnetic and gravitational fields. Note, the on-board accelerometers do not, however, measure only the gravitational acceleration. Rather, the accelerometers measure the quantity $\vec{g} - \vec{a}$, the specific force on the accelerometers. We assume that large pinnipeds do not maneuver at a large fraction of $\vec{g}$, and thus the error in the accelerometer measurement is in fact very small.

The basic steps of the attitude filter are as follows (again, see [8], [11], [12] for details):

1) Initialize the attitude quaternion estimate to:

$$\hat{q} = \left[\begin{array}{cccc} 1 & 0 & 0 & 0 \end{array} \right]^T \tag{4}$$

and the error quaternion to

$$q_e = \left[\begin{array}{cccc} 1 & 0 & 0 & 0 \end{array} \right]^T \tag{5}$$

2) Use the attitude estimate, $\hat{q}$, to map the measured magnetic field to the navigation frame. That is:

$$\hat{h}^n = \hat{q} \otimes h^b \otimes \hat{q}^* \tag{6}$$

where $\otimes$ is quaternion multiplication and $\hat{q}^*$ is the quaternion complement of $\hat{q}$.

3) Map the accelerometer measurements to the navigation frame as above.
4) Formulate the errors in the navigation from by subtracting the estimated values from the known ones:

$$\delta\hat{h}^n = \vec{h}^n - \hat{h}^n \tag{7}$$

and

$$\delta\hat{a}^n = \vec{a}^n - \hat{a}^n \tag{8}$$

5) Formulate the measurement matrix directly from the body measured quantities:

$$H = \begin{bmatrix} -2\left[\vec{h}^b\times\right] \\ -2\left[\vec{a}^b\times\right] \end{bmatrix} \tag{9}$$

6) Solve the Normal equations (i.e. by using the Moore-Penrose pseudo-inverse):

$$H^\dagger = \left[H^T H\right]^{-1} H^T \tag{10}$$

and use it to compute the new error quaternion.

7) The new error quaternion is computed from:

$$q_e = \alpha H^\dagger \begin{bmatrix} \delta\hat{h}^n \\ \delta\hat{a}^n \end{bmatrix} \tag{11}$$

where α is a smoothing parameter chosen between 0 and 1.

8) Update the quaternion estimate using quaternion multiplication:

$$\hat{q}\,(+) = \hat{q}(-) \otimes q_e \tag{12}$$

note that this is rotating $\hat{q}$ by the small correction q_e.

9) Repeat from step 2 until converged. Also note that the previous value for $\hat{q}$ can be used to initialize the algoritm the next time through.

This algorithm has been extensively tested using Monte-Carlo simulations and has shown to be quite stable and robust to wide band noise on the measured inputs. Fig. 5 shows the attitude errors from the simulated trajectories detailed in Subsection IV-F. Note that these errors are such that the standard deviation $(1-\sigma)$ of the error in heading is approximately 1.6° and the corresponding errors in pitch and roll are 0.8° and 0.4°, respectively.

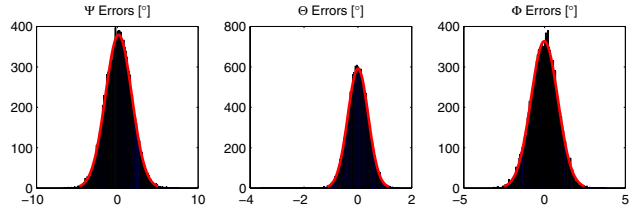

Fig. 5. Histogram of Simulated Attitude Errors (in degrees). Simulation of the attitude algorithm running on a diving seal shows excellent agreement with the true attitude. The pitch and roll histograms show error standard deviations of less than 1° and the yaw shows less than 2°.

B. Sensor Calibration

As shown above, in Fig. 5, measurement errors in the three axis magnetometers and accelerometers directly affect the attitude estimate. As will be discussed in Subsection (IV-C, *Dead Reckoning*, attitude errors corrupt the position solution as well.

Given that the accelerometer and magnetometer measurements make up an integral part of the trajectory reconstruction algorithms, care must be taken in calibration of these sensors to minimize the errors that are accumulated during the Dead Reckoning integration process.

In this work, we use a technique for calibrating the accelerometers and magnetometers directly from rotation using a two-step algorithm described in [10], [13]. The algorithm stems from the observation that when rotating the perfect sensor around through all angles, then the measurements plotted would trace out a circle for a 2D sensor, and a sphere for a 3D sensor.

The basic measurement equation for a given axis on a sensor is:

$$h_{meas} = \frac{1}{sf} h_{true} + b + \omega_n \tag{13}$$

where h_{meas} is the measured output, h_{true} is the true measurement, b is the bias or offset, and ω_n is the wide band noise on the sensor. In the two axis case, when the body fixed sensor is rotated around a circle, its components should be such that when plotted they have a center point of (0,0) and a radius of the value of the magnitude of the magnetic field. Bias errors will cause the circle to be shifted off of the origin, and scale factor errors will distort the circle into an ellipse.

The algorithm for calibration is a non-linear two step algorithm which first does a least squares estimation of a set of parameters which are then manipulated algebraically to extract the scale factor and bias errors. Note than non-orthogonality of the measurement axes will cause a distortion to the ellipse as well, and this can also be accounted for. The center of the ellipse is the bias error for both axes, and the semi-major and semi-minor axes of the ellipse are the scale factors.

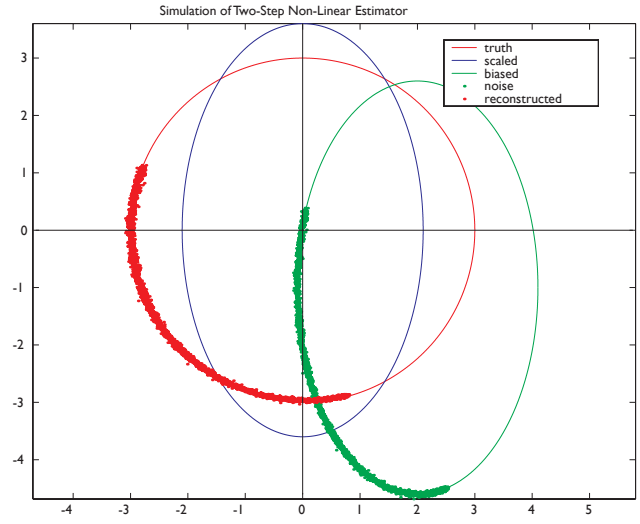

Fig. 6. Two-step Estimation Algorithm on Simulated Data. The red solid line is the true data, the blue solid line is the scaled (scale factor) data, and the green solid line includes the bias offsets. The green dots are the sampled points that are then used in the algorithm to reconstruct the red dots (true sampled points).

The same analysis is true for the 3D case, except that an ellipsoid is generated, rather than an ellipse. It the case of the 3D sensor, care must be take to traverse enough of the surface

of the ellipsoid to have good observability of the parameters. Again, note that the only thing required for the algorithm to work is a knowledge of the true magnitude of the magnetic field, and motion of the sensor.

Fig. 6 shows the 2D version of the algorithm working on simulated data. Here, no attempt is made to model non-orthogonal axes, and a significant portion of the circle is traversed in order to ensure good parameter estimation. The red solid line is the true data, the blue solid line is the scaled (scale factor) data, and the green solid line includes the bias offsets. The green dots are the sampled points that are then used in the algorithm to reconstruct the red dots (true sampled points).

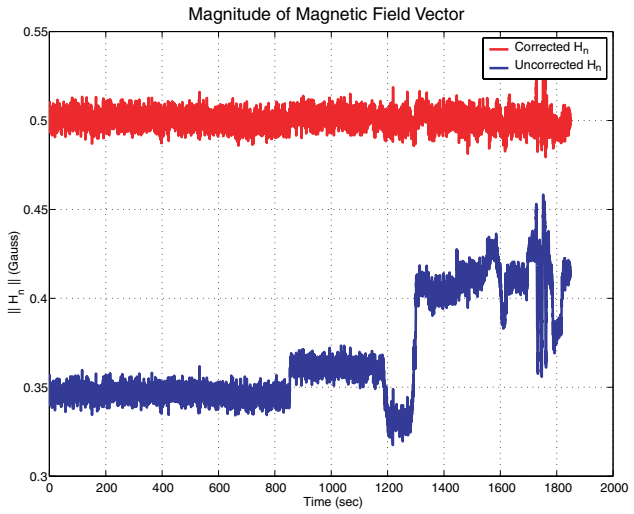

Fig. 7. Reconstruction of the Magnetic Field from Experimental Measurements. Before and after calibration measurements of the total magnetic field strength, as measured using a Honeywell HMC2300 sensor (similar to what is being deployed on the MAMMARK tag).

Fig. 7 shows the reconstruction of the magnitude of the magnetic field before and after the calibration algorithm. Note that this is on actual experimental data, and clearly the effects of the errors can be seen. This data is from a Honeywell HMC2300 integrated 3D magnetic field sensor, which uses the same sensing technology as the MAMMARK tag. It is not, however, taken from the tag itself. More details can be found in [8].

The other sensors are calibrated in a more conventional manner, taking data in a flume for the Velocity sensor, and using a calibrated pressure chamber for the depth sensor.

C. Dead Reckoning

The essence of the trajectory reconstruction is a dead reckoning filter that is used to keep track of the pinniped position. Dead Reckoning is a type of navigation scheme that propagates the position with a growing error over time. Thus, without periodic updates from some external measurement source, the solution will eventually grow without bound.

Dead Reckoning originates from the sailing world, and was basically the process by which the ships kept track of their positions. Speed was determined usually by timing a marker in the water (often an actual log) as it went by the ship, and plotting that velocity along the current heading to determine the new position. Of course, this could not account for surface currents, and without a celestial fix or a citing to a known landmark to reset the position, the errors could grow quite large.

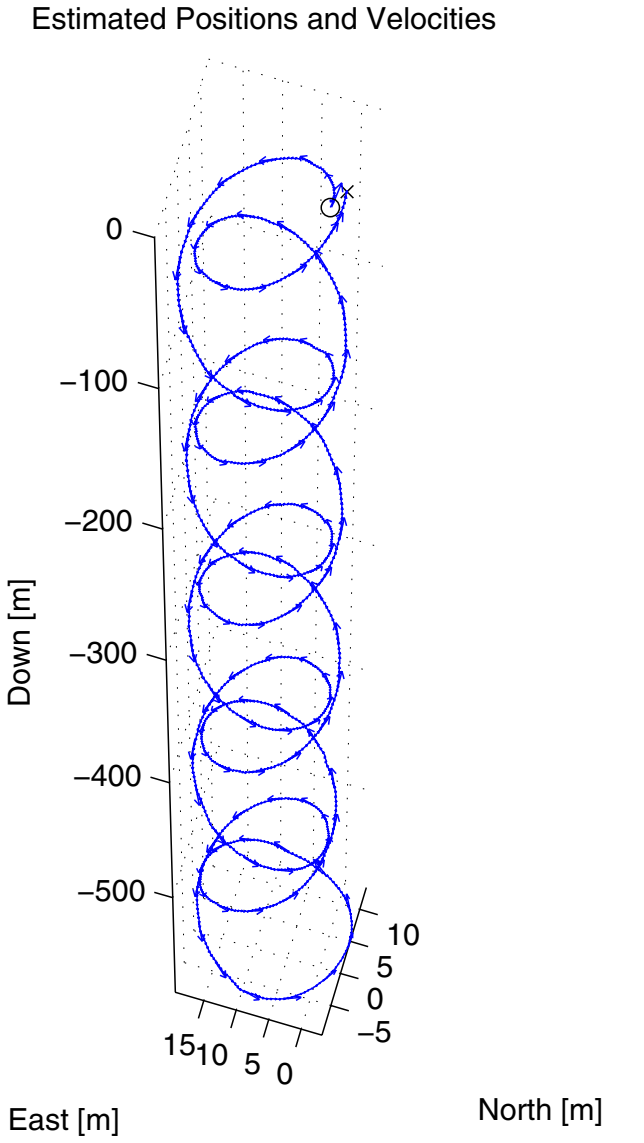

Fig. 8. Reconstruction of the spiral dive trajectory of a seal. The trajectory of a simulated diving seal is reconstructed from the measurements, suitably corrupted with wide band noise, using a Dead Reckoning Kalman filter. Position errors at the end of the run are within 4 meters after a path length of over a kilometer.

In the case of the MAMMARK tag, we use dead reckoning to propagate the position of the pinniped from its last known position (as determined by GPS when at the surface) to the time when the animal surfaces again. Having previously described the methodology for determining attitude from the sensor measurements, the body measured water velocity is

rotated into the navigation frame:

$$V^n = \hat{q} \otimes \begin{bmatrix} V_x^b \\ V_y^b \\ 0 \end{bmatrix} \otimes \hat{q}^* \tag{14}$$

There is no component of body measured velocity in the z-direction, as the sensor is a 2 axis device. Due to high water density and large induced drag from the low aspect ratio of the pinnipeds, it is expected that the animals will swim with little or no angle of attack.

The coordinate frame for integrating the velocity into position is a simple North-East-Down coordinate frame, with the origin placed at the last known GPS position. In this coordinate frame, the equations are linear, and can be expressed as:

$$\dot{x} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} x + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \vec{V}^n \tag{15}$$

where x is the vector $[N, E, D]^T$. The navigation frame velocity is the rate of change the navigation frame positions. Implemented in a discrete time system, these equations change to:

$$x_{k+1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} x_k + \begin{bmatrix} \Delta t_s & 0 & 0 \\ 0 & \Delta t_s & 0 \\ 0 & 0 & \Delta t_s \end{bmatrix} \vec{V}^n \tag{16}$$

where Δt_s is the time step between measurement updates in the attitude and velocity. What can be seen from Eq. 14 and Eq. 16 is that attitude errors corrupt the input to the integration (even if there were no noise on the velocity sensor) and cause the position error to grow.

D. Kalman Filter

In order to bound the position error growth from the dead reckoning filter, periodic position fixes are required. This, however, can be problematic as the pinnipeds dive very deep and surface for only a very brief amount of time (experimental results will demonstrate whether or not our GPS has time to lock on and track before the animal dives again).

We do, however, have a (noisy) measurement of depth that can be used to limit the error not just in z, but it can also be used to limit the error growth in the horizontal dimensions as well. This can be done both informally, and formally. The informal method is referred to as the "scaled" filter, and is the result of comparing the dead reckoning depth with the measured pressure depth. The position update (in all three axes) is scaled to match the measured depth.

$$x_{k+1} = x_k + \left(\frac{z_{pres} - z_k}{\Delta t_s V_z^n} \right) \Delta t_s \vec{V}^n \tag{17}$$

It is immediately obvious that this solution becomes numerically unstable when the z-velocity measurement is close to zero.

A more formal approach is to use the Kalman filter to feed back the measured depth into the equations. The measurement equation is very simple in this case:

$$H = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} x \tag{18}$$

Measurement of the depth is directly available. In order to implement a Kalman filter, the process noise and measurement noise must be defined. That is, the relative measures of trust in the process model as related to the measurements. In the case of Eq. 16, it must be noted that the attitude errors come in to the equation where process noise would normally enter. Additionally, an error in attitude causes the velocity vector to be mis-rotated, thus causing cross-correlation between input states (this is all a way of saying the the process noise matrix in not diagonal). Also note, the depth measurement will be very noisy compared to the attitude errors due to the much larger dynamic range required of the depth sensor.

The Kalman filter implementation is the standard one taken from [14], and is a time varying filter (though in practice the gains very quickly converge to steady state). The process noise covariance matrix is converted to its discrete time equivalent, and the time propagation of the state and covariance estimate is given by:

$$\begin{aligned} \hat{x}_{k+1} &= \Phi \hat{x}_k + \Delta T_s \cdot \hat{V}^n \\ P_{k+1} &= \Phi P_k \Phi^T + C_d \end{aligned} \tag{19}$$

where Φ is the identity matrix and C_d is the discrete equivalent process noise matrix. Note that this formulation does not require uniform sampling times for the sensors, and can easily handle variable length update steps. When a new measurement of depth is available, the update part of the Kalman filter is performed:

$$\begin{aligned} L &= PH^T \left[HPH^T + R_v \right]^{-1} \\ \hat{x}_+ &= \hat{x}_- + L \left(z_{pres} - \hat{x}_z \right) \\ P_+ &= \left[\Phi - LH \right] P_- \end{aligned} \tag{20}$$

Where R_v is the sensor noise variance and L is the resulting estimator gain. By varying the ratios of C_d to R_v, more emphasis can be given to the depth measurement or less as required. Again, due to the injection of the attitude errors into the process noise, the off diagonal elements of C_d are non-zero, meaning that some of the error in the z-axis is used to alter to x and y axes.

E. Simulation Results

In order to evaluate the different schemes for propagating the attitude and position of the MAMMARK tag for trajectory reconstruction, a simulation was set up in MATLAB in order to validate the different schemes. Two different dive trajectories were modelled: (1) a spiraling dive to 600 meters depth and spiraling ascent back to the surface, and (2) a long straight dive to 600 meters and back up with three 90° turns. Both simulations ran for a total dive length of 20 minutes.

The simulation begins at a known position in LLA coordinates and first matches the velocity and attitude to the trajectory, such that the noise free navigation frame values are computed. Using the noise free attitude previously computed, the true values of Earth's magnetic field and gravity are computed in the body frame, and then corrupted with wide-band noise at the appropriate levels (5 mg's for the accelerometers and 1 mGauss for the magnetometers). Additionally, the

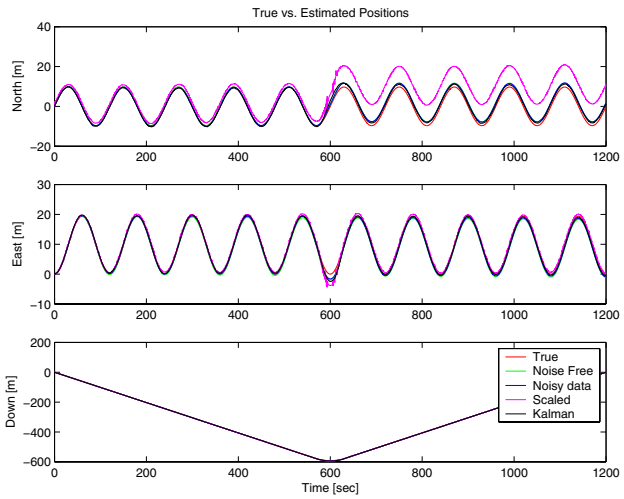

Fig. 9. Simulated diving trajectory and reconstruction vs. time. The red solid line is the true trajectory, the green line is the noise free sensor reconstruction, the blue is a straight dead reckoning filter, the magenta is the scaled solution, and the black line is the Kalman filtered data.

velocity sensor is also corrupted with wide-band noise (0.05 m/sec) and the depth measurement is corrupted (1.5 meters of noise). Lastly, the coordinates are all converted from LLA to NED with the origin at the starting point.

F. Simulation Results

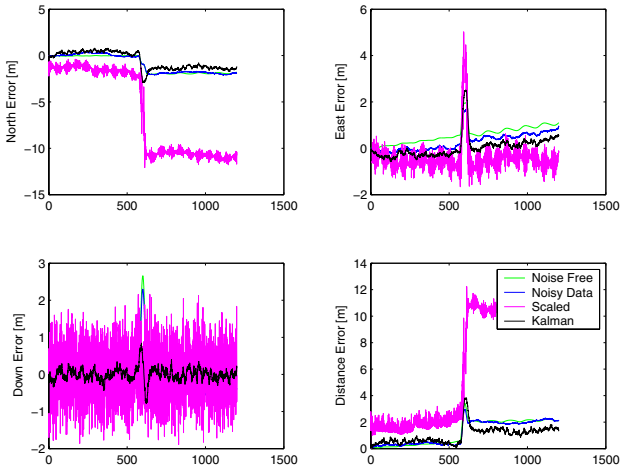

Fig. 10. Simulated spiral diving trajectory reconstruction errors vs. time. This shows errors in the North, East, and Down, as well as total errors in distance.

These corrupted measurements are fed to the various algorithms: Dead Reckoning, Scaled, and Kalman filter. In addition, a noise free case is also computed (which shows what would result with perfect sensors). Note that the noise free case still results in errors due to the fact that the velocity sensor is only 2D. The performance of the attitude estimation has already been shown in Fig. 5, and has been shown to have normally distributed errors of less than 1° in pitch and roll and approximately twice that in yaw ($1\text{-}\sigma < 2^\circ$).

Method	Run	North [%]	East [%]	Down [%]	RSS [%]
DR open loop	Spiral	0.13	0.15	0.1	0.13
	Square	0.18	0.1	0.03	0.29
DR "scaled"	Spiral	0.6	0.1	0.1	0.6
	Square	1	0.26	0.1	0.9
DR Kalman	Spiral	0.1	0.2	0.01	0.06
	Square	0.16	0.13	0.01	0.26

TABLE I
TRAJECTORY RECONSTRUCTION ERRORS

The overall trajectory reconstruction on the simulated data is excellent, with the worse method (scaled) showing errors of approximately 10 meters accumulated after a distance of over 1.5 km and 20 minutes of integration. The open loop dead reckoning filter shows an accumulated error of approximately 2 meters, and the Kalman filter approximately 1.5 meters. While the overall performance of the Kalman filter is the best, it achieves this result by integrating the very noise depth measurement into the solution, and thus trades off smoothness for accuracy. Fig. 8 shows the reconstruction of the spiral trajectory using the Kalman filter.

The specific performance on the other trajectory is different, but the trends are the same. The results are summarized in Table I, and the reconstruction of the second trajectory is pictured in Fig. 11. Note that the scaled method has problems at the bottom of the trajectory where the z-velocity changes sign and thus goes through 0, even though the scaled solution was ignored at that point.

V. CONCLUSION

Current understanding of the life cycle of the pinnipeds is limited by a lack of knowledge due to limited observation of the animals in the wild. In this work, we have detailed the progress on the prototype marine mammal marking tag, the MAMMARK. The main features of the tag are low-cost, long duration, and large storage capability. The MAMMARK is capable of operating at depths of 2000 meters, and to survive a two ton seal smashing it upon the rocks. Low-cost is achieved by using commercial off-the-shelf technology, utilizing low-cost MEMs sensors developed for the automotive market. Low-power is achieved by using a modern low-power microcontroller, and using it to power cycle most of the sensor and communication subsystems in order to last over a year on a small battery pack.

This ruthless power management, however, complicates the software and hardware structure as several of the sensors must be allowed to stabilize before the ADC can convert the values. Additionally, certain sensors, such as the magnetometers, require pre-sampling conditioning in order to eradicate permanent bias errors (in the case of the magnetometers, a set/reset pulse must be performed to demagnetize the sensing element). A simple, yet robust, software structure has been designed to maximize the longevity of the sensor, while at the same time giving good sensing performance.

Based on the sensor data, pre- and post-calibration may be possible. The magnetometers and the accelerometers are

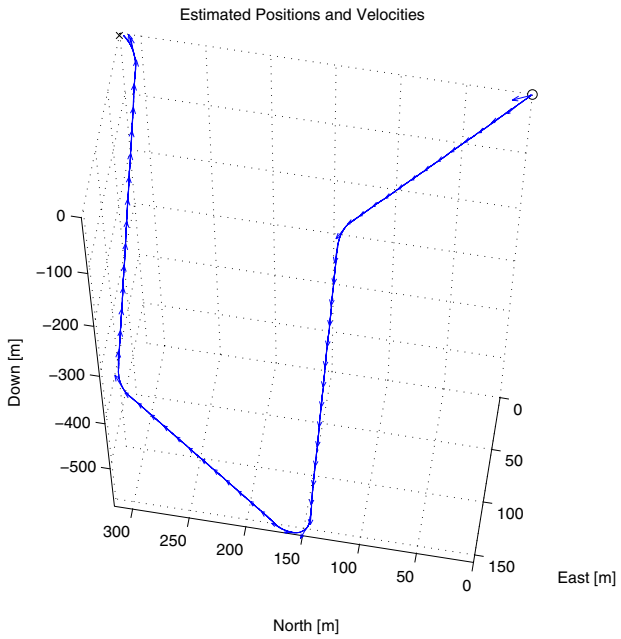

Fig. 11. Reconstruction of the square dive trajectory of a seal. The trajectory of a simulated diving seal is reconstructed from the measurements, suitably corrupted with wide band noise, using a Dead Reckoning Kalman filter. Position errors at the end of the run are within 4 meters after a path length of over a kilometer.

calibrated using a two-step process that requires only the motion of the sensor and the knowledge of the magnitude of the total gravitation (9.81 m/s^2) and the magnetic field (0.52 Gauss). Pre-calibration is accomplished by rotating the sensor through a diverse set of angles; post calibration will be accomplished if the animal has moved through a diverse enough set of angles. Other sensors such as depth and velocity are pre-calibrated in a conventional manner.

Simulation studies of two trajectories for a diving seal were performed, and three dead reckoning filters were tested to check robustness and performance. The simple dead reckoning open loop integration performed very well, and was made much worse by attempts to scale its output to correspond to noisy depth measurements. The dead reckoning Kalman filter performed best of all of both trajectories, showing most improvement in the depth measurement, and sacrificing smoothness for improved position estimates. All filters showed an error drift that ranged from 0.1% to 1% of the distance traversed, and should be more than sufficient for oceanographic study.

The prototype MAMMARK tag has been built, and currently testing and calibration is underway. The MAMMARK tag, when deployed, will represent a significant step forward for in-situ sensing capabilities for marine biologists. In addition to the large storage capacity, high precision sensing, and long-life, the MAMMARK tag will have a significantly lower cost than traditional Time/Depth Recorders (TDR's) enabling large scale deployments.

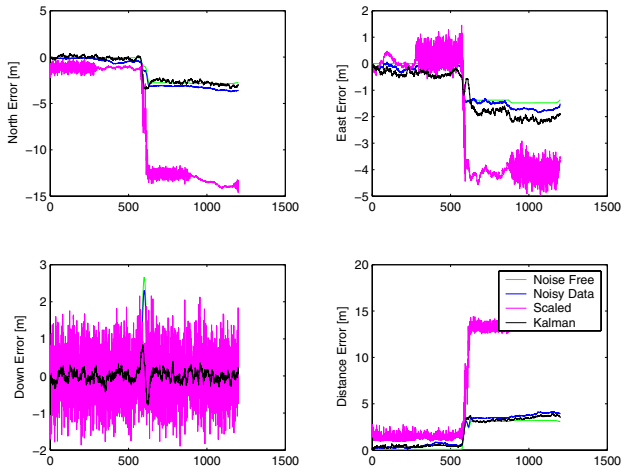

Fig. 12. Simulated square diving trajectory reconstruction errors vs. time. This shows errors in the North, East, and Down, as well as total errors in distance.

VI. Future Work

Current work is focused on the development of the prototype MAMMARK tag, and as such, leaves significant detail for future work. The current state of the tag is as a laboratory prototype. The major blocks of work for the future of this tag are:

1) Finish and test tag software
2) Calibrate sensors
3) Lab testing of finished tag
4) Field deployment in controlled environment
5) Field deployment for returning juvenile pinniped pups
6) Full deployment

References

[1] R. D. Andrews, D. R. Jones, J. D. Williams, G. W. Oliver, P. H. Thorson, D. P. Costa, and B. J. Le Boeuf. Heart rates of northern elephant seals diving at sea and resting on the beach. *Journal of Experimental Biology*, 200:2083–2095, 1997.

[2] B. J. Le Boeuf, D. P. Costa, A. C. Huntley, and S. D. Feldkamp. Diving behavior of female northern elephant seals, mirounga angustirostris. *Canadian Journal of Zoology*, 66:446–458, 1988.

[3] B. J. Le Boeuf, D. E. Crocker, S. P. Blackwell, P. A. Morris, and P. H. Thorson. *Marine mammals: advances in behavioral and population biology*. Oxford University Press, London, UK, 1993.

[4] B. J. Le Boeuf, D. E. Crocker, D. P. Costa, S. P. Blackwell, P. M. Webb, and D. S. Houser. Foraging ecology of northern elephant seals. *Ecological Monographs*, 70(3):353–382, 2000.

[5] J. L. Crassidis and F. L. Markley. Unscented filtering for spacecraft attitude estimation. *AIAA Journal of Guidance, Control, and Dynamics*, 26(4):536–542, 2003.

[6] G. Creamer. Spacecraft Attitude Determination Using Gyros and Quaternion Measurements. *The Journal of Astronautical Sciences*, 44(3):357 – 371, July - September 1996.

[7] R. L. DeLong, B. S. Stewart, and R. D. Hill. Documenting migrations of northern elephant seals using daylength. *Marine Mammal Science*, 8:155–159, 1992.

[8] G. H. Elkaim. *System Identification for Precision Control of a WingSailed GPS-Guided Catamaran*. PhD thesis, Stanford University, Stanford, CA, 2001.

[9] G. H. Elkaim, E. B. Decker, G. Oliver, and B. Wright. Development of a marine mammal marker (mammark) for in-situ environmental monitoring. In *Proceedings of the ION National Technical Meeting, ION-NTM 2006*, Monterey, CA, Jan 2006. ION.

[10] D. Gebre-Egziabher and G. H. Elkaim. Calibration of strapdown magnetometers in magnetic field domain. *ASCE Journal of Aerospace Engineering*, 19(2):1–16, 2006.

[11] D. Gebre-Egziabher and G. H. Elkaim. Gyro-Free Attitude Determination by Vector Matching. *AIAA Aerospace Electronic Systems Journal*, submitted for publication.

[12] D. Gebre-Egziabher, G. H. Elkaim, J. D. Powell, and B. W. Parkinson. A Gyro-Free, Quaternion Based Attitude Determination System Suitable for Implementation Using Low-Cost Sensors. In *Proceedings of the IEEE Position Location and Navigation Symposium, PLANS 2000*, pages 185 – 192. IEEE, 2000.

[13] D. Gebre-Egziabher, G. H. Elkaim, J. D. Powell, and B. W. Parkinson. A non-linear, two-step estimation algorithm for calibrating solid-state strapdown magnetometers. In *8th International St. Petersburg Conference on Navigation Systems, St. Petersburg, Russia*. IEEE/AIAA, 2001.

[14] A. Gelb. *Applied Optimal Estimation*. MIT Press, Cambridge, MA, 1974.

[15] R. D. Hill. *Elephant seals: population ecology, behavior and physiology*. University of California Press, Berkeley, CA, 1994.

[16] Jack B. Kuipers. *Quaternions and Rotation Sequences*. Princeton University Press, Princeton, New Jersey, 1999.

[17] G. W. Oliver. The MAP Tag: visualizing the tracking and diving behavior of marine mammals. In *Proceedings of the 11th Biennial Conference on Biology of Marine Mammals*, Orlando, FL, 1995.

[18] G. W. Oliver, P. A. Morris, P. H. Thorson, and B. J. Le Boeuf. Homing behavior of juvenile northern elephant seals. *Marine Mammal Science*, 14(2).

[19] Radiotronix. Wi.232dts user manual. http://www.radiotronix.com/products/proddb.asp?ProdID=1, 2005.

[20] Texas Instruments (TI). Msp430x161x mixed signal microcontroller data sheet. http://focus.ti.com/docs/prod/folders/print/msp430f1611.html, 2005.

[21] Grace Wahba. Problem 65-1 (Solution). *SIAM, Review*, 8:384 – 386, 1966.

[22] T. M. Williams, R. W. Davis, L. A. Fuiman, J. Francis, , B. J. Le Boeuf, M. Horning, J. Calambokidis, and D. A. Croll. Sink or swim:strategies for cost efficient diving in marine mammals. *Science*, 288:133–136, 2000.