

binocular and motion stereo matching. These passive range detection techniques will form part of a machine perception system for autonomous cross-country navigation. Using Datacube image processing boards, we have designed a prototype system for binocular stereo matching that can match $256 \times 256 \times 8$ -bit grayscale image pairs in approximately one second using a disparity range of 64 pixels. We have shown the modifications needed to perform color stereo matching. The current system can easily be condensed to a board and, perhaps, a single VLSI chip operating in real-time.

ACKNOWLEDGMENT

The development of the serial version of the binocular stereo algorithm was carried out at the Computer Vision Research Laboratory, The University of Michigan, Ann Arbor by Saied Moezzi and Terry E. Weymouth under a contract from General Dynamics. We also would like to acknowledge the assistance of Charles Cole and Carl Hobson for reviewing this manuscript and the assistance provided by Claude Bugenis and John Giampa during the early phases of this project.

REFERENCES

- [1] S. R. Sternberg, "Architectures for neighborhood processing," *Proc. IEEE Pattern Recognition and Image Processing Conf.*, 1981, pp. 374-380.
- [2] R. N. Loughheed, "A high speed recirculating neighborhood processing architecture," *Proc. SPIE*, vol. 534: Architectures and Algorithms for Digital Image Processing II, 1985, pp. 22-33.
- [3] A. P. Reeves, "Survey: Parallel computer architectures for image processing," *Computer Vision, Graphics, and Image Processing*, vol. 25, pp. 68-88, 1984.
- [4] W. D. Hillis, *The Connection Machine*. Cambridge, MA: MIT Press, 1985.
- [5] M. Annaratone *et al.*, "The Warp computer: Architecture, implementation, and performance," Carnegie Mellon University Robotics Institute Tech. Rep. CMU-RI-TR-87-18, 1987.
- [6] J. P. Hayes *et al.*, "A microprocessor-based hypercube supercomputer," *IEEE Micro*, pp. 6-17, Oct. 1987.
- [7] BBN Advanced Computers Inc., "The Butterfly GP1000 parallel processor," *AI Review, AAAI*, pp. 59-60, 1988.
- [8] S. T. Barnard and M. A. Fishler, "Computational Stereo," *ACM Computing Surveys*, vol. 41, Dec. 1982.
- [9] P. Anandan, "Measuring visual motion from image sequences," Univ. Massachusetts COINS Tech. Rep. 87-21, Mar. 1987.

Real-Time Algorithms and Data Structures for Underwater Mapping

DAVID N. OSKARD, TSAI-HONG HONG,
AND CLIFFORD A. SHAFFER

Abstract—As part of the multiple autonomous underwater vehicle (MAUV) project at the National Institute of Standards and Technology, a spatial mapping system has been developed to provide a model of the underwater environment suitable for autonomous navigation. The sys-

Manuscript received March 8, 1989; revised January 10, 1990. This work was partially presented at the SPIE Conference on Sensor Fusion in Cambridge, MA, November 1988.

D. N. Oskard is with the Robot Systems Division, National Institute of Standards and Technology, Gaithersburg, MD 20899.

T. H. Hong is with the Department of Computer Science and Information Systems, The American University, Washington, DC 20016 and the Robot Systems Division, National Institute of Standards and Technology, Gaithersburg, MD 20899.

C. A. Shaffer is with the Department of Computer Science, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061.
IEEE Log Number 9037608.

tem is composed of multiresolution depth maps designed to integrate sensor data with an a priori model, an object/attribute database for storing information about detected objects, and a set of flags to monitor abnormal or emergency conditions in the environment. The structure of the mapping system and the algorithms used to map terrain and obstacles detected by acoustic sonar are described.

I. INTRODUCTION

In an autonomous system, the ability to predict, estimate, and evaluate the state of the environment largely determines the capabilities of the system. Much work has been done in the area of sensor fusion to enable autonomous robots to operate effectively in the world. Levi has discussed sensor fusion techniques that have been applied to autonomous vehicle navigation [12]. Luo *et al.* have presented a good overview of multisensor integration and fusion in mobile robots and other intelligent systems [13], [14]. Multiple sensor systems have been used in highly structured applications such as road following [5], [29] as well as unstructured environments in both indoor [4] and outdoor settings [2], [7].

For the application of autonomous underwater vehicles, a means must be developed to model the world accurately enough to perform such tasks as obstacle avoidance, path planning, and long-range mission planning with sufficient speed to operate in real-time. Research at the National Institute of Standards and Technology (NIST) has led to the development of a real-time control system (RCS) for multiple autonomous underwater vehicles (MAUVs) performing cooperative tasks [8]. Based on the theory of hierarchical control [1], the system is designed to perform both high and low-level planning tasks with a world model that provides timely information at multiple resolutions. The RCS is divided into several logical levels of complexity; at each level, the world model contains six modules:

- 1) Update Module—interprets and adds incoming sensor data to the model,
- 2) Data Server—provides requested model information to the planner,
- 3) Sensor Prediction Module—provides expected sensor data estimates for use in validation by the sensor processor,
- 4) State Evaluator—monitors emergency conditions and the state of each RCS module,
- 5) Learning Module—categorizes information supplied by the database and other modules,
- 6) Database—for storage of information.

The emphasis of this paper is on describing the update module and portions of the database. For information on the other modules and the functionality of the world model, refer to [8].

In its current form, the MAUV world model mapping system consists of depth maps of lake bottom terrain, state variables, emergency flags, and a database of objects. We will focus primarily on the lake bottom and obstacle mapping functions of the world model and considerations in their design and implementation. The system uses sonar range data to update the terrain maps of the lake bottom. Stewart has done similar work in the underwater environment [27], while other researchers have used range data in various modeling applications. Grimson and Lozano-Perez have used range and tactile data for object identification and localization from a set of known objects [6]. Roggemann *et al.* [22] and Tong *et al.* [28] have used laser radar and forward looking infrared sensors to segment and enhance features of man-made objects in a cluttered environment.

In the following sections, we discuss requirements of the data structures used to represent the underwater environment, and then describe the mapping system. Algorithms used in the

update module are presented, with an emphasis on the concept of confidence-based mapping. Finally, we point out some of the strengths and weaknesses of the system and suggest directions for further research.

II. CRITERIA FOR DATA REPRESENTATIONS

The first and most important task in the development of an effective world model is choosing an appropriate data representation. In most mapping systems, a primary distinction can be made as to whether given information is spatial or non-spatial data. Depth maps or other models of a 3-D environment are referred to as spatial data structures, whereas structures representing objects and their attributes are primarily nonspatial data structures. Although most objects in the world contain a spatial component, namely a location and 3-D structure in some frame of reference, a suitable representation for the associated descriptive information is an object/attribute database. A simple object/attribute database has been developed to store data of this nature. Each object structure contains an identification number along with a 32-bit mask describing known capabilities. A linked list of attributes stores sensed characteristics for each object (e.g. X , Y , Z coordinates; velocity vectors; etc.) with the ability to store multiple entries over time in a circular queue. It is flexible enough to allow 3-D representations of objects to be stored as attributes themselves by storing pointers to the 3-D data structures in the attribute list.

Choosing a data representation to model the underwater environment is a more difficult problem. One solution is to implement a complete 3-D data structure to model the lake bottom and all detected obstacles. An example of such a representation is the octree [10], [15], which considers the working universe to be one large cubic region. The universe is recursively divided into 8 equal cubic sub-regions or "octants" until each sub-region is considered to be homogeneous, that is, either completely empty space or completely full, containing an object or the lake bottom. This approach has been used successfully to represent an automated factory environment [9]. However, in the application of autonomous underwater vehicles, an octree representation cannot keep up with the demands of the real-time environment when CPU resources must be shared by more than one RCS module. The cost of bookkeeping is high with octrees, as each new sonar reading may cause the tree to decompose or merge several nodes. With each decomposition, eight new nodes must be allocated for the newly created subregions, and every merge operation requires the reclamation of unused nodes. Considering that the CPU is shared and the sampling rate for the six sonar sensors on a prototype vehicle is 0.6 s [11], a full 3-D representation was not attempted.

III. GLOBAL AND LOCAL MAP CONFIGURATION

A less computationally intensive approach is to use a $2\frac{1}{2}$ -D representation, storing only depth information for any map location. The current autonomous underwater vehicle (AUV) world model uses two types of $2\frac{1}{2}$ -D data structures for its mapping scheme: a set of global maps, each of which contains data for the vehicle's operational domain, and region-of-interest maps, which only store a localized area around the AUV's current location. (Vehicle locations are obtained by triangulation using an underwater transponder system. Distances to each of the transponders are used to compute the global position of the vehicle.) The region quadtree [23], which is the 2-D analogue of an octree, is used to represent depth maps at a global level. The root node of the tree represents a 2-D square area that is recursively decomposed into 4 quadrants, sub-quadrants, and so on, until each node contains only one depth value. Because each branch of the tree contains half as many nodes as an octree branch, quadtree updates are generally faster than octree updates and the structure consumes less storage as well.




Fig. 1. A region quadtree representation for an example depth map. Only nodes containing more than one depth value are further decomposed. Each branch of the tree corresponds to northwest, northeast, southwest, and southeast quadrants of a given region.




Fig. 2. Illustration of quadtree implementation. Each GRAY node has four fields, labelled NW, NE, SW, and SE. Each field has 2 subfields: the value and the GRAY/Child discriminator.

As with octrees, a node is only further decomposed if it is not homogeneous. This property leads to storage savings over array data structures when the data is sparse, or when the data set contains large homogeneous regions such as islands or flat regions of a lake bottom (see Fig. 1). The quadtree approach also allows point and linear feature data to be represented using compatible data structures as described in the following [23]–[25].

A. Global Quadtree Implementation

The quadtree representation used for this project is a compact implementation of a pointer-based quadtree. The pointer-based quadtree explicitly stores pointers from each internal node I to I 's four children. This is in contrast to the linear quadtree which stores a list of leaf node values sorted by block location [3]. The linear quadtree has been found to be effective for disk-based applications. However, the MAUV system brings maps into primary memory when in use.

Our quadtree implementation reduces storage requirements over the traditional pointer-based quadtree implementation which stores four pointer fields and a value field with each node of the tree. Instead, we only store pointer fields with internal nodes, commonly referred to as GRAY nodes. A GRAY node G is represented by a block of four contiguous descriptor fields, each corresponding to one of G 's children. In the case where a child is itself a GRAY node, the value stored in the descriptor is a pointer to the block of storage representing that child. In the case where a child is a leaf node, the value stored is simply the depth value for the block. This representation is illustrated by Fig. 2, which shows the storage structure for the quadtree of Fig. 1. In Fig. 2, each of the four child fields has an extra bit field associated with it to distinguish between GRAY children (a "0" bit) and leaf children (a "1" bit). In actual implementation in the C programming language on Motorola 68020 processors, GRAY nodes are represented by four contiguous 32-bit words. The low-order bit of each 32-bit value distinguishes between a GRAY child and a leaf child.

This implementation avoids the wasted storage of four null pointers associated with each leaf node to indicate that no children are present. The data structure stores a single pointer to each GRAY node and a value for each leaf node (i.e., no explicit pointers to leaf nodes are stored). Thus the total storage required is a single 32-bit value for each node in the tree for both leaf nodes and GRAY nodes. Since a quadtree of L leaves

requires $(4L - 1)/3$ nodes, this is also the number of 32-bit words required.

Three region quadtrees are used to represent different aspects of the global depth maps. One of these is the a priori map, which contains data from a survey of the lake bottom and does not change during the mission. A priori data was collected for an approximate area of 1.6 km² in Lake Winnepesaukee, NH (the 1987 test site). The resulting set of irregularly spaced data points was mapped to a regular grid using Renka and Cline's triangulation algorithm [20], [21] and converted to quadtree format using a region quadtree construction algorithm [26]. The a priori quadtree in this instance does not yield much storage savings over a grid structure, as the lake survey data is of low resolution and could be represented by an array smaller than 512 by 512; however, using a quadtree provides a convenient means for integrating a priori map information with incoming sensor information.

A separate sensor quadtree is used to store higher resolution depth values collected from the AUV sonar sensors during vehicle test runs. Both downward- and forward-looking (obstacle avoidance) sonars are used in refining the sensor map, which is overlaid onto the a priori data to give the most accurate view of the world. A third quadtree stores a depth confidence value for each node in the tree. The confidence map supports the important function of distinguishing spurious sonar readings caused by debris or signal inconsistencies from actual obstacles that must be detected and avoided. Additional bit fields in the confidence map node value are used as flags to mark the vehicle's track through the water, as well as to record whether a given node has been updated with downward or forward-looking sonar. This extra information could be useful in classifying newly discovered objects and in examining the AUV's complete path after a mission. The region quadtree is particularly efficient for sensor and confidence map representation, since unexplored portions of those maps are empty. Such areas can be represented by a small number of nodes in the tree.

B. Point and Linear Feature Representation

In addition to the quadtrees used to represent region data, point and line storing quadtrees have been implemented to provide locations of known objects and topographic features of the lake bottom used in high-level path planning. This simplifies tasks such as locating the nearest other vehicle to a given location or plotting a course along linear topographic features like ravines or underwater pipelines. The planner can then retrieve information in a region of interest without having to search the object/attribute database directly. The global maps provide an efficient method for storing large quantities of data while maintaining the spatial relationships between different types of features.

Point objects are represented by means of a point region (PR) quadtree [23]. The PR quadtree is similar to the region quadtree except that its decomposition rule is based on the number of point objects located within a block. If a block contains no points or a single point, then this block is represented by a single leaf node. Otherwise, the block is decomposed into quadrants and sub-quadrants recursively, until each block contains at most a single point. Implementation of the PR quadtree is identical to that of the region quadtree illustrated in Fig. 1(a). Each point object is identified in the object/attribute database by means of an index number. Leaf nodes of the PR quadtree store this index. Operations such as finding the nearest object to a specified location and locating all objects within a specified distance of a location have been implemented.

Linear features are represented by means of the PMR (or polygonal map, random) quadtree [17]. Upon insertion of a line segment into a block B, the PMR quadtree decomposes B exactly once if the number of line segments already contained

within B exceeds some threshold (in our case, the threshold is four segments). This does not guarantee that a block will contain four or fewer segments, since segments lying very close together may pass through the same block even after decomposition; however blocks rarely contain much more than four segments. The main objective of this decomposition rule is to reduce the number of line objects contained within a single block so as to make processing efficient. Note that varying numbers of line segments may be associated with a single block. Our PMR quadtree implementation is identical in structure to that of the region and PR quadtree implementations, except that leaf nodes store a pointer to a linked list of the line segments contained within that node. Operations such as finding the nearest line segment to a point have been implemented. The PMR quadtree has yet to be integrated into the map representation and planning portions of the MAUV system. However, availability of the PMR quadtree implementation will enable future versions of the planner to operate directly on the higher level representation of ridge and trench lines rather than on the low level data currently provided by the sensor and confidence maps.

C. Local Maps

The AUV world model makes use of grid or array data structures for storing local region-of-interest maps of varying resolutions. Arrays are used for their fast, constant access and update time and for ease of implementation. Local maps are generated from the global quadtree database, first by extracting a priori map data for the region, then overlaying the data stored in the sensor and confidence quadtrees, which are presumed to be more accurate than the lake survey information since they reflect recent sensor input. In fusing these three sets of data, all three quadtrees are traversed over the local map region. The value stored in the confidence quadtree for any map location is used to determine whether or not there is sensor data for that area. Because the updating algorithm only stores data in the sensor quadtree if it has a confidence measure above the level assigned to a priori data, any node for which there is sensor data uses the sensed value. The local map uses a priori knowledge only if insufficient sensor data has been collected for that node. Confidence quadtree values are also copied into the local map to be used extensively in the updating algorithm.

Each map has an offset coordinate which defines its position relative to a global origin. As the vehicle moves to within a predefined distance of the current local map's edge (in this case 64 meters defined as 1/4 the width of the map), a new shifted map is generated in a second buffer so as to keep the AUV nominally centered. Before this is done, any newly acquired sensor and confidence data in the old local map must be written to the respective quadtrees, making it available for the new map. During the generation process, incoming sonar data is used to update the old buffer (still being used by the planner), but is also held in a data queue to be applied to the new buffer after it has been extracted. When the new map has been completed, the software simply swaps the two map pointers and updates the global coordinate offsets of the new local map. With this double buffering arrangement, the planner can still access current data while a new map is being generated. This provides the capability for asynchronous world model queries as required in the AUV's multiprocessing environment.

The algorithm that copies data from the quadtree representing the global map to the local map array is a simple directed traversal of the quadtree beginning at the root node. If the root is a leaf node, then the value of the root is copied to the portion of the array overlapped by the root. If the root is a GRAY node, the algorithm is recursively called for only those children of the root that overlap the local map. Thus, only that portion of the global quadtree overlapping the local map is processed.

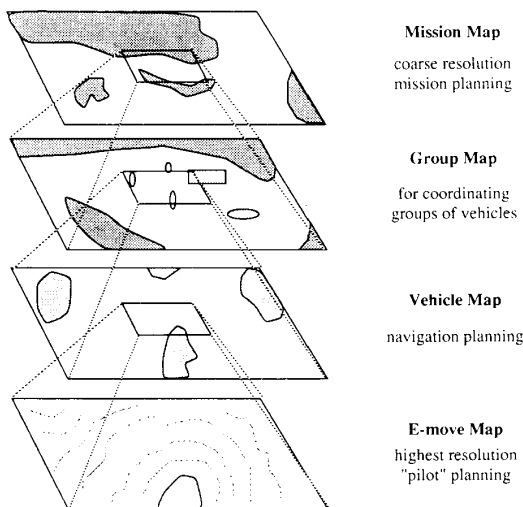


Fig. 3. The MAUV mapping hierarchy.

Updates of a global quadtree from the local map take place when the AUV moves beyond the center region of the local map. The quadtree update algorithm is analogous to the array update algorithm, except that one can view the quadtree update algorithm as being driven by a traversal of the array rather than the quadtree. Each pixel of the array is visited not in raster scan order, but in the order that the pixels would be visited by a preorder traversal of the corresponding quadtree. The pixel addresses for such a traversal are calculated by interleaving the bit representation of the x and y coordinates for each pixel, and visiting the pixels in ascending order; this is known as Morton order, N -order, or Z -order [3]. As each local map pixel is visited, it is compared against the corresponding node in the global quadtree, whose node value is updated if required. Processing in Morton order ensures that all pixels of the current quadtree node will be checked before the next node must be located.

D. The Mapping Hierarchy

As mentioned in Section 1, the RCS is a hierarchical control system, meaning that tasks to be performed are functionally decomposed into distinct levels of complexity. Various levels of the control hierarchy require different local map resolutions. Different types of data may be needed at each level. Generally, the resolution of maps at each level differs by an order of magnitude or more. All local maps are implemented as array data structures and only the lowest level (highest resolution) local map updates the global quadtrees. Fig. 3 shows the mapping hierarchy for a generalized data set. In the current implementation, the mission level map divides the area into a coarse grid of approximately 25×25 picture elements or pixels per square mile, each pixel storing the average depth of the corresponding area, along with probabilities of detection and destruction of the AUV at each location based only on *a priori* knowledge. It functions primarily as a guide in tactical planning.

The next two levels in the hierarchy, the group and vehicle levels respectively, share the same local map for this data set. Each pixel of the local map stores the minimum and maximum known depths over a 4×4 meter area. It serves the purpose of providing information for high level navigation tasks, such as determining the probability that an area is traversable by one or more vehicles. This map is updated as new information is added to the lowest level map, known as the elemental move or "E-move" level map.

The *E*-move local map has the highest resolution, and is used in determining the traversability of a path between two specified points. The world model returns a probability that the path is traversable based on the information in this map and the needs of the planner. For example, the output may be a percentage of pixels for which the AUV clears the lake bottom over the hypothesized path. In the simplest case, the world model can provide a probability of 1 if all of the pixels are traversable, or 0 if any are obstructed. Typically, the pilot planner will query the world model for the traversability of several paths, using a modified A^* search algorithm [19] to choose the safest, most efficient path available. The *E*-move map is also the level updated directly by sensor readings; its modifications are propagated up through the hierarchy.

In addition to the world model maps, a set of emergency flags are constantly monitored for conditions which require immediate action. These flags are raised by the lowest level of the sensory processing module when abnormalities are detected in sensor data. Each bit in the emergency flag register corresponds to a different condition, such as low fuel, water leakage, sensor failure, etc. One bit is reserved for an imminent collision condition, meaning that the system has not responded quickly enough to avoid an obstacle, an oncoming ship, or possibly even another AUV. This guards against the inherent delay in detecting obstacles, adding them to bottom maps, and planning paths around them. It also provides a backup safety mechanism for the case that a bug in the software has caused an object not to be detected. While the mapping system is designed for use in dynamic path planning, its performance in terms of object detection delay is ultimately tied to sensor cycle times, CPU efficiency, and resource allocation, in addition to the algorithm design and implementation.

IV. THE MAPPING SYSTEM ALGORITHM

For the first implementation of the mapping system, the confidence assignment algorithm has been greatly simplified. A complete implementation would include a stochastic analysis of the sonar data, taking into account the increased likelihood of error as the range to the target increases, as well as uncertainties in the AUV's navigational data [27]. Here we have assumed that the incoming data has been filtered by a sensory processing module to minimize such inaccuracies.

At the beginning of a mission, the AUV initializes the global and local maps, reading available *a priori* knowledge from a secondary storage device. In our prototype, the lake bottom survey map is read from an optical disk storage medium into the *a priori* quadtree. The sensor quadtree is initially composed of a single, empty node, though it could also contain sensor data stored from previous missions if available. Likewise, the confidence quadtree is initialized as a single node containing a base confidence value, unless there is confidence data from a previous mission. In general, the AUV world model starts up in a state of total dependence on *a priori* knowledge, gradually becoming more reliant on the current sensor map as data is collected.

A. Confidence-Based Mapping

The concept of confidence-based mapping is central to the map updating algorithm. By modifying values in the confidence map, the interpretation of incoming sensor data can be controlled, which in turn influences the vehicle's behavior. Confidence values are incremented or decremented from an initially assigned base value as confirming or conflicting information is received. The range of values is set by the designer at compile time, the default being in the range 0 to 20, where a value of 20 corresponds to 100% confidence. The upper bound on a pixel's confidence guards against the problem of not detecting a moving object due to overconfidence in data that no longer is

accurate. When the system receives conflicting sensor data, the local map update module waits until a pixel's confidence drops below a threshold before changing its depth value. The threshold value is a prespecified constant determined by the maximum change a single sensor reading can affect.

The other constants in the algorithm are the base confidence values. One is assigned to a priori data; the other, to pixels whose depth values have just been modified with sensor information. These constants directly affect the behavior of the mapping system by determining the amount of conflicting sensor data required to cause a change in the depth map. If the base confidence given to the a priori map is relatively high, it will take several conflicting sonar readings before a pixel's depth is updated with new sonar information. The higher the a priori confidence, the longer the AUV will take to "see" and subsequently avoid a new obstacle. On the other hand, higher initial confidence values reduce the effects of erroneous sensor readings.

Another feature of the confidence algorithm is the ability to weight the input of each type of sensor differently. In our current implementation, depth sonar readings are able to modify a pixel's confidence by 10% per reading, while obstacle avoidance sonars only modify confidence values by 5% per reading. Although these particular values are somewhat arbitrary—based on the assumption that depth sonar is more accurate than obstacle avoidance sonar due to fluctuations in the AUV's orientation—the basic idea has been to create a mapping system whose behavior can be modified by changing its parameters, without requiring substantial effort on the part of the programmer.

B. Evaluating Sonar Data

When presented with new information, a world model may classify it into one of three categories. The new data may affirm knowledge already in the model, causing the confidence in that knowledge to be increased; it may present conflicting information, in which case an intelligent world model would decrease its confidence in that knowledge; or the incoming data may be classified as irrelevant, requiring no action. Our ability to function in the world as humans depends heavily on our ability to make such judgements. The same is true in autonomous navigation. In the AUV prototype, inputs are limited to two types of range data; downward-looking sonar and forward-looking obstacle avoidance sonar. Because of this simplification, data classification of this sort is fairly straightforward.

Although the return from a sonar source can be somewhat eccentric [16], a sonar reading is assumed to represent a conic beam of approximately six degrees in diameter, with a coordinate transformation describing its position with respect to the vehicle. The map updating routine currently handles each sensor separately, using different approaches for the downward and obstacle avoidance data to modify the local map. As a first step in processing, a sonar beam is projected into the X - Y plane, transforming its coordinates with respect to the vehicle into world model coordinates with respect to the global origin.

1) *Depth Sonar*: In updating the map from the downward-looking sonar, the algorithm first computes an approximate neighborhood size of pixels to be updated around the current AUV location depending on the width of the beam and the distance to the lake bottom. This approach is possible because the pitch and roll properties of the vehicles are negligible, causing the depth sonar to always point straight downward. Given that the beam width of the sonar is fixed and the range is returned by the sensor, a closed-form trigonometric solution can be performed using a lookup table. Although the 2-D projection would be best represented as a circular region, for our purposes, a square neighborhood is sufficiently accurate and more efficient to update. The depth stored at each pixel of the neighbor-




Fig. 4. Three stages of a map update from depth sonar: (A) Incoming sonar data conflicts with the world model. (B) Confidence values for the corresponding world model pixels are decremented (depth is not modified yet). (C) After repeated conflicting readings, confidence drops below the threshold, the depth is updated, and a new confidence assigned.




Fig. 5. Updates from obstacle avoidance sensors remove false obstacles in the world model by increasing depth values in the map. Obstacles are added by decreasing the depth, in effect, raising the bottom of the lake model.

hood in the local map is compared to the observed sonar reading. If the two values are not within an acceptable margin of error, the conflicting data causes the pixel's confidence to be lowered. If the two depth values are in agreement, the confidence value is incremented unless it has already reached the maximum allowed. Whenever a pixel's confidence value drops below the predefined threshold, it takes on the new depth reading and is assigned a base confidence value (see Fig. 4). For the depth sonar, all information is classifiable as either conflicting or agreeing with the knowledge already in the model. None of the data is irrelevant in this case.

2) *Obstacle Avoidance Sonar*: The obstacle avoidance mapping algorithm is more complicated. Here the projection of the cone into the two-dimensional plane approximates a triangular region. Interior points in the projection are determined using an algorithm presented by Newman and Sproull [18] given the vertices of the triangle. The cone itself is approximated by two planar surfaces representing the top and bottom surfaces of the cone, calculated at three degrees above and below the central axis of the cone respectively. Due to the relatively coarse resolution used in the obstacle avoidance algorithm (0.5 m^3 per pixel) and the narrow width of a sonar beam, this does not introduce significant error into the calculations. As with the depth sonar algorithm, each pixel in the 2-D projection is examined and updated if its confidence value drops below the threshold. Forward-looking sonar readings provide two types of information: a given pixel may be clear, or it may be obstructed by an obstacle. When the AUV detects an obstacle, the mapping algorithm adds the information to the local map by raising the modeled bottom of the lake at that location (i.e. making it shallower) (see Fig. 5). It is also an essential function of the world model to be able to remove hypothesized obstacles in the local map as well as add them. For each pixel in the triangular projection, if the three dimensional distance (measured along the sonar cone trajectory) from the sonar source to the current pixel being examined is less than the range returned by the sensor, the pixel is assumed to be clear. No obstacle was detected there, so the depth at that location in the local map should reflect this information. Its value should be greater than or equal to the depth of the bottom surface of the sonar cone at

that location, since any object obstructing the beam would presumably cause the sensor to return the range to that object. If the local map value is shallower than the beam, it conflicts with the new sensor data and the confidence value is decremented. If this results in a confidence lower than the threshold, the pixel is reassigned the depth value of the bottom surface of the cone and a new base confidence value. Note however that a local map value in agreement with sonar information does not necessarily increase its confidence. The sonar beam may be projected in front of the vehicle when the AUV is near the surface, and a clear reading near the surface would not yield any information about the depth of the lake bottom if we already have some a priori knowledge that the lake is approximately N meters deep. In this case it would be considered irrelevant data.

The same is not true for pixels in the projection whose distance from the sonar source is greater than or equal to the range returned by the sensor. These pixels correspond to detected obstacles and the depth values in the local map are compared to the top surface of the cone.¹ Here the local map data should be at least as shallow as the top surface of the beam to be in agreement with the sensor reading. If the map data does agree, it represents confirmation of an existing obstacle and the confidence value should be incremented. It should be noted that this confirmation only supports the hypothesis that there is an obstacle at the depth it was detected; no conclusions can be drawn as to the true height of the object or whether it extends all the way to the lake bottom. Although this may lead to some inaccuracies in the depth map, an array of multiple sensors pointing in different directions minimizes this effect.

In a similar manner, if the model continually disagrees with the sensor reading, the confidence is decremented until the depth value is reassigned to the depth of the top surface of the cone, making the model shallower. Its confidence is again initialized to a base value.

V. MAPPING SYSTEM PERFORMANCE

Processing time for a sonar input is roughly proportional to the number of pixels that must be examined in the depth map. On the software development system used at NIST, a Sun-3 workstation, obstacle avoidance processing ranges in speed from 30 milliseconds of CPU time for short distances of approximately 10 meters, to 100 milliseconds for the maximum accepted reading of 50 meters. Depth sonar updates are practically instantaneous, due to the lack of computation required. Since the sensor cycle time is approximately 0.6 s and the vehicles used are equipped with an array of 5 obstacle avoidance sensors [11], it is not possible to process each sonar reading separately while sharing the CPU with several other processes. To improve performance, a preprocessing filter has been added to the front end of the system that combines similar sonar readings from the same sensor into a single, more heavily weighted reading. The filter passes the average range returned by the sensor, the average position of the vehicle, and the number of combined readings to the mapping system. Tolerances for what constitute similar readings are parameters to the system. Only the confidence assignment routine must be modified to accommodate the filtered data, multiplying the weighting factor for the current sensor by the number of combined readings. This filtering process enables the system to operate well within hardware timing constraints.

Another time-consuming operation is the transfer of data between local maps and global maps. Each time the AUV crosses the boundary of the current local map (really 64 m from

the map's edge as described in Section III-C), the world model must examine the entire array for pixels that have been updated, modify the sensor quadtree, and extract the new local map. This process takes on the order of several CPU seconds. Due to the 1.5 m per second maximum speed of the vehicle, this occurs once every 42 s in the worst case. For best results, a separate process should perform this task in the "background" while processing of the sensor data continues using the old map buffer. As previously mentioned, data collected while the new window is being extracted must be stored in a queue to be applied to the new local map. Although this feature has not been fully implemented yet, it is expected that the queued data will not cause the system to fall behind in its processing if filter parameters are set properly. If necessary, the algorithm can be further modified to only store readings which caused the old local map to be updated.

VI. CONCLUSION

The MAUV mapping system has been developed to provide an effective means for autonomous navigation. During the first year of the project, much work has been done on the problem of short term planning for obstacle avoidance. The $2\frac{1}{2}$ -D depth maps using the confidence-based mapping algorithm have proven to be an adequate solution. A dynamic planner continuously queries the world model for the probability of safe traversal over the planned trajectory and is able to modify the vehicle's path to avoid obstacles as soon as they are detected. It typically takes 3-4 supported readings from the obstacle avoidance sonar before the map is modified to show the obstacle, though this property can be adjusted by altering the confidence parameters of the system. For example, a low base confidence value would make the mapping system more responsive to sonar readings, decreasing obstacle detection time, but increasing the probability of error due to sensor noise.

Because of the $2\frac{1}{2}$ -D nature of the data structures, the current implementation is unable to model underwater formations such as caves and ledges. All obstacles are modeled as elevated portions of the lake bottom. With this representation, the path planning module does not consider routing alternatives which would go under detected obstacles. Another drawback of the current map updating scheme is that it does not take into account the fact that most underwater terrain features are smooth. In general, neighboring pixels in the depth map will have similar depth values. This information could be used in the assignment of confidence or depth values stored at a map location, for example, assigning low confidence values to pixels in areas of high depth variance. The system also does not yet include facilities for representing the 3-D shape of objects; however, the locations of objects can be stored in the PR quadtree with associated 3-D data maintained in the object/attribute database. Because the database can be configured to store arbitrary parameters for each data type, descriptive information such as volume, centroid location, and number of holes may accompany an entry in the PR quadtree. As mentioned earlier, one of the entries in the database may even be a pointer to an extensive 3-D representation of the object.

The next generation mapping system will include the full integration of the PMR quadtree with the planning modules to store route and topographic information. This will enable high level planners to follow contours of the lake bottom and find paths of minimal risk to the vehicles. Additional sensor data from laser ranging devices and underwater cameras for inspection of objects may be included as well. Another topic for future research is the use of spherical representations of the underwater environment in real-time world modeling. Much more work is needed to fully test and develop the system, but progress so far indicates that this approach will produce a fast and reliable system for real-time underwater mapping.

¹One detail worth mentioning is that the 2-D projection of the cone is extended by two pixels so that it will not only include the projection of the beam, but also a portion of the detected obstacle.

ACKNOWLEDGMENT

The authors would like to thank Dr. Tsung-Ming Tsai and Shujen Chang for their contributions to this work.

REFERENCES

- [1] J. Albus, "Hierarchical control for robots and teleoperators," IEEE Workshop on Intelligent Control, Albany, NY, Aug. 26-27, 1985, pp. 39-49.
- [2] M. Daily *et al.*, "Autonomous cross-country navigation with the ALV," *Proc. IEEE Int. Conf. Robotics Automat.*, Apr. 1988, pp. 718-726.
- [3] I. Gargantini, "An effective way to represent quadrees," *Commun. ACM*, vol. 25, Dec. 1982, pp. 905-910.
- [4] G. Giralt, R. Chatila, and M. Vaissset, "An integrated navigation and motion control system for autonomous multisensory mobile robots," *Robotics Research: The First International Symp.*, Cambridge, MA: MIT Press, 1984, pp. 191-214.
- [5] Y. Goto and A. Stenz, "The CMU system for mobile robot navigation," *Proc. IEEE Int. Conf. Robotics Automat.*, Mar. 1987, pp. 99-105.
- [6] W. E. L. Grimson and T. Lozano-Perez, "Model-based recognition and localization from sparse range or tactile data," *Int. J. Robotics Research*, 1984, pp. 3-35.
- [7] S. Y. Harmon, "The Ground Surveillance Robot (GSR): an Autonomous Vehicle Designed to Transit Unknown Terrain," *IEEE J. Robotics Automation*, 1987, pp. 266-279.
- [8] M. Herman and J. Albus, "Overview of the multiple autonomous underwater vehicle (MAUV) project," in *Proc. IEEE Int. Conf. Robotics Automat.*, Apr. 1988, pp. 618-620.
- [9] T. H. Hong and M. O. Shneider, "Describing a robot's workspace using a sequence of views from a moving camera," *IEEE Trans. Pattern Anal. Machine Intell.*, Nov. 1985, pp. 721-726.
- [10] C. L. Jackins and S. L. Tanimoto, "Oct-trees and their use in representing three-dimensional objects," *Computer Graphics and Image Processing*, vol. 14, Nov. 1980, pp. 249-270.
- [11] J. Jalbert, "Low level architecture for the new EAVE vehicle," in *Proc. Fifth Int. Symp. Unmanned Unethered Submersible Tech.*, Jun. 1987, vol. 1, pp. 238-244.
- [12] P. Levi, "Principles of planning and control concepts for autonomous mobile robots," in *Proc. IEEE Int. Conf. Robotics Automat.*, Mar. 1987, pp. 874-881.
- [13] R. C. Luo and M. G. Kay, "Multisensor Integration and Fusion in Intelligent Systems," *IEEE Trans. Syst. Man Cybern.*, pp. 901-931, Sept./Oct. 1989.
- [14] R. C. Luo, M. Lin, and R. S. Scherp, "The issues and approaches of a robot multisensor integration," *Proc. IEEE Int. Conf. Robotics Automat.*, Mar. 1987, pp. 1941-1946.
- [15] D. Meagher, "Geometric modeling using octree encoding," *Computer Graphics and Image Processing*, vol. 19, Jun. 1982, pp. 129-147.
- [16] H. P. Moravec and A. E. Elfes, "High resolution maps from wide angle sonar," in *Proc. IEEE Int. Conf. Robotics and Automation*, Mar. 1985, pp. 116-121.
- [17] R. C. Nelson and H. Samet, "A consistent hierarchical representation for vector data," in *Proc. SIGGRAPH'86 Conf.*, Dallas, Aug. 1986, pp. 197-206.
- [18] W. M. Newman and R. F. Sproull, *Principles of Interactive Computer Graphics*, 2nd ed. New York: McGraw-Hill, 1979.
- [19] N. J. Nilsson, *Problem-Solving Methods in Artificial Intelligence*. New York: McGraw-Hill, 1971.
- [20] D. J. Orser and M. Roche, "The extraction of topographic features in support of autonomous underwater vehicle navigation," *Fifth Int. Symp. Unmanned Unethered Submersible Tech.*, Jun. 1987, vol. 2, pp. 502-514.
- [21] R. J. Renka and A. K. Cline, "A triangle-based C^1 interpolation method," *Rocky Mountain J. Mathematics*, vol. 14, Winter 1984, pp. 223-237.
- [22] M. C. Roggermann, J. P. Mills, S. K. Rogers, and M. Kabrisky, "Multisensor information fusion for target detection and classification," *Proc. SPIE*, Apr. 1988, pp. 8-13.
- [23] H. Samet, "The quadtree and related hierarchical data structures," *ACM Computing Surveys*, vol. 16, Jun. 1984, pp. 187-260.
- [24] H. Samet and R. E. Webber, "Storing a collection of polygons using quadtrees," *ACM Trans. Graphics* 4, July 1985, pp. 182-222.
- [25] H. Samet, A. Rosenfeld, C. A. Shaffer, and R. E. Webber, "A geographic information system using quadtrees," *Pattern Recognition*, Nov./Dec., 1984, pp. 647-656.
- [26] C. A. Shaffer and H. Samet, "Optimal quadtree construction algorithms," *Computer Vision, Graphics, and Image Processing*, vol. 37, Mar. 1987, pp. 402-419.
- [27] W. K. Stewart, "A model-based approach to 3-D imaging and mapping underwater," in *Proc. ASME Conf. Offshore Mechanics and Arctic Eng.*, Houston, TX, Feb. 1988.

- [28] C. W. Tong, S. K. Rogers, J. P. Mills, and M. K. Kabrisky, "Multisensor data fusion of laser radar and forward looking infrared (FLIR) for target segmentation and enhancement," *Proc. SPIE*, May 1987, pp. 10-19.
- [29] M. A. Turk, D. G. Morgenthaler, K. D. Gremban, and M. Marra, "Video road-following for the autonomous land vehicle," in *Proc. IEEE Int. Conf. Robotics Automat.*, Mar. 1987, pp. 273-279.

Modeling and Control of Underwater Robotic Vehicles

J. YUH

Abstract—Remotely operated, underwater robotic vehicles have become the important tool to explore the secrete life undersea. They are used for various purposes: inspection, recovery, construction, etc. With the increased utilization of remotely operated vehicles in subsea applications, the development of autonomous vehicles becomes highly desirable to enhance operator efficiency. However, engineering problems associated with the high density, nonuniform and unstructured seawater environment, and the nonlinear response of the vehicle make a high degree of autonomy difficult to achieve. The dynamic model of the untethered vehicle is presented, and an adaptive control strategy for such vehicles is described. The robustness of the control system with respect to the nonlinear dynamic behavior and parameter uncertainties is investigated by computer simulation. The results show that the use of the adaptive control system can provide the high performance of the vehicle in the presence of unpredictable changes in the dynamics of the vehicle and its environment.

I. INTRODUCTION

A large portion of the earth is covered by seawater and has not been fully explored, so plenty of resources still remain in a natural condition. In a recent report [1] to the National Science Foundation, seven critical areas in ocean system engineering were identified as follows: system for characterization of the sea bottom resources; systems for characterization of the water column resources; waste management systems; transport, power and communication systems; reliability of ocean systems; materials in the ocean environment; analysis and application of ocean data to develop ocean resources. It was also concluded in the report that the area of underwater robotics should be supported in all of the above areas. It is obvious that all kinds of ocean activities, including both scientific ocean related research, and commercial utilization of ocean resources, will be greatly enhanced by the development of an intelligent, robotic underwater work system.

Current underwater working methods include scuba, remotely operated vehicle (ROV), submarine, etc. During the last few years, the use of ROVs has rapidly increased since such a vehicle can be operated in the deeper and riskier areas where divers cannot reach. In the undersea environment, ROVs are used for various work assignments. Among them are: pipelining, inspection, data collection, construction, maintenance and re-

Manuscript received March 6, 1989; revised February 22, 1990. This paper includes part of the results of the research being conducted at the Robotics Laboratory of the University of Hawaii. Funding for this research ("Automatic optical stationkeeping for a subsea remotely operated vehicle," Project R/OE/11) was provided in part by UH Sea Grant College Program under Institutional Grant No. NA85AA-D-SG082 from NOAA Office of Sea Grant, Department of Commerce, and in part by the Pacific International Center for High Technology Research under Grant No. PO-6116003. This is a Sea Grant publication UNHI-SEAGRANT-JC-90-15.

The author is with the Department of Mechanical Engineering, University of Hawaii, Honolulu, Hawaii 96822.
IEEE Log Number 9037609.