IEEE Instrumentation and Measurement
Technology Conference
Budapest, Hungary, May 21–23, 2001

Ultrasonic Surface Profile Determination by Spatial Voting

Billur Barshan
Department of Electrical Engineering
Bilkent University, TR-06533 Ankara, Turkey
Phone: +90 312 290-2161, Fax: +90 312 266-4192
Email: billur@ee.bilkent.edu.tr, URL: http://www.ee.bilkent.edu.tr/~ billur

Abstract – A novel spatial voting scheme is described for surface profile determination based on multiple ultrasonic range measurements. Spatial voting relies on the number of votes accumulated in each pixel of the ultrasonic arc map but ignores neighboring relationships. This approach is extremely robust, flexible, and straightforward. It can deal with arbitrary numbers and configurations of sensors as well as synthetic arrays, with the intrinsic ability to suppress spurious readings, crosstalk, and higher-order reflections, and process multiple reflections informatively. The performance of the method is investigated on various examples involving both simulated and experimental data. The effect of varying the surface roughness is also considered.

Keywords – Ultrasonic range measurement, Sonar, Voting, Feature extraction, Surface profile determination, Map building.

I. INTRODUCTION

An inexpensive, yet effective and reliable approach to machine perception is to employ multiple simple range sensors coupled with suitable data processing. Geometrical or analytical modeling approaches to map building are often limited to elementary target types with constant or piecewise-constant curvature, or simple sensor configurations [1], [2]. A novel spatial voting scheme is applied to ultrasonic range measurements to reconstruct the profiles of *arbitrarily* curved surfaces that are encountered in unstructured environments such as mines, rough terrain, or underwater. Although voting strategies have been previously used in other areas [3], [4], [5], this is their first application to ultrasonic surface profile extraction. The method presented is extremely flexible and can easily handle arbitrary sensor configurations as well as synthetic arrays obtained by moving a relatively small number of sensors.

Ultrasonic sensors have been widely used in intelligent systems since they are robust, light, and inexpensive. They are able to reveal a great amount of useful information when coupled with appropriate data processing and interpretation. A commonly noted disadvantage of ultrasonic ranging systems is the difficulty associated with interpreting spurious readings, crosstalk, higher-order, and multiple reflections. The proposed method is capable of effectively suppressing the first three of these, and has the intrinsic ability to process echoes returning from surface features further away than the nearest (i.e. multiple reflections) informatively.

This work was supported by TÜBİTAK under grants 197E051, EEEAG-116, and EEEAG-92.

0-7803-6646-8/01/$10.00 ©2001 IEEE

Spatial voting is a novel way of processing range data in the form of an arc map which represents angular uncertainties. The method can also be viewed as a new way of solving a class of nonlinear reconstruction problems that arise when a large number of sensors produce range measurements. These nonlinear inverse problems do not seem amenable to efficient solution by standard analytical or numerical techniques.

It is important to underline that spatial voting is employed here to process the ultrasonic arc map of the surface being reconstructed, rather than conventional camera images. The use of multiple ultrasonic range measurements combined with spatial voting and thresholding, can be applied to different physical modalities of range sensing of vastly different scales and in many different areas. These may include radar, sonar, optical sensing and metrology, remote sensing, ocean surface exploration, geophysical exploration, robotics, and acoustic microscopy. Although the present paper deals with the determination of two-dimensional surface profiles, the method can be readily generalized to three-dimensional environments with the arcs replaced by spherical or elliptical caps and the arc map processing rules extended to three dimensions.

II. ULTRASONIC RANGING SYSTEMS

Although the method presented in this paper can be applied to other kinds of range sensors or measurements, here we concentrate on ultrasonic range sensing from which our experimental results are derived. We consider simple ultrasonic transducers that measure time-of-flight (TOF) t_o, which is the round-trip travel time of the pulse between the transducer and the point of reflection on the surface. Given the speed of sound c, the range r can be easily calculated from $r = ct_o/2$. Many ultrasonic transducers operate in this pulse-echo mode [6]. The same transducer can function as both the receiver and the transmitter. The major limitation of ultrasonic transducers comes from their large beamwidth. Although these devices return accurate range data, usually they cannot provide direct information on the angular position of the object from which the reflection is

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. In this naive approach, a range reading of r from a transmitting/receiving transducer is taken to imply that an object lies along the line-of-sight of the transducer at the measured range. Consequently, the angular resolution of the surface profile measurement is limited by the rather large beamwidth, which is a major disadvantage. Our approach, as will be seen, turns this disadvantage into an advantage. Instead of restricting oneself to an angular resolution equal to the beamwidth by representing the reflection point as a coarse sample along the line-of-sight, circular or elliptical arcs representing the uncertainty of the object location are drawn. By combining the information inherent in a large number of such arcs, angular resolution far exceeding the beamwidth of the transducer is obtained.

When the same transducer transmits and receives, all that is known is that the reflection point lies on a circular arc whose radius is determined by $r = ct_o/2$, as illustrated in Fig. 1(a). More generally, when one sensor 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 (Fig. 1(b)). Notice that the arcs are tangent to the reflecting surface at the actual point(s) of reflection. The angular extent of these arcs is determined by the sensitivity regions of the transducers. For the same transmitting/receiving pair, if multiple echoes are detected at the receiver, circular or elliptical arcs are drawn to correspond to each echo. As a result of this process, we obtain the *arc map*, comprised of circular and elliptical arcs. An example of an arc map is given in Fig. 2(a).

We distinguish between multiple reflections and higher-order reflections as follows: Multiple reflections are those which are caused by one time (or first-order) reflection of the emitted wave from different points in the environment, resulting in multiple returns at the detector. For instance, along the curved surface, there may be multiple points falling within the sensitivity region of the transducer at which the incoming beam is perpendicular to the surface tangent. In general, such points will be at different distances from the transducer, creating echoes at locations corresponding to these distances. This type of multiple reflections have been modeled in the simulations and also properly registered by the experimental detection circuitry so that they are processed informatively. Some simpler systems detect and process only the first echo received, corresponding to the nearest reflection point.

A higher-order reflection refers to a single return detected after bouncing off successively from more than one point in the environment before returning to the detector. The key idea of the method is that a large number of data points coincide with the actual surface (at least at the tangent points of the arcs) and the data points off the actual curve are more sparse. Those

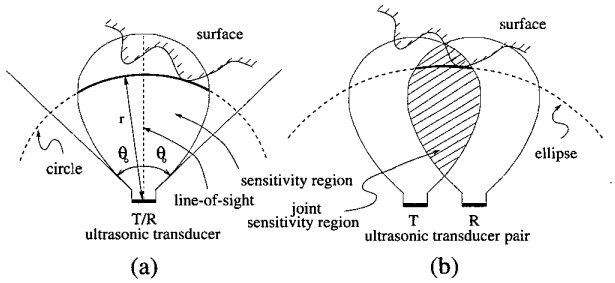

Fig. 1. (a) For the same transducer transmitting and receiving, the reflection point is known to be on the circular arc shown. (b) The elliptical arc if the wave is transmitted and received by different transducers.

spurious arcs caused by higher-order reflections and crosstalk also remain sparse and lack reinforcement most of the time. The algorithm eliminates these spurious arcs together with the sparse arc segments resulting from the angular uncertainty of the sensors. Although higher-order reflections have not been generated in the simulated data sets due to the complexity of modeling them, they naturally exist in the experimental data, providing the opportunity to verify the ability of our method to eliminate them.

We initially consider the case where the surfaces reflect the ultrasonic waves specularly (like a mirror). Since most airborne ultrasonic systems operate at resonance frequencies below 200 kHz, the propagating waves have wavelengths well above several millimeters. Thus, the features on the objects which are smaller than the wavelength cannot be resolved, resulting in specular reflections [2]. Later, we also consider rough surfaces that do not reflect the waves specularly. The ultrasonic devices modeled in the simulations and used in the experiments are Polaroid 6500 series transducers [7], operating at a resonance frequency $f_o = 49.4$ kHz. This corresponds to a wavelength of $\lambda = c/f_o = 6.9$ mm and a half beamwidth angle of $\theta_o = \pm 12.5°$ at room temperature.

III. SURFACE PROFILE DETERMINATION

Structured sensor configurations such as linear and circular arrays, as well as irregularly configured, moving, and synthetic arrays have been considered. As an illustrative example, Fig. 2(a) shows the arc map obtained from a surface using an irregular sensor configuration. A considerably large number of arcs can be obtained with a moderate number of sensors because each sensor can receive pulses transmitted from all the others, provided a reflection point lies in the joint sensitivity region of that sensor pair. For sensors with large beamwidth, the number of arcs obtained approaches the square of the number of sensors. Near the actual reflection point(s), several arcs intersect with small angles. The many small segments of the arcs superimposed in this manner tend to coincide with and cover the actual surface, creating the darker features in Fig. 2(a). Although each arc represents considerable uncertainty as to the

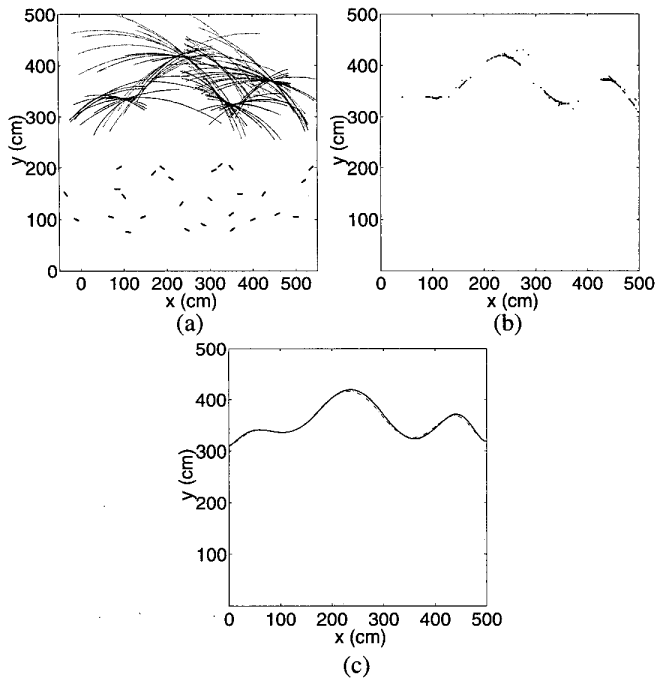

Fig. 2. (a) The arc map and sensor configuration. (b) The result of spatial voting and subsequent thresholding. (c) The curve fitted to part (b) (solid line), and the actual surface (dashed line): $E_1 = 2.74$ cm (simulated data).

angular position of the reflection point, it is often possible to visually perceive the surface profile from the darker features in the arc map (see for example Figs. 3, 4, and 5). The remaining parts of the arcs, not actually corresponding to any reflections and simply representing the angular uncertainty of the sensors, remain more sparse and isolated. Similarly, those arcs caused by higher-order reflections, crosstalk, and noise also remain sparse and lack reinforcement. The purpose of the method described in this paper is to formalize and automate this process. The dark features of the arc maps which eventually reveal the surface profile are essentially caused by two effects: Several arc segments slightly displaced with respect to each other, and several arcs crossing at the same pixel. The method described below takes into account the number of arcs intersecting at each pixel.

A. Spatial Voting

The information contained in the arc map is processed by employing a spatial voting scheme followed by thresholding. The number of arcs which cross at each pixel is kept track of while generating the arc map. A matrix is created which represents the number of arcs crossing each pixel. The values of pixels which have not been crossed by any arcs remain zero. The values of other pixels are equal to the number of arcs crossing them. A pixel size of 1 cm is used in our examples.

The most important parameter in the voting method is the threshold value. A suitable threshold level τ is chosen to select those pixels which have been crossed more frequently. If the number of crossings for a given pixel is less than τ, the value of that pixel is set equal to zero. If the number is greater than or equal to τ, the pixel value is set equal to one. Reconsidering the arc map in Fig. 2(a), the result of applying the above procedure is presented in part (b) of the same figure.

B. Curve Fitting

As a last step, a least-squares polynomial fit is obtained to compactly represent the surface. Since the error distribution after voting is zero mean and symmetric, application of least-squares fit is appropriate. The curve fitted to the processed map given in Fig. 2(b) is displayed in Fig. 2(c). Two error measures, both comparing the final polynomial fit $p(x_i)$ with the actual curve, are introduced:

$$E_1 = \sqrt{\frac{1}{N}\sum_{i=1}^{N}[p(x_i) - y(x_i)]^2} \quad (1)$$

$$E_2 = \frac{E_1}{\sigma_y} \quad (2)$$

The first is a root-mean-square absolute error measure, whereas the second is a dimensionless relative error measure with respect to the variation of the actual curve. N is the total number of columns in the arc map matrix, and $\sigma_y^2 = \frac{1}{N}\sum_{i=1}^{N}[y(x_i) - \frac{1}{N}\sum_i y(x_i)]^2$ is the variance of the actual surface profile $y(x_i)$.

C. Examples from Simulations and Experiments

First, we consider a linear array of sensors. This array has an horizontal extension of 5.0 m with 50 cm spacings between the 11 sensors, as shown in Fig. 3. Given the narrow beamwidth of the sensors (25°), the number of arcs obtained with this many sensors turns out to be insufficient to reconstruct the surface profile [8]. Whereas narrow beamwidths are esteemed for their higher resolving power in conventional usage of ultrasonic transducers, here it would have been desirable to have sensors with larger beamwidths. This would have enabled a greater number of the sensor pairs at hand to produce elliptical arcs, better revealing the surface profile. Instead, we have considered the alternative strategy of rotating the sensors and refiring them several times in order to collect a sufficient number of arcs.

A further consideration is that in practice the number of sensors may be limited. One way to overcome this limitation is to move a smaller array much in the same spirit as synthetic aperture radar (SAR) techniques [9]. However, this is not exactly equivalent to the full array since those elliptical arcs corresponding to pairs of sensors not contained within the actually existing array will be missing.

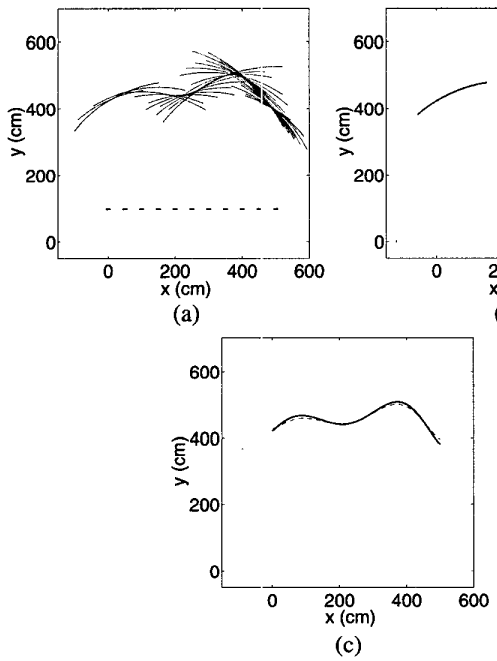

Fig. 3. (a) The arc map obtained from a linear array of 11 sensors with 50 cm spacing where the sensors are individually rotated from 40° to 140° in 10° steps. (b) The result of spatial voting and subsequent thresholding. (c) Curve fitted to part (b) (solid line), and the actual surface (dashed line): $E_1 = 5.42$ cm, $E_2 = 0.227$ (simulated data).

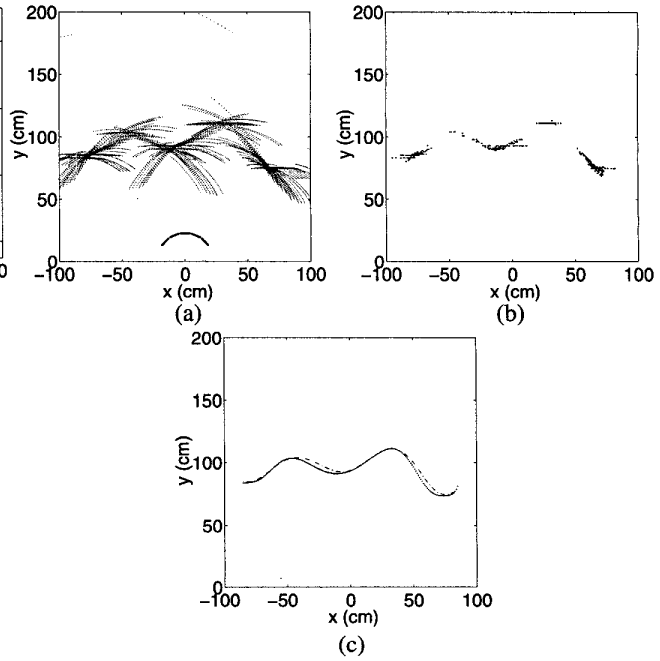

Fig. 4. (a) The arc map and sensor configuration. The data are collected by translating the circular array of five sensors from $(-75, 0)$ to $(75, 0)$ and refiring every 2.5 cm. (b) The result of spatial voting and subsequent thresholding. (c) Curve fitted to part (b) (solid line) and the actual surface (dashed line): $E_1 = 2.12$ cm, $E_2 = 0.206$ (experimental data).

We now return to Fig. 3(a) where the sensors are individually rotated around their shown positions. The result of spatial voting is presented in part (b), and the fitted polynomial is shown in part (c) of the same figure, with $E_1 = 5.42$ cm and $E_2 = 0.227$.

As another example, consider the experimentally obtained arc map shown in Fig. 4(a). These data were collected with a real ultrasonic ranging system, from a cardboard surface constructed in our laboratory. The circular configuration corresponds to the arrangement of ultrasonic sensors on the Nomad 200 mobile robot [10] in our laboratory. An array of five ultrasonic sensors has been moved horizontally over a distance of 1.5 m to increase the total number of arcs, collecting data every 2.5 cm. In the resulting arc map, there are some arcs which are not tangent to the actual surface at any point. (e.g. the isolated arcs faintly visible in the upper part of Fig. 4(a)). These correspond to spurious data due to higher-order reflections, readings from other objects in the environment, or totally erroneous readings. Such points are readily eliminated upon processing (Fig. 4(b)). The polynomial fit to part (b) is shown in Fig. 4(c), with $E_1 = 2.12$ cm, $E_2 = 0.206$. For this example, the actual surface profile was determined by using a very accurate (and much more expensive) laser structured-light system [10].

In the next example, the locations and the line-of-sight orientations of the sensors are generated randomly and do not conform to any special structure. Fig. 5(a) illustrates an arc map obtained with such a sensor configuration. In Fig. 5(b), the result obtained after applying spatial voting and thresholding to the arc map in part (a) are shown. The curve fitted are presented in part (c) of the same figure. The errors obtained by spatial voting and subsequent thresholding and curve fitting are $E_1 = 5.34$ cm, $E_2 = 0.224$.

In the above examples, the value of the threshold employed is that which yields the smallest error E_1 (and consequently E_2). For the examples given, threshold values between 2–6 were found to give the best results. The optimal value of the threshold depends on the amount of measurement noise, the resolution of the arc map, and the total number of arcs, which further depends on the number and configuration of the sensors used. In the simulations, where the actual surface profile is known, it is possible to choose the optimal value of the threshold, minimizing E_1 or E_2. In practice, this is not possible so that one must use parameter values judged appropriate for the given system and the class of surfaces under investigation, based on previous simulations and experiments.

Although structured arrays such as linear or circular ones are

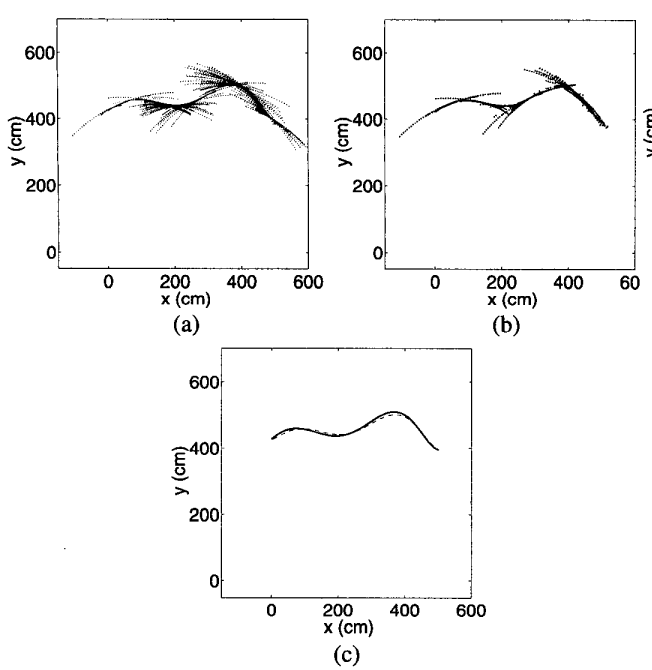

Fig. 5. (a) The arc map obtained by using 150 sensors positioned and oriented randomly (not shown). The x and y coordinates of each sensor are independent and uniformly distributed in the intervals $[0, 500]$ and $[0, 360]$ respectively. The orientation is uniformly distributed in $[40°, 140°]$. (b) The result of spatial voting and subsequent thresholding. (c) Curve fitted to part (b) (solid line), and the actual surface (dashed line): $E_1 = 5.34$ cm, $E_2 = 0.224$ (simulated data).

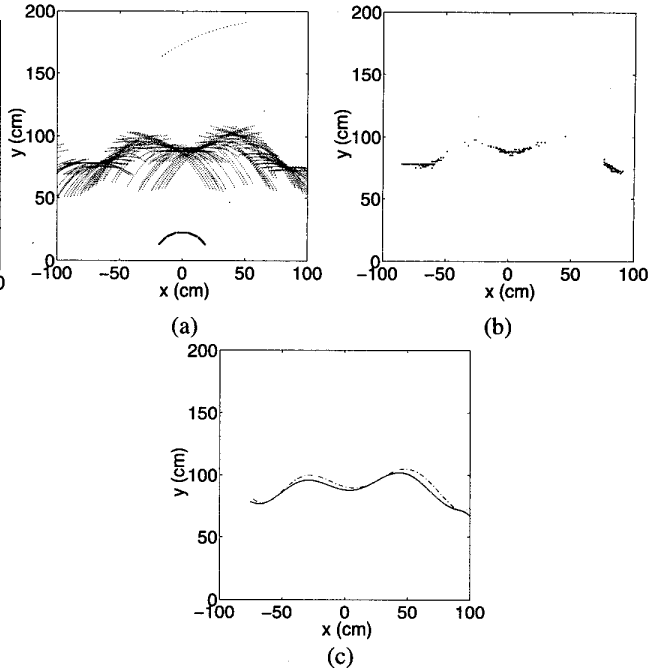

Fig. 6. Spatial voting results from a rough surface covered with packing material with small bubbles. (a) The arc map and the sensor configuration. (b) The result of spatial voting and subsequent thresholding. (c) The curve fitted to part (b) (solid line), and the actual surface (dashed line): $E_1 = 3.45$ cm, $E_2 = 0.357$ (experimental data).

often preferred in theoretical work for simplicity and ease of analysis, the method presented here can handle irregular arrays equally easily. In fact, the large number of simulations we have undertaken indicate that arrays consisting of irregularly located and oriented sensors tend to yield better results. This stems from the fact that the many different vantage points and orientations of the sensors tend to complement each other better than in the case of a structured array. Although the problem of optimal complementary sensor placement is a subject for future research, the large number of simulations performed indicate that it is preferable to work with irregular arrays rather than simple-structured arrays such as linear or circular ones.

A detailed study of the effect of using structured sensor configurations such as linear and circular arrays, as well as irregularly configured and moving sensors in conjunction with morphological processing can be found in [11]. A comparison of the voting method with morphological processing is made in [8]. Although the elements of the irregular transducer array may obscure each other under certain circumstances, such configurations may also correspond to a synthetic array (as described after the first example in this section). Even when it is the case that SAR techniques are not used, the perturbation on the fields caused by other transducers will be negligible if the transducers are relatively small. In those cases where the obstruction cannot be neglected, the only modification required in the algorithm is to simply omit those arcs corresponding to such transducers.

D. The Effect of Surface Roughness

Although the method was initially developed and demonstrated for specularly (mirror-like) reflecting surfaces, subsequent tests were performed also with Lambertian surfaces of varying roughness. Arc maps obtained from rough surfaces generally resemble those from specularly reflecting surfaces, except that the arcs are slightly more dispersed.

We now reconsider the example of Fig. 4. The original surface which produced the arc map in Fig. 4(a) was made of smooth, thin cardboard. In Fig. 6, the results obtained from the same surface when its roughness is varied by covering it with with bubbled packing material are illustrated. Table I provides a summary of the errors obtained with surfaces of varying roughness. The packing material with small bubbles has a honeycomb pattern of uniformly distributed circular bubbles of diameter 1.0 cm and height 0.3 cm, with a center-to-center separation of 1.2 cm. The packing material with large bubbles has the same pattern with diameter, height, and center-to-center separation 2.5 cm, 1.0 cm, and 2.8 cm, respectively.

TABLE I
RESULTS FOR SURFACES OF VARYING ROUGHNESS.

surface material	E_1(cm)	E_2
thin cardboard	2.12	0.206
small-bubbled packing material	3.45	0.357
large-bubbled packing material	3.58	0.448

The results indicate that the method still works for rough surfaces, with slightly larger errors than for specularly reflecting surfaces, as exemplified by Fig. 6(c) and Table I.

IV. DISCUSSION AND CONCLUSION

A novel method is described for determining arbitrary surface profiles by applying spatial voting and subsequent thresholding to data acquired by ultrasonic range sensors. The method is extremely flexible, versatile, and robust, as well as being simple and straightforward. It can deal with arbitrary numbers and configurations of actual and synthetic arrays of sensors. Accuracy improves with the number of sensors used and can be as low as a few centimeters. The method is robust in many aspects; it has the inherent ability to eliminate undesired range readings arising from higher-order reflections, crosstalk, and noise, as well as processing multiple echoes informatively.

From a map-building perspective, the method can be considered as the hybrid of feature-based and grid-based methods: Initially, the environment is discretized into a rectangular array of square cells. After accumulation of a sufficient amount of measurements and their processing using the method, a curve-fitting procedure is employed to extract the geometry of the surface under investigation. The voting method essentially relies on the accumulation of multiple measurements in each pixel so that it does not allow the use of very small grid sizes.

The average CPU times for arc map processing, starting with the raw TOF data, are in general fractions of a second, indicating that the method is viable for real-time applications. The total processing time consists of the data collection time plus the arc map processing time. Map processing operations are implemented in the C programming language and the programs are run on a 200 MHz Pentium Pro PC. On the other hand, the time it takes for an array of 16 ultrasonic sensors to collect all the TOF data is 16×40 ms $= 0.64$ s which is of the same order of magnitude as the processing time. It should be noted that 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 extract the surface profile even faster, bringing the computation time much below the data collection time.

The method extends the types of objects that can be handled by intelligent systems from simple primitives such as planes, corners, edges, cylinders, to also include those with arbitrarily curved surfaces. This is a substantial generalization, increasing the recognition and navigation capability of such systems. This approach, based on the use of multiple range sensors combined with ultrasonic arc map processing, is very basic and general so that it can be applied to different physical modalities of range sensing of vastly different scales and in many different areas. These may include radar, sonar, optical sensing and metrology, remote sensing, ocean surface exploration, geophysical exploration, robotics and intelligent systems, and acoustic microscopy. In our case, the application motivating our research was map building for mobile robotics. For instance, the system demonstrated can be used for continual real-time map building purposes on a robot navigating in a given environment. The robot can continually add to and update its collection of arcs and reprocess them as it moves, effectively resulting in a synthetic array with more sensors than the robot actually has. Apart from indoor mobile robotics, the method can also find application in outdoor land vehicles and other intelligent systems operating in underwater, underground or outer space, and dealing with rough terrain or other curved surfaces.

The method can be readily generalized to 3-D environments with the arcs replaced by spherical or elliptical caps and the spatial voting rules extended to three dimensions. In certain problems, it may be preferable to reformulate the method in polar or spherical coordinates. Some applications may involve an inhomogeneous and/or anisotropic medium of propagation. It is envisioned that the method can be generalized to such cases by constructing broken or non-ellipsoidal arcs.

References

[1] B. Barshan and R. Kuc, "Differentiating sonar reflections from corners and planes by employing an intelligent sensor," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 6, pp. 560–569, June 1990.

[2] M. K. Brown, "The extraction of curved surface features with generic range sensors," *Int. J. Robot. Res.*, vol. 5, no. 1, pp. 3–18, Spring 1986.

[3] B. Parhami, "Voting algorithms," *IEEE Trans. Reliab.*, vol. 43, no. 4, pp. 617–629, December 1994.

[4] L. Lam and C. Y. Suen, "Application of majority voting to pattern recognition: an analysis of its behavior and performance," *IEEE Trans. Syst., Man, Cybern.*, vol. 27, no. 5, pp. 553–568, September 1997.

[5] S. W. Utete, B. Barshan, and B. Ayrulu, "Voting as validation in robot programming," *Int. J. Robot. Res.*, vol. 18, no. 4, pp. 401–413, April 1999.

[6] P. Hauptmann, *Sensors: Principles and Applications*, Prentice-Hall, Englewood Cliffs, NJ, 1993.

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

[8] B. Barshan and D. Başkent, "Comparison of two methods of surface profile extraction from multiple ultrasonic range measurements," *Meas. Sci. Technol.*, vol. 11, no. 6, pp. 833–844, June 2000.

[9] M. I. Skolnik, *Introduction to Radar Systems*, McGraw-Hill, New York, NY, 1981.

[10] Nomadic Technologies, Inc., Mountain View, CA, *Nomad 200 Manual*, 1997.

[11] D. Başkent and B. Barshan, "Surface profile determination from multiple sonar data using morphological processing," *Int. J. Robot. Res.*, vol. 18, no. 8, pp. 788–808, August 1999.