A COMPARATIVE STUDY ON THE PROCESSING OF ULTRASONIC ARC MAPS

Billur Barshan

Department of Electrical and Electronics Engineering, Bilkent University
Bilkent TR-06800 Ankara, Turkey
phone: + (90) 312-290-2161, fax: + (90) 312-266-4192, email: billur@ee.bilkent.edu.tr
web: www.ee.bilkent.edu.tr/∼ billur

ABSTRACT

The directional maximum (DM) technique for processing ultrasonic arc maps is proposed and compared to previously existing techniques. The method processes ultrasonic arc maps directionally to extract the map of the environment and overcome the intrinsic angular uncertainty of ultrasonic sensors. It also eliminates noise and cross-talk related misreadings successfully. The comparison is based on experimental data and three complementary error criteria. The DM technique offers a very good compromise between mean absolute error and correct detection rate, with a processing time less than tenth of a second. It is superior to existing techniques in range accuracy and in eliminating artifacts, resulting in the best overall performance. The results indicate several trade-offs in the choice of ultrasonic arc-map processing techniques.

1. INTRODUCTION

Sensing and becoming aware of their environment is an essential feature of intelligent systems. This awareness can be realized by processing the sensory signals acquired from simple sensors according their perceptive needs. Due to the limited resources of autonomous systems, the available resources need to be maximally exploited. Before more expensive sensing modalities with higher resolution and resource requirements are considered for a given task, it makes sense to first exploit the potential of simple and inexpensive sensors to extract information about the environment as much as possible. One of the aims of this study is to explore the limits of simple and low-cost ultrasonic sensors through intelligent processing of their raw echo signals. Ultrasonic sensors have been widely used in intelligent systems due to their accurate range measurements, robustness, low cost, and simple hardware interface. When coupled with intelligent processing, they provide a useful alternative to more complex laser and camera systems. Furthermore, it may not be possible to use the latter in some environments due to surface characteristics or insufficient ambient light. Despite their advantages, the frequency range at which air-borne ultrasonic transducers operate is associated with a large beamwidth that results in low angular resolution and uncertainty in the location of the echo-producing object. Thus, having an intrinsic uncertainty of the actual angular direction of the range measurement and being prone to various phenomena such as multiple and higher-order reflections and cross-talk between transducers, a considerable amount of modeling, processing, and interpretation of ultrasonic data is necessary.

This paper provides a valuable comparison between the performances of one newly proposed and six existing techniques for processing ultrasonic arc maps (UAMs). We have also modified one of the existing techniques and included it in the comparison. The comparison is based on experimentally acquired signals and three complementary error criteria. The newly proposed DM technique introduces a sense of direction in sensor data processing and offers a very good compromise between two of the error criteria, higher range accuracy, and effective elimination of spurious arcs, resulting in the best overall performance. Some of the existing techniques also have certain advantages that make them suitable for mapping under different constraints and requirements.

2. REPRESENTING ANGULAR UNCERTAINTY BY UAMS

Simple ultrasonic range sensors are employed that estimate the time-of-flight (TOF), which is the round-trip travel time between the transducer and the object. The received echo is usually contaminated by noise and the time at which the reflection is received can be estimated by means of simple thresholding. Alternatives to simple thresholding such as adaptive, variable or double thresholding [1] and curve-fitting [2] techniques have been proposed. Once the TOF value is estimated, the range is calculated from $r = \frac{c\,t_\circ}{2}$, where $t_\circ$ represents the TOF and c is the speed of sound.

Although ultrasonic sensors return accurate range data, typically they cannot provide direct information on the angular position of the object from which the reflection was obtained. Most commonly, the large beamwidth of the transducer is accepted as a device limitation that determines the angular resolving power of the system, and the reflection point is assumed to be along the line-of-sight (LOS) of the transducer. According to this naive approach, a simple mark is placed along the LOS at the measured range, resulting in inaccurate maps with large angular errors and artifacts. This will be referred as the point marking (PM) method below. In specularly reflecting environments, the accumulation of such marks usually produces arc-like features called *regions of constant depth* (RCDs) [3]. Other researchers have attempted to fit line segments [4, 5] and use Hough transforms [6] for both air-borne and underwater ultrasonic data that proved to be unsuccessful.

As a more general approach which is not limited to specularly reflecting surfaces, the angular uncertainty in the range measurements can be represented by UAMs [7] that preserve more information (see Fig.1(b) for a sample UAM). Construction of UAMs is more general in that they can be generated for environments comprised of both specularly and diffusely reflecting (Lambertian) surfaces. However, RCDs are limited to specularly reflecting environments only. The arcs comprising the UAM are different than the arcs corresponding to RCDs. The UAMs are obtained by drawing arcs spanning the beamwidth of the sensor at the measured range, representing the angular uncertainty of the object location and indicating that the echo-producing object can lie anywhere on the arc. Thus, when the same transducer transmits and receives, all that is known is that the reflection point lies on a circular arc of radius r. More generally, when one transducer transmits and another receives, it is known that the reflection point lies on the arc of an ellipse whose focal points are the transmitting and receiving elements. The arcs are tangent to the reflecting surface at the actual point(s) of reflection. Arc segments near the actual reflection points tend to reinforce each other. Those not actually corresponding to any reflections and simply representing the angular uncertainty of the transducers remain more sparse and isolated. Arcs generated by spurious readings, cross-talk, higher-order reflections, and noise also remain sparse and lack reinforcement. These are not enhanced as much as the arcs resulting from actual reflections. The proposed technique effectively suppresses these effects, and, although not implemented here, it has the intrinsic ability to process echoes returning from surface features further away than the nearest (i.e., multiple reflections) informatively. By combining the information inherent in a

large number of such arcs, much improved angular resolution is obtained.

3. UAM PROCESSING TECHNIQUES

3.1 Point Marking (PM)

This is the simplest approach, mentioned above, where a mark is placed along the LOS at the measured range. This method produces reasonable estimates for the locations of objects if the arc of the cone is small. This can be the case at higher frequencies of operation where the corresponding sensor beamwidth is small, or at nearby ranges. Since every arc is reduced to a single point, this technique cannot eliminate any of the outlying TOF readings. The resulting map is inaccurate with large angular errors and artifacts.

3.2 Voting and Thresholding (VT)

Another simple technique for processing UAMs is a voting scheme where each pixel stores the number of arcs crossing that pixel, resulting in a 2-D array of occupancy counts for the pixels [8]. By simply thresholding this array and zeroing the pixels lower than the threshold, artifacts can be eliminated and the map is extracted.

3.3 Directional Maximum (DM)

This technique is based on the idea that in processing the acquired range data, there is a direction-of-interest (DOI) associated with each detected echo. Ideally, the DOI corresponds to the direction of a perpendicular line drawn from the sensor to the nearest surface from which an echo is detected. However, in practice, due to the angular uncertainty of the object position, the DOI can be approximated as the LOS of the sensor when an echo is detected. Since prior information on the environment is usually unavailable, the DOI needs to be updated while sensory data are being collected and processed on-line.

In the implementation, the number of arcs crossing each pixel of the UAM is counted and stored, and a suitable threshold value is chosen. Up to this point, the DM method is exactly the same as the VT method. The novelty of the DM method is the processing done along the DOI. Once the DOI for a measurement is determined using a suitable procedure, the UAM is processed along this DOI as follows: The array of pixels along the DOI is inspected and the pixel(s) exceeding the threshold with the maximum count is kept, while the remaining pixels along the DOI are zeroed out. If there exist more than one maxima, the algorithm takes their median (If the number of maxima is odd, the maxima in the middle is taken; if the number is even, one of the two middle maxima is randomly selected.) This way, most of the artifacts of the UAM can be removed.

3.4 Morphological Processing (MP)

The processing of UAMs using morphological operators was first proposed in [7]. This approach exploits neighboring relationships and provides an easy to implement yet effective solution to ultrasonic map building. By applying binary morphological processing operations, one can eliminate the artifacts of the UAM and extract the surface profile.

In MP, we applied the thinning operation with parameter m to the UAM. The thinning parameter is varied between 0 and 8 ($m = 1$: pruning, $m = 8$: erosion, $m = 0$: no effect). Only the pixels with value 1 are considered, indicating that there exists at least one arc crossing that pixel. The number of nonzero neighbors of each pixel is counted. If the number of nonzero neighbors of a pixel is less than m, this means that the pixel does not have sufficient support and its value is equated to zero; otherwise, its value remains as 1.

3.5 Bayesian Update Scheme for Occupancy Grids (BU)

Occupancy grids were first introduced by Elfes, and a Bayesian scheme for updating their probabilities of occupancy and emptiness was proposed in [9] and verified by ultrasonic data. Starting with a blank or completely uncertain occupancy grid, each range measurement updates the grid formation in a Bayesian manner. In the following, r is the range measurement, $r_{\min}$ is the near-field limit, r_ε is the maximum range measurement error, and $\theta_\circ$ is sensor half-beamwidth angle. The distance between the ultrasonic sensor at $S = (x_s, y_s)$ to the pixel $P(x,y)$ is σ. The p.d.f. for the emptiness of a pixel is given by:

$$p_E(x,y) = p[P(x,y) \text{ is empty}] = E_r(\sigma) \cdot E_a(\theta) \tag{1}$$

where $E_r(\sigma)$ and $E_a(\theta)$ are functions that model the range and angle dependence of $p_E(x,y)$ as follows:

$$E_r(\sigma) = \begin{cases} 1 - [\frac{\sigma - r_{\min}}{r - r_\varepsilon - r_{\min}}]^2 & \text{for } \sigma \in [r_{\min}, r - r_\varepsilon] \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

$$E_a(\theta) = 1 - \left(\frac{\theta}{\theta_\circ}\right)^2, \text{ for } \theta \in [-\theta_\circ, \theta_\circ]. \tag{3}$$

Likewise, the p.d.f. for occupancy is:

$$p_O(x,y) = p[P(x,y) \text{ is occupied}] = O_r(\sigma) \cdot O_a(\theta) \tag{4}$$

where $O_r(\sigma)$ and $O_a(\theta)$ are functions that model the range and angle dependence of $p_O(x,y)$ as:

$$O_r(\sigma) = \begin{cases} 1 - [\frac{\sigma - r}{r_\varepsilon}]^2, & \text{for } \sigma \in [r - r_\varepsilon, r + r_\varepsilon] \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

$$O_a(\theta) = 1 - \left(\frac{\theta}{\theta_\circ}\right)^2, \text{ for } \theta \in [-\theta_\circ, \theta_\circ]. \tag{6}$$

Initially, the map-building process starts with a maximally uncertain map where all the pixel values are set to zero, corresponding to the mid-point of the interval $[-1, 1]$ for pixel values. For each range reading r, p_E values are updated for those pixels within the sensitivity region of the transducer that fall in the range interval $[r_{\min}, r - r_\varepsilon]$ using Eqns.(1)–(3). Similarly, p_O values are calculated for those pixels within the sensitivity region that fall in the range interval $[r - r_\varepsilon, r + r_\varepsilon]$, using Eqns.(4)–(6). The following BU rules are employed to update the existing values in the pixel array:

$$\text{updated value of } p_E(pixel) = p_E(pixel) + p_E(reading) - p_E(pixel) \times p_E(reading) \tag{7}$$

$$\text{updated value of } p_O(pixel) = p_O(pixel) + p_O(reading) - p_O(pixel) \times p_O(reading) \tag{8}$$

The map is constructed by iteratively updating the contents of each pixel using all the available range measurements comprising the UAM. In the end, p_E and p_O arrays contain modified p.d.f.s whose values vary between –1 and 1. These are then transformed to the interval $[0, 1]$ and thresholded with a value close to 1 (≥ 0.997) to get the map.

3.6 Arc-Transversal Median (ATM)

The ATM algorithm requires both extensive bookkeeping and considerable amount of processing [10]. For each arc in the UAM, the positions of the intersection(s) with other arcs, if they exist, are recorded. For arcs without any intersections, the mid-point of the arc is taken to represent the actual point of reflection (as in PM). If the arc has a single intersection, the algorithm uses the intersection point as the location of the reflecting object. For arcs with more intersections, the median of the positions of the intersection points with other arcs is chosen to represent the actual point of reflection. In [10], the median operation is applied when an arc has *three or more* intersection points. If there is an even number of intersections, the algorithm uses the mean of the two middle values (except that arcs with two intersections are ignored). It can be considered as a much improved version of the PM approach.

To improve the performance of this technique in eliminating the artifacts, we have also implemented a modified version of the algorithm (ATM-mod) where we ignored arcs with no intersections. Furthermore, since we could not see any reason why arcs with two intersections should not be considered, we took the mean of the two intersection points.

3.7 Triangulation-Based Fusion (TBF)

The TBF method is primarily developed for accurately detecting the edge-like features in the environment based on triangulation [11]. The triangulation equations involved are not suitable for accurately localizing planar walls. All of the methods compared in this study aim at extracting the strongly supported, darker features in the UAM. A distinguishing feature of TBF is that it realizes this by using a geometric model suitable for edge-like features, whereas the previously described techniques accomplish this by dividing the environment into grids and processing the information in each grid.

Unlike the previously introduced grid-based techniques, the TBF method extracts the features of the environment by using a geometric model suitable for edge-like features. In addition, TBF considers a sliding window of sonar scans where the number of rows of the sliding window corresponds to the number of sonars fired, and the number of columns corresponds to the number of most recent sonar scans to be processed by the algorithm. TBF is focused on detection of edge-like features located at ≤ 5 m. The other methods consider all of the arcs in the UAM corresponding to all ranges, and are suitable for detecting all types of features.

We have implemented the TBF method as described in [11] with exactly the same parameters except that we used sliding windows of dimension 3×20 or 5×20 instead of 16×10. This is because our ultrasonic data were collected every 2.5 cm (which is half of what is stated in [11]) using three or five ultrasonic transducers.

An important parameter of the TBF algorithm is the number n_t of successful triangulations performed for a given range reading. If the maximum deviation between successful triangulations associated with a range reading is larger than a preset threshold, the triangulation point is classified as belonging to an object which is not well-represented by an edge, for instance a wall. At the end of the algorithm, features which are not edge-like may or may not be shown in the map. In our comparison, if we extracted and showed only the edge-like features, the resulting map would have been very sparse. To improve the performance of this technique, we have also considered including all the successful triangulation points. In addition, we have varied the threshold value for n_t from 1 to 8, considered sliding window column sizes of 10 and 20, and presented the best results possible.

3.8 Error Criteria

We defined and used three complementary error criteria to evaluate and compare the techniques described above. The first and maybe the more important one is the mean absolute error (MAE) between the true and extracted features of the environment. To calculate the MAE, the distance of each nonzero point in the resulting map to the nearest point in the structured-light profile (taken as absolute reference) is found and these distances are averaged over all the nonzero points in the constructed map. The second error criterion is the correct detection rate (CDR). In evaluating the CDR, first, the full and empty directions of the actual map are considered separately. Then, combining these, an overall CDR is calculated, given by:

$$\mathrm{CDR_F} = \frac{\text{correctly detected full}}{\text{total number of full}} \tag{9}$$

$$\mathrm{CDR_E} = \frac{\text{correctly detected empty}}{\text{total number of empty}} \tag{10}$$

$$\mathrm{CDR_O} = \frac{\text{correctly detected full and empty}}{\text{total number of full and empty}} \tag{11}$$

We have evaluated the CDR using the concept of DOI described above, instead of evaluating it pixel by pixel. For a given DOI, if the actual profile revealed by the structured-light data indicates that that direction is full (empty), and the resultant map agrees with this, correct detection of a full (empty) DOI occurs.

The final criterion stands for the computational cost of the techniques in terms of the CPU time t_{CPU}. These error criteria, one standing for quality, the second for quantity, and the last for implementability are more meaningful when considered together.

4. EXPERIMENTAL VERIFICATION

The techniques briefly described above are experimentally verified using the ultrasonic and structured-light sensing systems on the Nomad 200 mobile robot in our laboratory [12]. The structured-light system is much more expensive and complex, requiring higher-power and sufficient ambient light for operation. Furthermore, this mode of sensing does not work in all environments, such as those with dark-colored upholstery and glass. Since it reveals a very accurate surface profile, the profile detected by this system is taken as an absolute reference in the experimental calculation of the MAE and CDR using ultrasonic data.

The device that is used in this study is the Polaroid 6500 series transducer [13] operating at a resonance frequency of $f_\circ = 49.4$ kHz. In the experiments, the ultrasonic transducers are fired at 40 ms intervals, which is sufficient to eliminate most of the cross-talk.

4.1 Experimental Results

We have conducted experiments with curved surfaces as well as indoor environments comprised of basic target types. In each experiment, the same UAM is processed by each technique to provide a uniform comparison. The results of the experiments with curved surfaces can be found in [14].

Here, we present experimental results obtained by using the front three ultrasonic sensors of the Nomad 200 robot, following the walls of the indoor environment in Fig.1(a). The environment comprises of smooth wooden (top and left) and painted (right) walls, and a window shade with vertical slats of 15 cm width (bottom). Some of the corners of the room were not perfect (e.g., where the shade and the right wall make a corner). There is a cylinder of radius 15 cm at a distance of 30 cm from the center of the right wall. A total of 702 arcs are obtained, 659 of which fall in the 525×525 region (Fig.1(b)). The UAM includes many artifacts, especially exterior to the surrounded region where the robot can move. The results of the different processing techniques are shown in Fig.1(c)-(j). In this experiment, we took the DOI as the direction of the currently followed wall. For VT and DM, the variable parameter is the pixel threshold value, for MP, the thinning parameter m, for BU, the threshold used for probability of occupancy, and for TBF, the number n_t of successful triangulations associated with a range reading. Results of varying the parameters of VT, DM, MP, BU, and TBF techniques can be found in [14]. Suitable parameter values are selected such that the corresponding $\mathrm{CDR_F}$ and $\mathrm{CDR_E}$ are as close to each other as possible. This is needed to eliminate any tendency to over-detect pixels as full or empty and to ensure that full and empty regions are detected with approximately the same correct detection rate. Using the most suitable parameter values, a comparison of the seven UAM processing techniques is provided in Table 1 for one of the experiments. ATM-mod produces the smallest error but has the second lowest CDR. The maps extracted with this technique are quite accurate but sparse. The errors of DM and VT are comparable to that of ATM-mod, followed by TBF, ATM-org, MP, BU, and PM. The larger MAEs usually result from artifacts that a certain technique was unable to remove. The DM and ATM-mod techniques are the best in eliminating artifacts, followed by VT, MP, and TBF. The processing results of PM, BU, and ATM-org still contain artifacts. The highest overall CDR is obtained by DM and VT. The CDR of MP is also quite high, followed by BU and PM. The CDRs of PM and ATM-org are comparable and those of ATM-mod and TBF are the lowest among all seven techniques. The results of MP given in Fig.1(f) correspond to applying a single thinning operation of order $m = 7$. Artifacts have been removed to some extent and planar surfaces are satisfactorily represented. The MP results have larger range uncertainty indicated by thick solid line features at planar surfaces, while corners are too much eroded. Furthermore, a substantial amount of arc branches remain and applying further thinning operations not only removes these unwanted branches but also reduces the CDR considerably, while increasing CPU time.

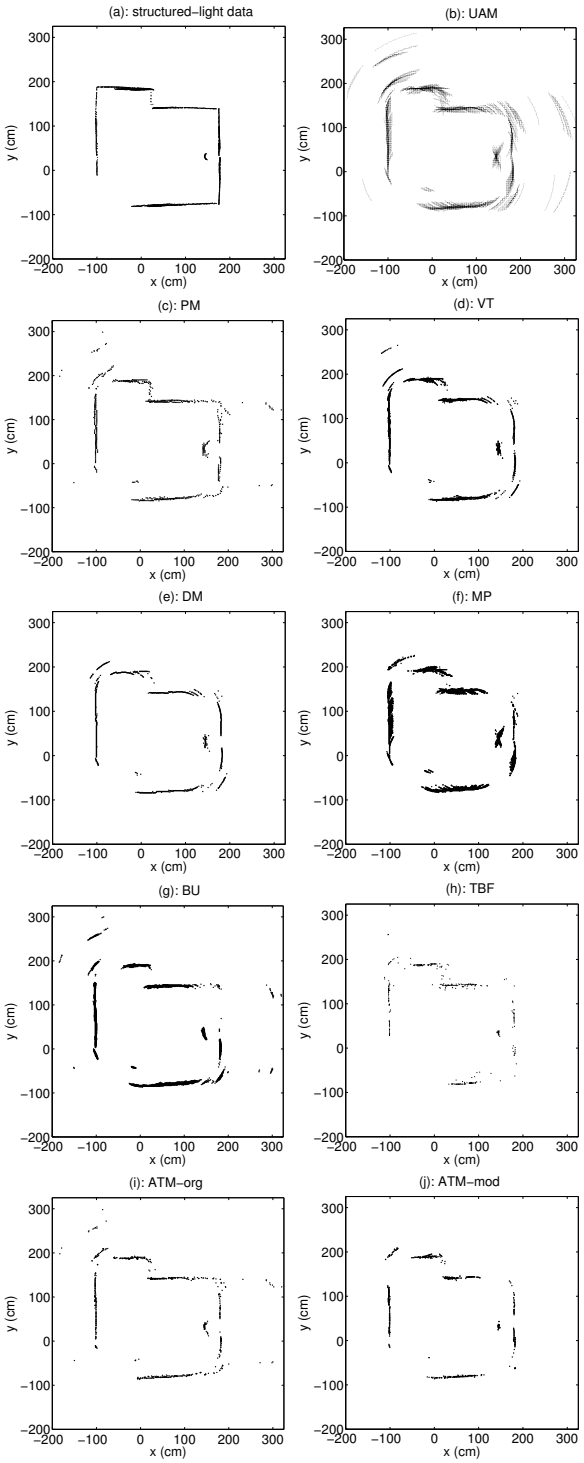

Figure 1: (a) The map acquired with structured-light system, (b) UAM; processing results of (c) PM, (d) VT, (e) DM, (f) MP, (g) BU, (h) TBF, (i) ATM-org, (j) ATM-mod.

From the given resulting maps, DM and ATM techniques are superior in terms of range uncertainty, indicated by thinner borders and more accurately placed features, as well as their lower MAE. Artifacts are also still present in the results of BU given in Fig. 1(g). Thresholding the probability of occupancy with smaller values is found to result in more artifacts, while larger thresholds reduce the CDRs to too low values.

Table 1: Experimental results for the indoor environment.

method	MAE (cm)	CDR_O	CDR_F	CDR_E	t_{CPU} (s)
PM	6.06	0.631	0.579	0.955	0.001
VT (thld=5)	2.63	0.883	0.881	0.896	0.074
DM (thld=5)	2.37	0.883	0.881	0.896	0.078
MP ($m = 7$)	4.89	0.820	0.824	0.799	0.082
BU (thld=0.998)	5.20	0.773	0.768	0.799	0.563
TBF ($n_t \geq 2$)	2.70	0.414	0.320	1.000	0.010
ATM-org	4.78	0.606	0.549	0.955	1.054
ATM-mod	1.68	0.522	0.449	0.947	0.871

4.2 Discussion

In general, if a processing technique cannot eliminate artifacts well, the resulting MAE is larger. DM, VT, and ATM-mod are superior to the other techniques in eliminating artifacts, therefore they result in smaller errors. Note that after thresholding but before applying any directional processing in DM, the VT and DM pixel values are exactly equal if the same threshold value is used in both algorithms. If there are multiple maxima along a given DOI, DM directionally processes and selects one of them, whereas VT keeps all. For this reason, DM usually performs better than VT in terms of MAE, especially at smaller values of the threshold between 1 and 4. At larger values of the threshold, they perform comparably. The CDRs of VT and DM are always equal because if a given DOI is empty, it will remain as empty after directional processing; if it is full, it will be still full after the directional maximum is taken along that direction.

ATM-mode is also quite accurate and eliminates the artifacts better than ATM-org. However, comparing Figs.1(i) and (j), there are more gaps in part (j), especially around corners and edges. This is because in those regions of the UAM, there were arcs with no intersections that were removed with ATM-mod but not removed by ATM-org. Generally speaking, ATM technique creates accurate yet sparsely filled maps with a tendency towards under-filling. Since the same number of arcs are processed by each technique, ATM is found to be requiring a higher number of arcs in order to produce a map with similar CDR as the other techniques.

The PM technique reduces each arc to a single point in the middle of the arc. ATM-org places a more accurately placed point mark on arcs with transversal intersections (except those with 2), reducing many of the arcs to single points also. It treats arcs with no intersections in the same way as PM. For this reason, the CDR of this technique is quite close but a little lower than that of PM.

Edge locations obtained with TBF are very accurate as expected. The accuracy of TBF usually falls between ATM-org and ATM-mod. However, due to the large number of gaps in the resulting map, the CDR obtained with TBF is the lowest among all the techniques compared. This is expected because apart from the fact that a smaller number of arcs is used at a given time to begin with (due to the sliding window), TBF eliminates those arcs without any meaningful and accurate correspondence. In addition, planar wall locations are not very accurate. On the other hand, a major advantage of TBF is that it is very fast for the given window sizes and takes about the same time as the simplest PM method. This is because it does not divide the environment into grids and processes the information geometrically, instead of grid by grid.

Among the seven approaches considered, DM produces one of the lowest MAE and CPU time. Its performance according to the CDR criterion is also very good, where VT and DM usually rank the best in different examples. Considering that DM also has very high range accuracy and is superior in eliminating artifacts and outliers of the UAM, it can be concluded that it results in the best overall performance.

In the works [15, 16], VT and MP were investigated in detail based on simulations and experimental studies for different transducer configurations (linear, circular, random), different beamwidths (5° to 105°), different surface curvature, roughness,

distance, and different noise levels on the TOF measurements. The best results were obtained with a random configuration of transducers, followed by circular and linear ones. For both methods, the errors were shown to increase with increasing beamwidth, increasing surface distance, curvature, and roughness. Although such detailed studies for the other methods have not been performed, we expect similar results for the remaining techniques. This is because the impact of varying these parameters is primarily to affect the quality of the information inherent in the ultrasonic arc map. This also leads us to expect that for a given choice of these parameters, the comparison of the methods will not be altered significantly.

4.3 Computational Cost of the Techniques

The average CPU times are of the order of fractions of a second, indicating that processing methods are viable for real-time applications. These represent the total time the computer takes to realize the processing techniques starting with the raw TOF data (i.e., the UAM). (The processing techniques have been implemented in the C language and run on a Intel Pentium 4 PC with 3.00 GHz Hypertreading processor and 1 GB memory. Internal C commands are used for time keeping.)

PM being the simplest technique, it has the smallest processing time as expected (less than hundredth of a second). It is closely followed by TBF. The CPU times of VT, DM, and MP are comparable to each other and slightly less than tenth of a second in the last two experiments. CPU times of BU and ATM are larger but still suitable for real-time applications. Since the total CPU time depends on the total number of arcs in the UAM, dimensions of the environment, and the grid size, CPU time per arc is a better indicator of the computational load of the processing techniques. For the BU method, the CPU time per arc is less than 1 ms, whereas for the ATM method, the CPU time per arc is about 1.5 ms. The CPU times of original and modified ATM are comparable.

The computational complexity of TBF highly depends on the calculation of the triangulation points and their optional refinement using a local grid map. It is also dependent on the dimensions of the sliding window. Since we omitted the optional refinement step and used a relatively small window size, computational complexity was not high in our implementation. Inclusion of the refinement step would considerably increase the cost but further improve the accuracy of the triangulation points, consequently reducing the MAE.

For comparison, the time it takes for an array of 16 ultrasonic sensors to collect TOF data is 16×40 ms = 0.64 s which is of the same order of magnitude as the processing time for BU. The actual algorithmic processing time is a small fraction of the CPU time, as most of the CPU time is consumed by file operations, reads and writes to disk, matrix allocations etc. Thus, it seems possible that a dedicated system can perform the mapping even faster.

Among the techniques considered, ATM is the one that requires the largest amount of bookkeeping and storage. PT and BU can be implemented on-line while data are still being acquired by the mobile robot. The thresholding part of VT, thresholding and directional processing of DM, MP, and ATM work better if employed off-line, after the complete UAM becomes available. However, it may be possible to develop regional or on-line versions of these techniques.

5. CONCLUSIONS AND FUTURE WORK

A valuable comparison between the DM technique and existing techniques for processing UAMs is provided. The results indicate that DM technique has some advantages over existing techniques: having a small MAE, one of the largest CDR and a CPU time less than tenth of a second, it provides a good compromise between the MAE and CDR criteria in many cases. It also has low range uncertainty and is superior in eliminating artifacts. Some of the existing techniques also have characteristics that may be suitable for certain situations or conditions. For example, if a large number of ultrasonic arcs is conveniently available and processing time and memory is not a major issue, ATM algorithm would be a good choice. On the other hand, if simple and fast on-line processing is desirable, and high accuracy is not crucial, one could use the PM method, as done in many cases. TBF is more suitable for environments in which many edge-like features are present and extracts the positions of these features quite accurately. It is also very fast for moderate sliding window sizes.

Our current work is focused on combining two of the error criteria (MAE and CDR) used in this paper into a single criterion by fitting active snake contours and Kohonen's self-organizing maps to the map points for their efficient and compact representation. Such efficient representation makes it easier to compare the performances of different techniques. In future work, 3-D versions of the processing techniques and their compact representation may be developed. Even though the results demonstrated in this paper are based on processing ultrasonic echo signals, the techniques presented and compared here can be conveniently extended to other sensing modalities such as radar, laser, and infrared.

Acknowledgement: This work is supported in part by TÜBİTAK under grant number EEEAG-105E065.

REFERENCES

[1] W. G. McMullan, B. A. Delanghe, and J. S. Bird, A simple rising-edge detector for time-of-arrival estimation, *IEEE Trans. Instr. Meas.*, 45(4):823–827, Aug. 1996.

[2] B. Barshan and R. Kuc, A bat-like sonar system for obstacle localization, *IEEE Trans. Syst. Man Cybern.*, 22(4):636–646, July/Aug. 1992.

[3] R. Kuc and M. W. Siegel, "Physically-based simulation model for acoustic sensor robot navigation," *IEEE Trans. Pattern Anal. Mach. Intell.*, PAMI-9(6):766–778, Nov. 1987.

[4] M. Drumheller, "Mobile robot localization using sonar," *IEEE Trans. Pattern Anal. Mach. Intell.*, PAMI-9(2):325–332, March 1987.

[5] J. L. Crowley, "Navigation for an intelligent mobile robot," *IEEE Trans. Robot. Automat.*, RA-1(1):31–41, March 1985.

[6] D. Ribas, P. Ridao, J. Neira and J. D. Tardós, "Line extraction from mechanically scanned imaging sonar," *Pattern Recogn. Image Anal., LNCS*, 4477:322–329, July 2007.

[7] D. Başkent and B. Barshan, "Surface profile determination from multiple sonar data using morphological processing," *Int. J. Robot. Res.*, 18(8):788–808, Aug. 1999.

[8] B. Barshan, "Ultrasonic surface profile determination by spatial voting," *Elect. Lett.*, 35(25):2232–2234, Dec. 1999.

[9] A. Elfes, "Sonar based real-world mapping and navigation," *IEEE Trans. Robot. Automat.*, RA-3(3):249–265, June 1987.

[10] H. Choset, K. Nagatani, and N. Lazar, "The arc-transversal median algorithm: a geometric approach to increasing ultrasonic sensor azimuth accuracy," *IEEE Trans. Robot. Automat.*, 19(3):513–522, June 2003.

[11] O. Wijk and H. I. Christensen, "Triangulation-based fusion of sonar data with application in robot pose tracking," *IEEE Trans. Robot. Automat.*, 16(6):740–752, Dec. 2000.

[12] Nomadic Tech., Inc., Mountain View, CA, *Nomad 200 Manual*, 1997.

[13] Polaroid Corp., Ultrasonic Components Group, 119 Windsor St., Cambridge, MA, *Polaroid Manual*, 1997.

[14] B. Barshan, "Directional processing of ultrasonic arc maps and its comparison with existing techniques," *Int. J. Robot. Res.*, 26(8):797–820, Aug. 2007.

[15] B. Barshan and D. Başkent, Morphological surface profile extraction with multiple range sensors, *Pattern Recogn.*, 34(7):1459–1467, July 2001.

[16] B. Barshan and D. Başkent, Comparison of two methods of surface profile extraction from multiple ultrasonic range measurements, *Meas. Sci. Tech.*, 11(6):833–844, June 2000.