SENSOR SELECTION AND BEHAVIOR STRATEGIES FOR OBSTACLE AVOIDANCE IN SMALL-DIAMETER AUVS

E. Gil Jones
Bluefin Robotics Corp.
Cambridge, Mass.
gil@bluefinrobotics.com

Nathaniel Fairfield
Bluefin Robotics Corp.
Cambridge, Mass.
than@bluefinrobotics.com

Abstract

Avoiding obstacles is an essential element of most small-diameter autonomous underwater vehicle missions, though specific obstacle avoidance requirements for different missions vary widely. Obstacle avoidance is one of the most important safety considerations for many missions, and designing an optimal avoidance system for a given mission deserves careful consideration. This paper attempts to give an overview of design options for obstacle avoidance systems. Obstacle avoidance sensor characteristics are examined, including processing techniques for generating obstacle data. The paper then examines obstacle avoidance algorithms and behavior system design, the systems that use the obstacle data to avoid obstacles within the context of a larger mission. Finally, recommendations are made for optimal obstacle avoidance systems for several important AUV missions.

1 Introduction

The domain of Autonomous Underwater Vehicles (AUVs) requires that the vehicles be capable of safely and efficiently coping with the full complexity of the underwater environment. As there is no human involved in the control of an AUV, the AUV must contain all sensors and algorithms for processing and acting on its sensor data autonomously. Thus the AUV must both be capable of accurately sensing the environment around it and intelligently reacting to that information.

The problem of obstacle avoidance in AUVs has been an active research topic recently [6] [1] [10] [5] [13], especially as mission requirements have become more extensive, requiring increased capabilities in sensors and control algorithms. Previously, sensing obstacles was especially difficult on small-diameter AUVs - defined as AUVs less than 1m in diameter - which have hard limitations on space, power, and processing capabilities. Older sensors were low resolution, low definition, or slow to scan the target area, drastically reducing their effectiveness for the reactivity required for obstacle avoidance. In the last several years, newer high definition sensors have become available that are high rate and fit the space and power budget of a small-diameter AUV.

This paper gives a holistic view of optimal obstacle avoidance system design on AUVs. The first section examines a host of sensors capable of obstacle detection. The next section focuses on high-level processing of sonar data for succinct obstacle definition. The paper then considers the use of the obstacle data in both high- and low-level obstacle avoidance. The final section groupsxs sensor selection, data processing techniques, and avoidance algorithms into optimal obstacle avoidance systems for several important AUV missions.

2 Obstacle Avoidance Sensors (OAS)

Acoustic sensor systems are well suited to the underwater environment. Although sound travels slowly in water (about 1500m/s), high levels of acoustic energy can be transmitted through water with relatively minimal attenuation [8]. This gives acoustic systems much greater ranges than optical systems, which can see perhaps 15m in clear water and which are useless in turbid water. Radar and other high-frequency radio waves are attenuated within a few feet. Thus, this discussion of long-range OAS will focus entirely on sound-based detection systems.

Active sonar systems operate by emitting some sort of "ping", a short burst of acoustic energy, and then recording and analyzing the echos which result when the ping reflects off of objects in the water. There are a number of sonar systems which expand on this basic principle. This section gives a brief overview of these systems and their operating principles, especially as they impact the type of obstacle avoidance data that the sonar returns.

2.1 Echosounder

An acoustic transducer is the basic element of transmission and detection of sound waves in water. The essential property of the transducer is that it can convert the mechanical energy of the acoustic compression wave into an electrical signal, andxs vice versa. The transducer may be constructed such that it does not transmit/receive sound energy from all directions with equal sensitivity. The relative sensitivity of a transducer as a function of angle is known as a beam pattern (Figure 3(d)). This directional sensitivity is often intentionally cultivated during the manufacturing process in order to form a concentrated narrow beam. In fact, the simplest OAS is a single narrow-beamed transducer pointed out in front of the vehicle [17]. If the OAS just returns the range of the primary echo, it is known as an echosounder (Figure 3(a)). An OAS that returns a plot of the intensity of the return as a function of time is known as a scanning sonar, as the echo 'scans' across a sector of the bottom (Figure 3(c)). This simple OAS demonstrates several principles which hold true for all OASs: 1) A narrower beam provides higher angular resolution, and 2) The refresh rate is inversely proportional to the range.

Resolution of an OAS is described in terms of range resolution and angular resolution. The range resolution is the minimum change in the distance of an obstacle that the OAS can detect. In general, sonar systems have a range resolution of 0.1% of the maximum range. Angular resolution refers to the smallest angle separating two obstacles such that the OAS can resolve them as separate obstacles. In terms of the beam pattern of the sonar, the edge of a beam is usually considered the 3-dB down, or half-power point [8] (Figure 3(d)). The resolution of the size of an obstacle depends on the angular resolution of the sonar, as well as the distance to the obstacle. There is a similar trade-off between range and refresh rate, since the ping must travel out to the maximum and range and back for each refresh cycle.

2.2 Multibeam

Since the purpose of an OAS is to alert the vehicle to potential obstacles, the range value returned by a single transducer amounts to a minimum of information. For example, the vehicle cannot tell how wide or high the obstacle is, whether it is perpendicular to the vehicle's direction of travel, or whether it is slanted. While it is possible to scan a region by mechanically aiming a single transducer, the time required to make a full sweep would keep the vehicle from taking timely evasive action. For this reason, most OASs use several narrow beams to sweep the area ahead.

The simplest improvement over a single transducer is to add more transducers, each aimed in a slightly different direction, which can then be fired in sequence to map out the area covered by their beams. This is known as a multi-preformed beam sonar because no beamforming is required. As most transducers have a beam width of around 10 degrees, multi-preformed sonars have low resolution, and usually have only a few beams. Either each echosounder can fire in turn, which lends to rather slow refresh-rates, or a separate projector transducer can ensonify the entire region, and each receive transducer record the echo along its particular aimed beam [9].

In order to achieve higher resolution, the signals from these seperate elements must be combined using a technique known as beamforming. The same beamforming techniques can be used both for transmission and reception. High-resolution beamforming sonars can achieve beam widths of 0.25 degrees, compared to the 5-10 degree beams achieved by a multi-preformed beam sonar [9]. The three commonly used techniques to form narrow beams are analog beamforming, digital beamforming and acoustic lenses (Figures 3(e) and 3(f)).

Both analog and digital beamforming work by combining the signals from multiple transducers with incremental time-delays or phase shifts, such that the reception from the steered direction sums together coherently. Several commercial OASs use analog beamforming, but the future seems to be oriented toward digital beamforming because of its superior precision and flexibility. In beamforming, the desired beam is formed by time-delaying the signals from each transducer as the echo sweeps across the array [3]. This is useful when the array size is large relative to the wavelength. For example, if the desired beam is directly perpendicular to a linear array, the echo will arrive at all elements of the array at the same time (time-delay of zero). If the the desired beam is at a 45 degree angle to the array, the echo will sweep across the transducers at a given speed, and the data from each element must be delayed accordingly. Alternatively and equivalently, beams can be formed using phase shift by converting to the frequency domain using an FFT. Time-delay processing in the frequency domain corresponds to

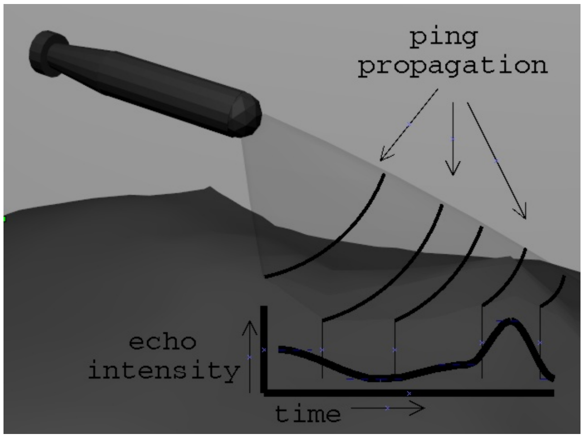

Figure 1: As the transmitted ping sweeps across the bottom, the scanning sonar records the intensit of the return over time, and converts the values to a row of an image, see Figure 2.

phase-shifts in the time domain.

If the hydrophones are arranged in a linear array they can be operated as echosounders, which ping and then report the time of the main reflection. This yields a simple list of ranges along the axis of array. More information can be extracted by recording the intensity of the entire return as it scans across the bottom (Figure 1). This is just an extension of the scanning principle mentioned earlier, but when there are sufficient beams available the quality of the data returned is sufficiently good to reconstruct a 2D image (Figures 2 and 3(i)). In order to scan a sector of the bottom with high resolution, the horizontal linear array of transducers have narrow horizontal beam patterns but wide vertical beam patterns so that they can receive the return as it sweeps away in range. A sonar capable of forming high resolution imagery is called an imaging sonar. The DIDSON imaging sonar has 96 beams, which it can ping in sequenced groups. It uses an acoustic lens to form beams in parallel, thus resulting in a relatively high refresh rate of 5-20Hz, depending on range [4]. Scanning sonars such as the DIDSON present their data as an image, so the sample bin size determines the resolution of that image. Since the range and position data of obstacles must be extracted from this image, the bin size actually determines the range resolution of the sonar, as well.

A 2-dimensional array of transducers, either planar or conformal, using 3-dimensional beamforming, can provide a full two dimensional array of ranges (Figure 3(b)). Since each of the beams is very narrow both vertically and horizontally, like a

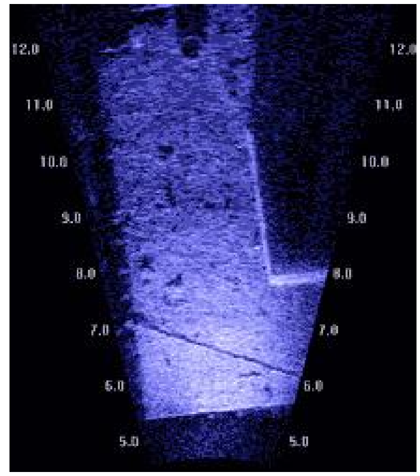

Figure 2: A sample image from the DIDSON imaging sonar, showing a bridge abutment at about 5 meters [4].

pencil beam, the scanning principle can't really be applied, and so 3-D multibeam sonars are inevitably echosounders, eg [16]. If the resolution of the 3-D multibeam is sufficiently high to yield details about the area in front of the OAS, it can also be considered an imaging sonar.

3 Obstacle Avoidance Sonar Data Interpretation

Most of the above sonars contain substantial onboard signal processing associated with recording the returned pings and filtering to obtain useful data. The could vary from a simple range in the case of an echosounder to the complex and human-interpretable images created by the high-resolution sonars. This data must then be processed into a form that can be passed to an obstacle avoidance algorithm. This section will discuss the process of moving from the native data format of the OAS to a highly compact data form that has identified obstacle threats. Some processors designed especially for obstacle avoidance may have the capabilities to do this processing, or it may be the task of the sonar integrator to perform the processing. Many of the techniques discussed below require substantial processing power for successful real-time operation, a limited resource on a small-diameter AUV. Processing requirements should be considered when considering the viability of a given sensor.

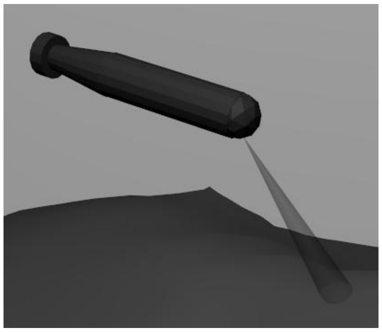

(a) Echosounder scene

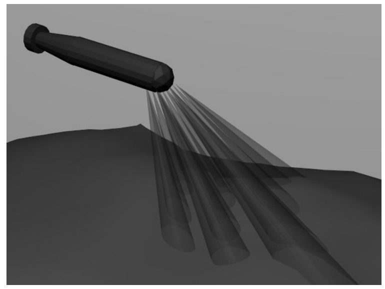

(b) Multibeam sonar scene

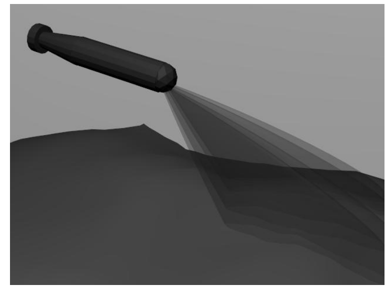

(c) Scanning sonar scene

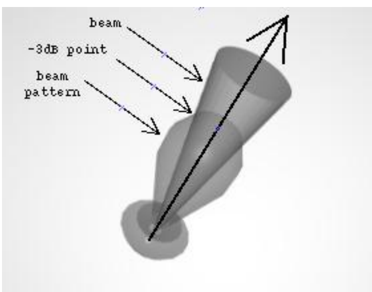

(d) Echosounder beampattern

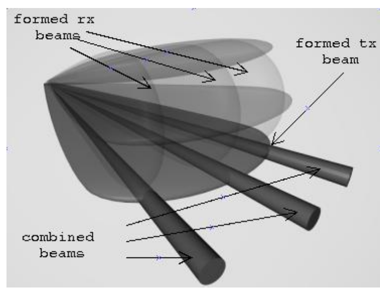

(e) Multibeam sonar beampattern

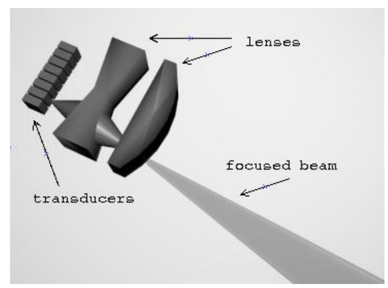

(f) Scanning sonar with acoustic lense beampattern

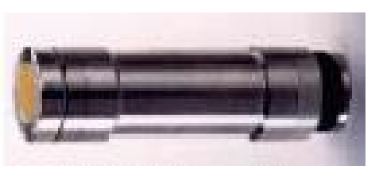

(g) Tritech PA500

Range: 50m
Freq: 500kHz
Beamwidth: 6° conical

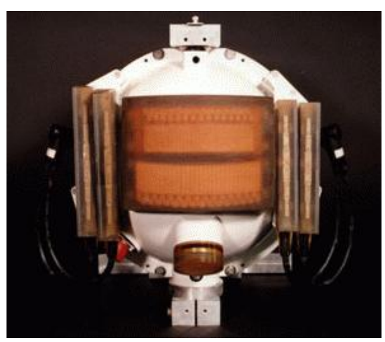

(h) Sonatech STA013

Range: 180m
Freq: 200/230KHz
Beamwidth: 11°H by 3°V

(i) DIDSON

Range: 23.25m
Freq: 0.7MHz
Beamwidth: 0.8°H by 12°V

Figure 3: Summary of several commercial sonars

3.1 Echosounders and Plane Filtration

As discussed above, simple echosounders will provide a range to the nearest object in the ensonified region. As there is no available corroborating information for the single echosounder, any time the sensor returns several continuous readings of approximately the same range, it should be assumed a real obstacle is being detected.

The mounting angle for echosounders is a substantial issue. See section 5.1 for a more detailed discussion of the trade-offs associated with pointing angle of echosounders. Whatever the mounting angle, the vehicle may assume an attitude that leaves the ocean floor or surface within the ensonified area of the sonar. In these cases, the sensor may return a valid sounding that originates not from an obstacle but from the plane of the ocean floor or surface. This sounding is useful information for topology-following, but should be classified as such and not as obstacle data. The AUV can use a simple algorithm that relies on attitude, altimetry, and bathymetry data obtained from other sensors to filter out echosounder range returns that probably indicate the surface or floor. The algorithm should be symmetric for either surface or floor. The following discussion will refer only to the floor.

Figure 4 shows the computation. The computation should give an approximate range value R that the sensor would report if the return came from a flat bottom. If this value is greater than the range of the sensor, then any valid return from the sensor be assumed as return from an obstacle. If R is within the range of the sensor, then a comparison to the real range should yield some useful information. If the obstacle avoidance algorithm is attempting to do topology following, the expected range, the returned range, and attitude and altimetry information can be used to compute a slope angle of the terrain. Other applications might be more interested in an obstacle/non-obstacle determination. For this purpose, a scaling coefficient should be set such that if the reported range is less than the predicted range multiplied by the scaling coefficient an obstacle is reported. Setting a coefficient near 1 will result in a very sensitive obstacle reporting module that reports any elevation change as an obstacle. Lower coefficients will reduce sensitivity, increasing the size of bottom elevation change that will be ignored but also allowing the vehicle to get closer to an obstacle before avoidance is initiated.

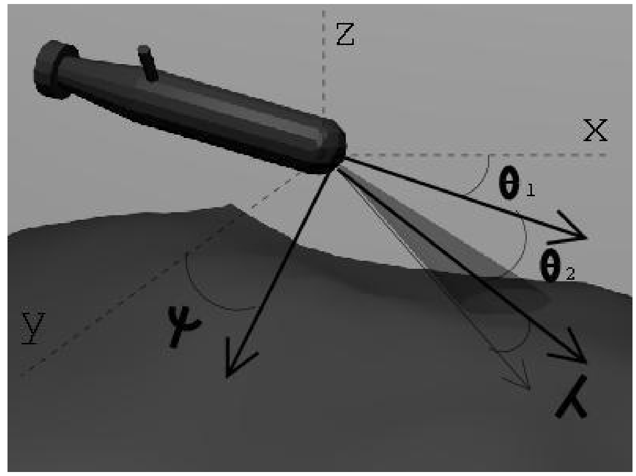

Figure 4: Range prediction computation over a flat bottom, where λ is radius of the beam of the echosounger, θ^1 is the pitch of the vehicle, θ^2 the pointing angle of the echosounder, and ψ the roll of the vehicle. The estimated range is given by: $\texttt{Range(m)} = \frac{\texttt{Alt(m)}}{\cos(\theta^1+\theta^2-\lambda)\cos(\psi)}$

3.2 Linear Arrays and Image Processing

Linear arrays return a 2-D array of information. If the array is horizontally mounted it will directly sense range and horizontal extent; if vertically mounted it will sense range and vertical extent. As many linear arrays are used for imaging, an imaging sonar will generally return an MxN pixelated data swath, where M is the number of beams of the sonar, and N is the number of range returns from each beam. Each cell will have an intensity associated with it - these intensities can be interpreted as color for the purposes of image creation. Initially, the data is filtered for noise; multibeam data is often extremely noisy, with backscatter and reverberation from the ocean floor and surface as the primary sources of noise [13] [15]. A threshold is then applied to the data to determine which pixels are judged occupied. The activated pixels are then grouped into distinct obstacles. The range to these obstacles can be set to the range to the nearest occupied cell of those that makeup the obstacle. If the linear array is horizontally mounted, then horizontal extent can be determined by the width of the cells with high intensity. A 2-D centroid can be taken that locates the object precisely in XY [13] [7] [12].

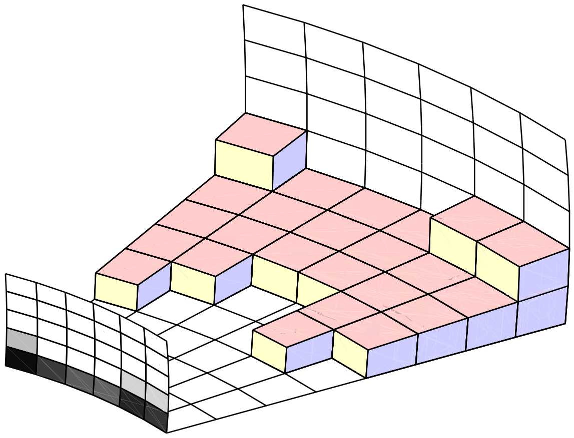

Figure 5: Representation of obstacle occupancy in a 3-D grid as generated from a 2-D multibeam sonar

3.3 3-D Data Obstacle Extraction

A 2-D transducer array returns actual 3-D data, since it creates beams in both the horizontal and vertical planes (Figures 3(b) and 3(e)). Thus, the vertical extent of an obstacle is directly sensed. The data from a 2-D array can be represented as a three-dimensional grid with dimensions MxNxR, where M and N represent the number of horizontal and vertical beams, respectively, and R is the number of ranges for each beams as determined by the sampling interval.

For a given ping, each of the MxNxR cells should have an associated intensity. This value can be normalized, and then compared to a detection threshold to determine a binary occupancy for each cell. The algorithm in [14] then travels down each beam, defined as the R returns associated with a given MxN position. It begins with the closest cell, continuing to further cells until one or more occupied cells are reached. The algorithm identifies all contiguous cells past the first several occupied cells as a occupied 'tube' - all subsequent occupied cells past a several non-occupied cell gap are discarded. Tubes are grouped together to form inter-beam obstacles. A centroid is taken to get the 3-D position of each obstacle, and information as to range, horizontal extent, and vertical extent are all available. Figure 5 shows a processed occupancy frame, with one large obstacle in the back of the ensonfied region.

3.4 Obstacle Tracking and Mapping

Both 2-D and 3-D processing yield an XY and possibly a Z location in space for each obstacle. While this information is useful for not hitting obstacles in the short term, building a map of the space or accurately determining characteristics of the obstacles requires another step: the obstacles must be 'tracked' from ping to ping. Tracking means that newer information is used to refine the position and characteristic measurements of the obstacles from previous ping frames. Some tracking techniques, prominently optical flow, attempt to interpolate pixel intensities from frame to frame to relate features [7]. Unfortunately, the noise of multibeam data is often too high to make this approach viable. Another approach uses the processed data on position and a Kalman filter to update positions based on new information. The Kalman filter attempts to identify whether the obstacles in the new ping cycle are consistent with previous measurements and come from an already identified obstacle, whether the measurements come from a new and as yet unidentified obstacle, or whether new data may call into question already- identified obstacles. The resulting data should be more accurate than using data from a single ping cycle. See [13] for a detailed implementation.

Tracked XYZ obstacle locations and estimations of obstacle characteristics constitute a map - this map can be used for higher level planning in situations where reactive obstacle avoidance is not sufficient to complete a given task. Having knowledge of the vehicle's own perceptions of its movement as acquired from LBL, inertial sensors, and a Doppler Velocity Log becomes essential for map generation - the less accurately the vehicle knows its own position, the less accurate estimates of obstacle XYZ will be. In general, position estimate drift rates will be low enough that the map remains fairly accurate for short-term path-planning.

4 Obstacle Avoidance Behavior Design

The previous sections have discussed the basics of obstacle avoidance sonar and the processes involved in getting obstacle data from those sensors. This section discusses the basic building blocks of the algorithms that actually take the obstacle data and use it to generate the motor commands that will cause the vehicle to move through the space.

4.1 AUV Dynamics

4.1.1 Rear-propelled vehicles

AUVs fall into two major classifications in terms of motion dynamics. The vast majority of

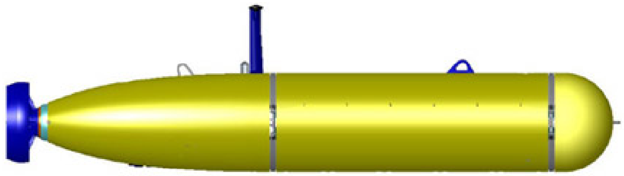

Figure 6: Bluefin21 Rear-propelled AUV

small-diameter AUVs fall into the first category (Bluefin21, MARTIN, RAIS, REMUS) - rear-propelled vehicles (Figure 6). A rear-propelled vehicle has actuators that all line up along the forward axis of the vehicle. Lateral or vehicle motion is caused by physical actuation of the thruster for single-propeller vehicles, or differential thrust in multi-propeller vehicles. Rear-propelled AUVs are often well-controlled only in the forward direction; running the thrusters backward can lead to uncontrolled motion. This configuration constrains the movement of the vehicle in several important ways. A rear-propelled vehicle has a well-defined turning radius, the radius of the circle that the vehicle makes when executing a lateral 360° turn at cruising speed. Vehicle dynamics in the vertical plane are limited by either the elevator deflection or the stall angle of the vehicle, constraining the depth delta the vehicle can achieve over a given horizontal distance.

4.1.2 Hovering vehicles

A much smaller group of vehicles belong to the hovering vehicle dynamic group (Figure 7). Hovering vehicles contain many separate thrusters at a number of different pointing angles. These thrusters should give the ability to hold a position ('hover') in XYZ, even against current, and to translate in any direction in the vertical and horizontal frames. Hovering vehicles have far fewer restrictions on movement than rear-propelled AUVs - the turning radius should be 0, as the hovering vehicle can turn in place, and the depth delta over a given distance should be infinite, as the vehicle can go straight up or straight down. Unfortunately, hovering vehicles are generally slow and not capable of going long distance quickly, and thus are not useful for certain mission specifications. While some of the following concepts are also applicable to control algorithm design on hovering vehicles, the rest of this discussion concerns the algorithms specifically for rear-propelled vehicles.

Figure 7: Lockheed-Martin's CETUS Hovering AUV (MIT AUV lab web page)

4.2 Autonomous Vehicle Obstacle Avoidance Concepts

All AUVs will have control modules in the vertical and horizontal frames. General control requirements include maintaining depth, altitude, or pitch in the vertical frame and following a heading or a trackline in the horizontal frame. These controls will be designed for function in unobstructed water.

The designer must decide how to incorporate obstacle avoidance into the basic control architecture. Obstacle avoidance is of paramount importance to land-based robotics, and substantial work has been done on the arbitration between obstacle avoidance and other vehicle mission goals. In recent years the literature has settled on hybrid systems as a way to accommodate both low-level and high-level goals of a system. A hybrid system is generally implemented in two layers. The lower-level layer is generally called the reactive component. The reactive component of a hybrid system is a stateless module that is generally responsible for implementing small steps in the plan formulated by the planning layer and reacting quickly to incoming sensor data. It's the responsibility of the higher-level, or deliberative, component of the hybrid system to keep a notion of the state of the vehicle as determined by accumulated sensor data, and to use this state to implement the higher-level goals of the system.

There are both reactive and deliberative approaches to obstacle avoidance. Reactive avoidance techniques are designed to ensure the safety of the vehicle and to perform basic maneuvers to negotiate obstacles. Deliberative avoidance can be used to increase the complexity of obstacle that can be successfully navigated.

4.3 Reactive Avoidance

4.3.1 Single Echosounder Avoidance

A single echosounder returns a minimum of information, and can be used only for simple obstacle avoidance. The nature of the information returned is not detailed enough to accumulate into useful state, so any avoidance algorithms based on a single echosounder are going to be primary reactive. For example, an algorithm may just turn the vehicle in a new direction when an obstacle is detected. The sensor doesn't return enough information to be useful in determining what direction to turn, and thus the direction selected to turn the vehicle should just be mission-dependent. Avoidance of floating obstacles like surface ice can be accomplished by commanding descent when an obstacle is encountered, and bottom-feature avoidance by commanding ascent.

Another primary use of a single echosounder can be detecting situations where the vehicle is physically endangered. If an obstacle is detected withing a given range, the vehicle could enter some sort of safety violation maneuver, like killing thrust entirely or doing a tight spiral ascent to the surface.

4.3.2 Basic Avoidance using 2-D and 3-D Data

With the exception of a single echosounder all sensors discussed will provide some measure of two dimensional data - at least range and bearing to obstacles. This enables additional capabilities in obstacle avoidance. The nature of the data allows the vehicle to make decisions about the direction to turn when encountering an obstacle, and even how strongly to react to a given obstacle.

The avoidance algorithm must define the manner in which the obstacle data is used to generate the motor commands that will avoid the obstacle. One simple approach uses a rule-based logic to determine the vehicle's actions. A rule could be "If there's an obstacle on the left and none on the right turn right" or "if there's an obstacle directly in front of the vehicle, ascend". This approach allows the designer to specify behavior exactly using expert knowledge - this approach should yield a high level of determinism, as the vehicle should act predictably. This approach does require the designer to transform the data from a low-level obstacle representation, perhaps ranges and bearings, to high-level determinations like "there's an obstacle on the left". Additionally, the designer can't anticipate all possible scenarios the vehicle may encounter - the approach's predictability also limits it's adaptability to novel situations.

Other approaches use a direct and continuous approach to translate obstacle data into commands. An example is the potential fields approach, which represents the obstacles as repulsive forces acting on the vehicle. An obstacle's location will cause a force to be exerted that will turn the vehicle away from the obstacle. If multiple objects are sensed, the forces are summed. Generally, the nearer an obstacle is to a vehicle the larger the magnitude of force it will exert [1]. It is relatively simple to create a function that creates a motion vector based directly on the ranges and bearings to the obstacles.

A number of factors may affect the magnitude of the action recommended by the chosen approach. For instance, in [6], the turn rate is limited in accordance with the vertical and horizontal swath of the vehicle's obstacle avoidance sensor - turning faster than a certain rate could cause the vehicle to move into territory that had not yet been sensed, and thus could contain an obstacle.

4.3.3 Arbitrating Among Multiple Goals

Obstacle avoidance will never be the sole function of the reactive system. Each goal of a given mission can be represented by a behavior, which will create a motor recommendation based on the goal of the behavior and the incoming data. These recommendations must then be arbitrated to yield the final motor outputs of the reactive system. Two prominent arbitration methods are the schema method and the subsumption method. The schema method takes all recommendations, scales them according to gains, and sums the result, creating an output that is a blend of all goals of the system. A goalpoint-achieving behavior and the obstacle avoidance behavior could be blended, yielding avoidance that should have an awareness of the goalpoint. The gains can be set such that a higher priority has a higher gain and will thus factor more heavily in the final weighted sum.

In contrast, the subsumption approach operates as a winner-take-all approach. Behaviors are organized in a strict hierarchy, and the highest-priority behavior that decides to act gets complete control of the motor commands. If a goalpoint-achieving behavior and a obstacle avoidance behavior operate simultaneously, the goalpoint behavior would govern the vehicle completely until an obstacle was encountered. At that point the higher-priority obstacle avoidance behavior would take complete control of the vehicle and execute an avoidance maneuver.

See [2] for a detailed dicussion of the trade-offs of these approaches and several other less used arbitration methods.

4.4 Deliberative Avoidance

4.4.1 Feature-based Avoidance

Feature-based avoidance seeks to expand on the limitations inherent in a purely stateless reactive system. Feature detection uses a very limited view of state to negotiate tricky or dangerous obstacles. An example of a feature that feature detection could avoid is the so-called "canyon" scenario, where the vehicle enters an area that has walls on three-sides. A stateless reactive system will probably just make a loop in the middle of the canyon. If a simple system can identify that there are obstacles not easily avoided, it can try to wall-follow around the obstacles, an extremely useful technique that would probably get the vehicle out of the canyon. Wall-following can be initiated by changing the weightings of parameters within the existing behaviors or activating entirely new behaviors. Another useful feature could be to keep a record of areas that the vehicle had previously occupied , and discourage the vehicle from entering these areas. This simple method can keep the vehicle from the cycles in which reactive systems get stuck.

4.4.2 Map-based Avoidance

Map-based avoidance attempts to use path-planning algorithms on map data to generate a safe path that is consistent with the goals of the system. The format of the map will likely determine what kind of path-planning algorithm is optimal. Path-planning algorithms can be generalized into two primary groups - grid-based and geometry-based. Geometry-based models represents the obstacle space as a list of geometric figures (e.g. spheres, polygons, or convex solids). See [13] for a geometry-based optimal path-planning approach. Grid-based models lend themselves well to maps that can be easily translated into 2-D or 3-D occupancy grids. An occupancy grid contains either a binary value associated with whether an obstacle is judged to be in a given grid cell, a percentage likelihood that the cell has an obstacle, or a cost associated with the cell, where cells judged to contain an obstacle have a higher cost. An algorithm attempts to select a path through the cells of the grid from the current location to a goal point that is judged to be the best. The best may be determined by lowest cost, shortest, or some combination of the two. The relative merits of the grid-based A* approach and a breadth-first search approach are discussed in [6].

Some planning techiques are guaranteed-optimal, meaning that they will find the optimal path in a map if allowed enough time. Unfortunately, these techniques are very processing intensive and may not be able to run in anything approaching real-time given the processing constraints on the vehicle. It may not be required that an optimal path be generated though, as new information could be discovered at any point which will change the optimal path. Thus it may be a good design decision to just concentrate on generating a decent but not necessarily optimal path.

5 Mission-level Obstacle Avoidance Strategy Recommendations

Many concepts have been introduced that will be part of selecting an optimal obstacle avoidance strategy for a given mission AUV and mission. An obstacle avoidance strategy will consist of an obstacle avoidance sensor, processing and filtration techniques that will extract obstacle data from the sensor, and the algorithms that will use the obstacle data to avoid obstacles and achieve the goals of the given mission.

A review of the AUV literature produced three main categories of missions being conducted today by AUVs - large-scale survey, route tracking, and reacquisition and identification. Each of these missions will be examined to determine what the obstacle avoidance requirements for the mission are, and what avoidance strategy will best negotiate the trade-offs associated with strategy selection.

5.1 Obstacle Avoidance on Survey-oriented AUVs

Requirements A primary mission of AUVs is large-scale survey. Survey-class AUVs are specifically designed to achieve bottom-coverage over a given area for high-detail sonar mapping, bathymetry, or scientific experiments (Bluefin21, REMUS, MARTIN). The primary payload on a survey will generally be a side-scan or multibeam sonar. Achieving full sonar coverage of an area will require running long low-altitude tracklines spaced at a range determined by the payload. Standard altitudes for acquiring high-definition side-scan data range from 5-15 m - low-altitude operation as far as sonar manufacturers are concerned [15]. Operating depths range from very shallow water to several thousand meters. While recording sonar data on tracklines, stability is of prime concern. If there is too much attitude or altitude instability it can compromise the sonar data and track coverage.

While a survey-oriented AUV should certainly not hit an obstacle, the avoidance requirements for the general survey mission are minimal compared to the other missions examined in Sections 5.2 and 5.3. A major caveat for design is that avoiding obstacles in the horizontal plane would force the vehicles from the survey tracklines, compromising the sonar data for that section. Thus the algorithm can be designed to avoid only in the vertical plane; sensing extent in the horizontal plane is not important. Additionally, obstacle avoidance behavior should not be initiated casually, as any instability in the vehicle could compromise sonar data. On the other hand, the mission could profit considerably from anticipating the topology of the bottom and scheduling depth changes to stay at a smooth and constant altitude. A survey-class AUV is not attempting to reach a single point in the 3-D space - its goal is rather to smoothly and safely follow a track.

Sensor selection All of these considerations can be met with a relatively simple sensor - a vertical linear echsounder array. A single echosounder is too limiting, as it cannot sense vertical extent. Additionally, setting the pointing angle of the single echosounder requires a trade-off. A more upward pointing angle can detect obstacle at greater range but does not give information that will be useful for toplology following. An echosounder mounted pointed steeply downward can give topology- following information but cannot detect obstacles at enough range to enable avoidance. With an array containing a number of beams with different pointing angles both topology-following and long-distance obstacle sensing can be accomplished. A more complicated sensor is not necessary - the most important data is ranges to obstacles, and thus the extra processing to enable scanning is not worthwhile. As the vehicle does not need horizontal extent information, a 2-D array is also not necessary.

The vertical array should have beams spaced such that the vertical swath is large, with some beams pointing steeply downward primarily for use in topology following, and some beams at shallower pointing angles or even slightly upward-looking for long-range obstacle sensing and vertical extent determination. While the ideal would be a number of narrow vertical beams that give full coverage in the vertical swath, for lower cost implementations having actual gaps in the vertical coverage would be acceptable; in this case, the pointing angles should be evenly distributed along a large vertical swath. An array of 4-6 beams with different pointing angles should still give the ability to determine enough bottom and obstacle properties for good vertical avoidance.

Algorithm Selection On an algorithmic level, much of what the survey-class AUV will be doing will be topology-following. The lower several beams of the linear array will be of primary use in near-term topology following, for use in following gentle slopes or compensating for small ground disturbances. The rangings of the more downward pointing beams can be used in a computation much like that in Figure 4 to determine bottom slope. This information can be used to schedule the depth adjustments that will keep the AUV flying at a smooth and constant altitude.

If an obstacle is detected by the more upward-looking beams, topology following on a larger scale should occur. If the linear array is dense, the obstacle should be detected on more than one beam. The vehicle should be able to use the information from several beams to determine the vertical extent of the obstacle. This information can be used to schedule the depth delta necessary to fly over the obstacle with a comfortable clearance. This should allow the vehicle to minimize the pitch angle required to clear the obstacle, keeping the vehicle as close to level as possible, preserving the quality of the sonar data.

Another important role of the algorithm is to determine when an obstacle is unavoidable. An obstacle should be considered unavoidable when the difference between the vertical extent of the obstacle and the elevation of the vehicle is larger than the maximum depth delta the vehicle can achieve in the lateral range to the obstacle. When encountering an unavoidable obstacle, the algorithm should initiate a maneuver that will spiral ascend. A spiral ascent is a steep ascent at the maximum turn rate of the vehicle - this should give maximum depth delta, and only require a lateral area the size of the turning radius. Obviously, this will disrupt the sonar data. With a sensor that has substantial range and beams with shallow pointing angles, situations in which an unavoidable obstacle is encountered should be extremely rare. An obstacle that has a vertical extent outside the ensonified area may or may be unavoidable - pitching the vehicle up should make the vertical extent detectable and the obstacle classifiable. If the vehicle finds itself in a situation where the obstacle range is less than the vehicle's turning radius, it constitutes an actively unsafe condition, and thrust should be killed.

Ice avoidance is a similar proposition, except the beams should be primarily upward-facing. A linear array with 180 degrees vertical swath should be able to accomplish both topology following and ice avoidance using a slightly more complicated version of the above algorithm.

The above algorithm is well-expressed by rules. Either the algorithm schedules a gradual depth or issues the spiral command. For this mission, the scenarios are predictable - and the desired behavior highly deterministic. Thus, implementing the algorithm should be rule-based. The extra adaptability gained by a continuous response algorithm is not worth losing the determinism and predictability of a rule-based system.

5.2 Obstacle Avoidance on Route Tracking AUVs

Requirements Survey-class AUVs tend to map out a large area of ocean floor - other AUVs have missions more centered on following paths for long periods of time, whether laying cable in situations difficult for ROVs or other submersibles (Theseus), or performing optical pipeline surveys (RAIS).

Route tracking places a more stringent requirement on the vehicle's ability to follow a specific track and reach specified points. The route as specified along a pipeline or route for cable generates a specific requirement for following a particular path in both the vertical and horizontal frames, often at extremely low altitude. The route might be occluded and run extremely close to obstacles, and failing to track along the route could constitute mission failure.

These factors combine to require that the obstacle avoidance sensor selected for an AUV with this mission be substantially more complicated than on a survey-class AUV. This vehicle needs the capability to reactively avoid in both the vertical and horizontal frames, and the necessity of reaching specific points in space may require path planning to negotiate certain difficult situations.

Sensor Selection The optimal sensor for this application is a 2-D multibeam array. The requirements of this mission justify the selection of the most complicated sensor covered in Section 2. Resolution is not a huge concern - the application is not particularly concerned with specific features of the obstacle, but it does need to know to a high degree of specificity where they are located in the space. Additionally, as the 2-D array should give 3-D information, it should be the best suited for use in constructing a 3-D map of the area for use when the route takes the AUV along a particularly difficult stretch of bottom. Of the sensors discussed, the Sonatech STA-013 seems best suited to this task. A less complicated sensor could not provide the information necessary to optimally execute this task. A linear array could be used, but would not be optimal, and would constrain the vehicle in vertical obstacle avoidance.

Algorithm Selection This vehicle will need a hybrid reactive/deliberative system to give the best chance of mission success. The reactive system will need several behaviors. A goal behavior must be implemented to follow the trajectories determined by the path-planning or pipelining tracking algorithms in the deliberative module. This behavior should try to steer the vehicle along the route at whatever altitude is required. An obstacle avoidance algorithm based on the potential fields approach and using schema-blending should also be optimal for this task. The task requires the greater adaptability of continuous response rather than the discrete determinism of a rule-based approach. Schema-based blending seems better a better choice than subsumption in this application as well - if the obstacle avoidance behavior takes sole control of the vehicle, it could send it far away from the route or pipeline, compromising the primary mission objectives. Implementing using blending should mean that the vehicle will avoid in such a manner that keeps it near the route.

The deliberative layer should include path planning capabilities for navigating complex obstacles and incorporating any *a priori* knowledge of the route. Both feature-avoidance and map-based path-planning should be useful in allowing the vehicle to successfully complete the mission even in complicated environments. [1] presents a complete approach to the problem, though they use what may be a less than ideal mechanically scanned sensor.

This vehicle will have a high degree of non-deterministic autonomy - its behavior during avoidance manuevers will be emergent and thus difficult to predict. This is certainly an uncomfortable state for an extremely expensive piece of equipment operating deep in the ocean, but the emergent properties of the algorithms give the extra measure of adaptability that should enable the vehicle to complete its mission. Given the uncertainty of the domain, accurate simulation should be an essential part of algorithm testing. A simulator should combine realistic hydrodynamic models with an accurate sensor representation. This is a substantial problem in itself - [6] considers the problem briefly.

5.3 OAS for Reacquisition/ Identification (RI) AUVs

Requirements RI AUVs are generally used in conjunction with survey or route tracking AUVs. A location or obstacle of interest is identified in an extended survey, and an RI is dispatched to the location to do an extremely close range and small-scale survey to try to identify. Obstacle avoidance on RI vehicles is perhaps the most complicated domain, as the express purpose of an RI vehicle is find and move near obstacles, not necessarily just avoid them. The OAS on an RI AUV must serve dual purposes, enabling both the reacquisition and identification portions of the mission. For reacquisition, the vehicle will be given an approximate location of the obstacle and perhaps some amount of *a priori* knowledge of the environment. It then must move to the target, avoiding obstacles in a possibly cluttered space occupied by potentially hazardous obstacles, such as mines. If a survey-class vehicle hits a wall while moving 3 knots, it may result in a bit of damage to the vehicle, and will certainly compromise the data for that segment, but the vehicle will probably continue to function. If an RI vehicle activates an underwater mine not only will the vehicle be obliterated, but other assets could be harmed. Thus the vehicle must be capable of excellent obstacle sensing and avoidance.

Sensor selection It may seem like the trade-off between reacquisition requirements and identification requirements could be an impossible one, but some sensors actually can use two different frequencies, a lower one for longer ranges with lower resolutions, and a higher frequency capability for short range, high resolution sensing (see the description of the DIDSON in section 2). If a single sensor is to be used in this task, this dual-frequency capability is essential. As this sensor will doubtless be the primary payload of the vehicle, the greater expense to get a complicated sensor will pay off in terms of mission capabilities. For the task of identification, a high-resolution imaging linear array will probably be sufficient. The vehicle should move extremely close to the object, and either record data for later review by a human operator or pass the data through a computer-aided detection algorithm. If the vehicle is doing real-time image analysis, it should be easy to include a pass on the data for obstacle detection. A horizontal linear array will not provide good vertical extent information, and will limit avoidance capabilites in the vertical frame. The true ideal sensor is a dense planar array like those discussed in [11]. Beamforming techniques on a high density array should yield extremely high-resolution 3-D data useful for both reacquisition and identification.

Algorithm Selection Given the hazardousness of the environment, the vehicle must use an extremely careful reactive algorithm. Additionally, the vehicle must have the capability to achieve a given location, which requires path-planning and map accumulation. Finally, the vehicle must acheive the adaptability necessary to operate in cluttered areas with uncertain obstacle locations. An interesting rule-based approach for operation of a rear-propelled vehicle in mine fields is presented in [6]. The approach attempts to carefully control for the safety considerations associated with rear-propelled vehicles, and depends substantially on deliberative approaches to guide the vehicle through complicated mine fields.

A different possible solution might lie in the form factor of the vehicle. A rear-propelled vehicle has significant limitations associated with the dynamics of the vehicle. Turning radius requirements and the impossibility of moving backwards make it extremely difficult to create an obstacle avoidance approach that is both safe and sufficiently adaptable. A vehicle capable of hovering has many fewer limitations. A hovering vehicle should be able to hold position, rotate in place, and translate in any direction (see Section 4.1.2. These properties seem extremely well-suited for an RI vehicle: staying relatively motionless near a vehicle while recording sonar data should reduce distortion and generate a greater volume of good data, and the limitations on the vehicle behaviors for avoidance maneuvers are reduced substantially when the dynamics of the vehicle become so much more tolerant.

6 Conclusion

This paper presents the full scope of considerations that go into the design of an obstacle avoidance system for small-diameter AUVs. It has attempted to highlight some of the important trade-offs that must be made for sensor selection, data processing, and algorithm design.

Future missions will always require more functionality for obstacle avoidance: additional sensor performance, greater autonomy, higher safety guarantees. Additionally, future requirements may include new challenges, like moving obstacles avoidance, or that operation must occur in places that are extremely difficult to navigate, like underwater caves. Care should be taken to negotiate the trade-offs inherent in obstacle avoidance system design, from sensor level trade-offs in resolution, range, expense, and processing requirements to algorithmic trade-offs of determinism versus adaptivity.

References

[1] Gianluca Antonelli, Stefano Chiaverini, Roberto Finotello, and Riccardo Schiavon. Real-time path plannin and obstacle avoidance for RAIS: An Autonomous Underwater Vehicle. *IEEE Journal of Ocean Engineering*, 26(2), April 2001.

[2] M. Carreras, J. Batlle, P. Ridao, and G.N. Roberts. An overview on behaviour-based methods for AUV control. In *Conference on Manoeuvring and Control of Marine Crafts*, August 2000.

[3] Curtis techology technical paper. `http://www.curtistech.co.uk/papers/beamform.pdf`.

[4] DIDSON web page. `http://www.didson.com/`.

[5] L. Henriksen, A. Bjerrum, and A. Ishoy. Sea trials of MARTIN - a European survey AUV. In *OCEANS 1995 MTS/IEEE*, volume 3, pages 1439–1445, Oct. 1995.

[6] John C. Hyland and Fred J. Taylor. Mine avoidance techniques for underwater vehicles. *IEEE Journal of Ocean Engineering*, 18(3), July 1993.

[7] D.M. Lane, M. Chantlter, Dai Dong Yong, and I. Tena Ruiz. Tracking and classification of multiple objects in multibeam sector scan sonar image sequences. In *Proceedings of the 1998 International Symposium on Underwater Technology*, pages 269–273, 1998.

[8] Harrison T. Loeser. *Sonar Engineering Handbook*. Peninsula Publishing, Los Altos, CA, 1992.

[9] C.D. Loggins. A comparison of forward-looking sonar design alternatives. In *OCEANS 2001 MTS/IEEE*, volume 3, pages 1536–1545, Nov. 2001.

[10] A. Martin, E. An, K. Nelson, and S. Smith. Obstacle detection by a forward looking sonar integrated in an autonomous underwater vehicle. In *OCEANS 2000 MTS/IEEE*, volume 1, pages 337–341, Sept. 2000.

[11] Fred Nussbaum, Gerald T. Stevens, and James G. Kelly. Sensors for a forward-looking high resolution AUV sonar. In *Proceedings of the 1996 Symposium on Autonomous Underwater Vehicle Technology, AUV '96*, pages 141–145, 1996.

[12] Y. Petillot, I. Tena Ruiz, D. M. Lane, Y. Wang, E. Trucco, and N. Pican. Underwater vehicle path planning using a multi-beam forward looking sonar. `http://citeseer.nj.nec.com/petillot98underwater.html`.

[13] Y. Petillot, T. Ruiz, and D.M. Lane. Underwater vehicle obstacle avoidance and path planning using a multi-beam forward looking sonar. *IEEE Journal of Oceanic Engineering*, pages 240–251, April 2001.

[14] R.L. Pietsch. Autonomous vehicle obstacle detection and tracking. In *Proceedings of the 1998 International Symposium on Underwater Technology*, pages 198–202, 1998.

[15] M.T. Shaw. Obstacle avoidance sonar-an analysis of system requirements and detection performance during near-bottom operation. In *Proceedings of the (1990) Symposium on Autonomous Underwater Vehicle Technology*, pages 223–228, June 1990.

[16] Sonatech ATOAS brochure. `http://www.sonatech.com/products/product_areas/snr_sys/ATOAS/atoas_brochure.pdf`.

[17] Tritech internation web page. `http://www.tritech.co.uk/Products/Products%20-%20PA200-PA500.htm`.