Two (somewhat unlikely) conditions will cause errors in the IPEC system. These are as follows.

1) If sufficiently strong external forces are applied onto the OmniMate so that either the front or rear truck slip laterally, then this will cause an error in the measured orientation error $\Delta\theta_m$ (see Fig. 3).
2) If both wheels of truck A slipped by the same amount, then this slippage would not be detected. However, the result would only be a translational error, which has usually less impact than an orientation error.

V. CONCLUSION

This paper presents results of odometric accuracy tests performed with a new, commercially available mobile robot called "OmniMate." The OmniMate provides true omnidirectional (i.e., holonomous) motion and its kinematic design eliminates the excessive wheel-slippage often associated with omnidirectional platforms. One of the OmniMate's most unique features is its ability to employ internal position error correction (IPEC) to dramatically improve its odometric accuracy.

Using rigorous test procedures called "UMBmark" and "extended UMBmark," the OmniMate and its implementation of IPEC were carefully tested at our lab. The results show an improvement of one order of magnitude in odometric accuracy over the accuracy of a conventional odometry system.

The foremost strength of the IPEC method is its ability to reliably and accurately detect and correct nonsystematic odometry errors such as those caused by bumps, cracks, or other objects on the floor. In conventional mobile robots the encounter of one or more such irregularities could have a catastrophic effect on the performance of the robot, i.e., cause the mission to fail completely. With the OmniMate and IPEC, on the other hand, floor irregularities have virtually no detrimental effect on the odometric accuracy of the vehicle at all.

REFERENCES

[1] J. Borenstein, "Internal correction of dead-reckoning errors with the compliant linkage vehicle," *J. Robot. Syst.,* vol. 12, no. 4, pp. 257–273, Apr. 1995.
[2] J. Vaganay, M. J. Aldon, and A. Fournier, "Mobile robot attitude estimation by fusion of inertial data," in *Proc. IEEE Int. Conf. Robot. Automat.,* Atlanta, GA, May 10–15, 1993, pp. 277–282.
[3] B. Barshan and H. F. Durrant-Whyte, "Inertial navigation systems for mobile robots," *IEEE Trans. Robot. Automat.,* vol. 11, pp. 328–342, June 1995.
[4] K. Komoriya and E. Oyama, "Position estimation of a mobile robot using optical fiber gyroscope (OFG)," in *Proc. Int. Conf. Intell. Robot. Syst. (IROS'94),* Munich, Germany, Sept. 12–16, 1994, pp. 143–149.
[5] HITACHI-Hitachi Cable America, Inc., White Plains, NY.
[6] ANDREW-Andrew Corporation, Orland Park, IL.
[7] J. Borenstein, B. Everett, and L. Feng, *Navigating Mobile Robots: Systems and Techniques.* Wellesley, MA: A. K. Peters, Ltd., Feb. 1996.
[8] D. B. Reister, "A new wheel control system for the omnidirectional HERMIES-III robot," in *Proc. IEEE Conf. Robot. Automat.,* Sacramento, CA, Apr. 7–12, 1991, pp. 2322–2327.
[9] M. West and H. Asada, "Design of a holonomic omnidirectional vehicle," in *Proc. 1992 IEEE Int. Conf. Robot. Automat.,* Nice, France, May 1992, pp. 97–103.
[10] S. Hirose and S. Amano, "The VUTON: High payload efficiency holonomic omni-directional vehicle," in *Proc. 6th Int. Symp. Robot. Res. (ISRR-93),* Hidden Valley, PA, Oct. 2–5, 1993.
[11] F. G. Pin and M. Killough, "A new family of omnidirectional and holonomic wheeled platforms for mobile robots," *IEEE Trans. Robot. Automat.,* vol. 10, pp. 480–489, Aug. 1994.
[12] H. Sugiyama, "A method for an autonomous mobile robot to recognize its position in the global coordinate system when building a map," in *Proc. 1993 IEEE/RSJ Int. Conf. Intell. Robot. Syst.,* Yokohama, Japan, July 26–30, 1993, pp. 2186–2191.
[13] R. Kurazume and S. Nagata, "Cooperative positioning with multiple robots," in *Proc. IEEE Int. Conf. Robot. Automat.,* San Diego, CA, May 8–13, 1994, pp. 1250–1257.
[14] T. Hongo, H. Arakawa, G. Sugimoto, K. Tange, and Y. Yamamoto, "An automated guidance system of a self-controlled vehicle," *IEEE Trans. Ind. Electron.,* vol. IE-34, no. 1, pp. 5–10, 1987.
[15] J. Borenstein and L. Feng, "UMBmark: A benchmark test for measuring dead-reckoning errors in mobile robots," in *Proc. 1995 SPIE Conf. Mobile Robots,* Philadelphia, PA, Oct. 22–26, 1995.

Vision-Based Navigation by a Mobile Robot with Obstacle Avoidance Using Single-Camera Vision and Ultrasonic Sensing

Akihisa Ohya, Akio Kosaka, and Avinash Kak

Abstract— This paper describes a vision-based navigation method in an indoor environment for an autonomous mobile robot which can avoid obstacles. In this method, the self-localization of the robot is done with a model-based vision system, and nonstop navigation is realized by a retroactive position correction system. Stationary obstacles are avoided with single-camera vision and moving obstacles are detected with ultrasonic sensors. We will report on experiments in a hallway using the YAMABICO robot.

Index Terms— Mobile robot, non-stop navigation, obstacle avoidance, self-localization, ultrasonic sensing, vision.

I. INTRODUCTION

In what has become a fairly well-researched approach to vision based navigation for mobile robots, a robot is provided with an environmental map and a path to follow. The important function of vision-based processing in this case consists of "self-localization." For literature on this approach, the reader is referred to [6], [13], [18], and [19]. In a different approach, as reported on by [11] and [15], a robot is provided with a sequences of images of the interior space. By comparing these prerecorded images with the camera images taken during navigation, the robot is able to determine its location. Other previous research contributions that are relevant to mobile robot localization include [4], [5], [8], [12], [14], [16], and [17].
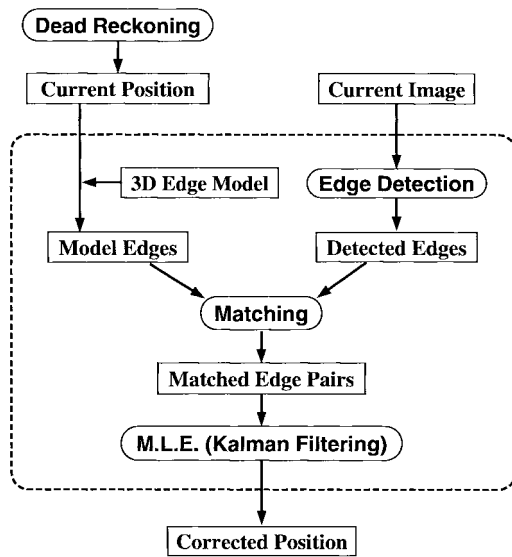
Fig. 1.   Flow of computations in the self-localization procedure.

An important adjunct to the problem of navigation is the problem of obstacle avoidance, another well-researched problem in the past [1], [2]. In the vision-based navigation work reported in the past, such as in [6], obstacle avoidance is carried out using ultrasonic sensors. These sensors take over the control of the robot as long as obstacles are detected in the vicinity. The control is then handed back to vision-based processing once the obstacles are no longer a factor. While it is expedient to use ultrasonic sensors in such a manner, we believe it is empirically more interesting to use vision-based processing for obstacle avoidance also. Since a primary focus of machine intelligence and advanced robotics is to capture human faculties in a computer and since humans use vision for obstacle avoidance while navigating, we are evidently interested in doing the same in a robot.

In this paper, we will present an integrated vision-based process for mobile robots that is capable of simultaneously navigating and avoiding stationary obstacles using monocular camera images. While the self-localization part of the process is the same as the FINALE system of [6], what distinguishes the work reported in this paper is that we are now able to give to the robot a vision-based obstacle avoidance capability at the same time. In the current implementation, this obstacle-avoidance capability is limited to the detection and avoidance of stationary obstacles. This is owing to the limitations of the computing hardware available to the robot. Therefore, moving obstacles must still be detected with ultrasonic sensors.

What is particularly noteworthy about our approach is that the self-localization and the obstacle avoidance are both carried out by processing the same image, thus eliminating what would otherwise be a redundancy in sensor data collection. While a model-based approach is used for self-localization, obstacles are detected by computing the difference between the edges estimated from the 3D environment model and the edges detected from the actual camera image. We should mention that is not the only approach to vision-based detection of obstacles. As reported by [3], an alternative approach consists of computing optical-flows from the images.

In the rest of this paper, in Section II we briefly explain the method used for self-localization and show also the concept of retroactive position correction for nonstop navigation. The method used for obstacle avoidance is presented in Section III. The system architecture is presented in Section IV and the experimental results obtained are presented in Section V. Finally the conclusions and the discussions are given in Section VI.

II. VISION-BASED NAVIGATION

A. Self-Localization

As mentioned earlier, the self-localization part of the overall navigation scheme is the same as the FINALE system of Kosaka and Kak [6]. However, the manner in which self-localization is used in the current system is different from that in [6]. Self-localization in FINALE kicks in whenever the variances associated with the positional parameters exceed a certain predetermined threshold.[1] In our current system, self-localization is carried out on a continuous basis. However, before we elaborate on what we mean by "self-localization on a continuous basis," we'd like to review briefly the computations that go into self-localization.

Fig. 1 shows the flow of the self-localization procedure. First, the robot renders an expectation image using its current best estimate of where its present location is. Next, the model edges extracted from the expectation image are compared and matched with the edges extracted from the camera image through an extended Kalman filter. The Kalman filter automatically then yields updated values for the location and the orientation of the robot.

To illustrate the process of self-localization, Fig. 2(a) shows a typical camera image. Shown in Fig. 2(b) is an expectation image rendered from the wire-frame model of the environment; this expectation map is overlaid on the camera image. As the reader can see, the discrepancy between the various edges in the underlying camera image and the highlighted edges in the expectation map is caused by the error between where the robot actually is and where the robot thinks it is. Shown in Fig. 2(c) are the edges extracted from the camera image. Note in particular that not all gray level variations in the camera image translate into edges. As explained in [6], this is due to the fact that the system only looks for those edges in the camera image that are in proximity—both spatially and in the Hough space—to the edges in the expectation map. Shown in Fig. 2(d) is a reprojection into the camera frame of those model edges that were successfully used for self-localization. The fact that these reprojected edges fall exactly where they should is a testimony to the accuracy of the result produced by the Kalman filter. Although not discernible, shown in Fig. 2(e) are two small icons, in close proximity to each other, the bright one corresponding to the updated position and orientation of the robot and the somewhat darkened corresponding to the old position and orientation. To help the reader discern these two icons, shown in Fig. 2(f) is an enlarged version of the image in Fig. 2(e).

By repeating the self-localization, the robot can correct its position error and navigate autonomously toward its destination.

B. Nonstop Navigation by Retroactive Position Correction

We now need to explain to the reader how the self-localization process described above fits into the overall framework for navigation. What we really want to do—and this has actually been implemented—is to navigate with dead-reckoning, meaning that we want the robot to update its position continuously on the basis of the information supplied by the wheel encoders. Unfortunately, there is always some differential slippage in the wheels that causes a purely dead-reckoning based approach to go awry if navigation is attempted over significant distances. So we want the robot to use vision-based self-localization to eliminate the errors accumulated during

---

[1] The uncertainty in positional parameters is caused primarily by differential slippage in the wheels of the robot, especially when the robot is engaged in collision-avoidance maneuvers. The two contributors to positional uncertainty are the translational uncertainty and the rotational uncertainty, the latter being the larger of the two.

Fig. 2. Sample images from the self-localization process: (a) camera image, (b) expectation map overlaid on the camera image, (c) edges extracted from the camera image, (d) matched model edges reprojected into the camera frame, (e) two small icons showing the robot's old and the updated positions in the hallway, and (f) enlarged version of (e). The gray icon is for the old position and the white icon is for the updated position.

dead-reckoning. Due to the time delays associated with vision-based processing, this turns out to be a fortuitous combination of dead-reckoning and vision-based position updating. Since we do not want the robot to come to a halt as the camera image is being processed, what is really needed is a retroactive approach to position correction, along the lines originally proposed in [7]. As explained below, in the current system we have implemented a slight variation on the retroactive updating idea reported in [10].

Fig. 3 shows the timing diagram used for retroactive position correction. At time t_0, an image is taken and the self-localization process started. The computations for self-localization come to an end at time $t_0 + n\Delta t$ where Δt is the sampling interval for dead reckoning, meaning that at the end of each Δt interval the position of the robot is recalculated based on the wheel encoder readings.

Since the self-localization results correspond to the image taken at time t_0, the dead-reckoning based estimate of robot's position at the current time (the time instant $t_0 + n\Delta t$) must be recalculated.

Since this recalculation is done within a sampling interval for dead reckoning, the robot is able to proceed in a nonstop manner without having to stop for processing the camera image.

III. OBSTACLE AVOIDANCE

Obstacle avoidance is carried out using both vision and ultrasonic sensing. While our ultimate goal is to use only vision for obstacle avoidance, due to the limitations of the computing hardware available to the mobile robot, at this time vision can only be used for the detection of stationary obstacles. So, in the current implementation, the detection of moving obstacles is left to ultrasonic sensors. Another

Fig. 3. The timing diagram for retroactive position correction. At time t_0, an image is taken and self-localization process starts. Self-localization results become available at time $t_0 + n\Delta t$ where Δt is sampling interval for dead reckoning. The corrected current robot's position corresponding to time $t_0 + n\Delta t$ is recalculated.

reason for using a combination of these two sensing modalities is that the ultrasonic sensors on the YAMABICO robot possess low directivity. Therefore, while they are able to discern the presence of an obstacle, they cannot be used to determine the direction of free passage. As we will show, the direction of free passage can be easily determined from the vision data.

A. Stationary Obstacle Avoidance

The detection of stationary obstacles is based on the premise that the camera image and the expectation map must be in near-perfect registration *immediately after self-localization.* So, any discrepancy between these two images can only be caused by the presence of obstacles in the environment, assuming of course that the edge-detection process does not result in any artifact edges in the camera image. Fortunately, as will be shown presently, the artifact edges, caused mostly by glare and other illumination dependent phenomena, can be eliminated by using adaptive thresholding.

1) Adaptive Thresholding: We will now explain how the adaptive thresholds are found for the edge detection operator. Recall, the edge-detection thresholds at the different locations of the robot must be such that when an edge detector is applied to a camera image immediately after self-localization, it should not yield any artifact edges. For a particular hallway with given illumination conditions, these thresholds are found through a learning procedure prior to any navigation in that hallway. The learning procedure consists of

1) clear the hallways of all obstacles;
2) manually drive the robot through the different sections of the hallway;
3) render expectation maps from the hallway model at regular intervals during these manual traversals by the robot;
4) record the camera images at the same locations where the expectation maps are rendered;
5) apply an edge detection operator to the camera images using a threshold T for the edge detection operator;
6) construct a difference image between the model edge map from Step 3 and the edge map from the previous step;
7) divide the difference image of Step 6 into five vertical regions and count the numbers of pixels, N_1 through N_5, in each region;
8) compute max N_m of these five numbers N_1 through N_5;
9) go back to Step 5 and by iterating Steps 5–8, construct a graph of the number of pixels N_m in Step 8 versus the threshold T;
10) choose the threshold T_0 for which the number of difference pixels N_m is a minimum. Designate this value of N_m by N_0.


Fig. 4. Determination of the optimum edge detection thresholds.

The reason why the index N_m is not just a total number of pixels in the whole difference image and is calculated in the manner mentioned above, is that it will also be used as another adaptive threshold value for determining the direction of safe passage. Fig. 4 is a pictorial depiction of these steps for the determination of the optimum threshold T_0 at each location of the robot.

To appreciate why there would be a threshold T_0 that would minimize the number of difference pixels, note that for high values of T both the real hallway edges and the artifact edges will be suppressed. As the threshold is increased, there will come a point where no edges will be extracted from a camera image. So large values of T will yield a high count for the number of pixels in the difference image. At the other extreme, when T is too small, the camera image will yield an excessively large number of edges, most corresponding to random gray level fluctuations in the camera image. So the number of difference pixels will again be large. In between these two extremes, there will be a value for T, designated T_0, for which the number of difference pixels N_m will be a minimum denoted by N_0. Note again that N_m is the maximum of the number of pixels N_1 through N_5 counted in the five vertical regions of the difference image. Shown in Fig. 5(a) is a camera image taken from a position where the rendered expectation map looks like what is shown by the overlaid white lines in Fig. 5(b).[2] Shown in Fig. 5(c) is a plot of the difference pixels N_m obtained in Step 8 above for different values of the threshold T. Also marked in this figure are the threshold T_0 and the number N_0 of pixels that corresponds to the threshold T_0.

---

[2]Only the vertical edges of the rendered expectation map are displayed as overlaid white lines. That's because we have found it sufficient to use just the vertical lines for the detection of stationary obstacles. However, the entire procedure can be readily extended to edges of arbitrary orientation, albeit at a higher computational cost.

(a)   (b)

(c)

Fig. 5. (a) Camera image, (b) vertical lines in the rendered expectation map overlaid on the camera image, and (c) plot showing the number of difference pixels N_m versus the threshold T.


Fig. 6. The robot's trajectory for the learning of adaptive thresholds.

Fig. 7. The edge detection thresholds T_0 and the minimum number of difference pixels N_0 for each reference point.

To illustrate a result of the ten-step procedure for determining the adaptive thresholds for an entire section of a hallway, shown in Fig. 6 is a section of the hallway immediately outside our laboratory. The robot was manually driven along the trajectory marked by x's. Along this trajectory, expectation maps were rendered at regular intervals and the corresponding camera images recorded. (In order not to jumble up the display in Fig. 6, only every fourth position where images were rendered and recorded is shown in the figure.) The edge detection thresholds T_0 and the numbers of pixels N_0 corresponding to T_0 are shown in Fig. 7. This graph constitutes the table used in the adaptive thresholding process during autonomous navigation.

*2) Obstacle Detection:* Fig. 8 shows the flow of computations for the obstacle detection procedure by vision. During autonomous navigation, in accord with the statements made at the beginning of Section III-A, obstacles are found from the difference of vertical edges in the camera image and the expectation map immediately after each exercise in self-localization. As shown in Fig. 8, model vertical edges in the scene are estimated from the 3-D edge model of the environment using the robot's position corrected by the self-localization first. Next, vertical edges are extracted from the camera image using for the detection threshold a value that corresponds to the nearest position in the table lookup. The vertical edges thus found
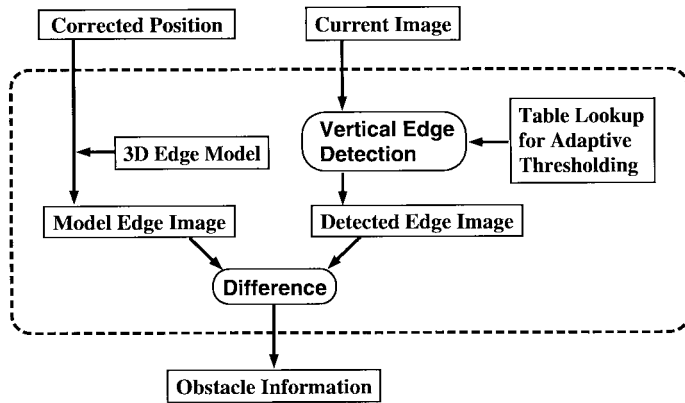
Fig. 8. The obstacle detection procedure using vision during autonomous navigation.

are compared with the model vertical edges. Any extraneous vertical edges must come from any obstacles in the environment.

Fig. 9 shows sample images for a typical exercise in obstacle detection by vision. Displayed by overlaid white lines in Fig. 9(a) are the model vertical edges as obtained from the expectation map rendered from the environment model. The tall dark object on the left is a wooden box that serves as an obstacle. Shown in Fig. 9(b) are the vertical edges extracted from the camera image using the nearest optimum edge-detection threshold T_0 from the table of such thresholds for the hallway in question. It is interesting to note that a part of vertical edge on the right edge of the obstacle has disappeared because of the low contrast between the door and the obstacle. Shown in Fig. 9(c) is the difference edge image for the vertical edges extracted from the camera image and the model vertical edges. Shown in Fig. 9(d) are robot's position, obstacle information and passage space in the hallway. The robot is shown as an inverted "T" icon, the obstacle by a small white rectangle, and the safe-passage direction by two closely spaced parallel lines. How the direction of safe passage is computed will be explained next.

3) Determining the Direction of Safe Passage: To find a direction for safe passage in the presence of obstacles, as pointed out in the ten-step procedure outlined in the previous subsection, the difference edge image [an example of which is Fig. 9(c)] is divided into five vertical regions, as shown in Fig. 10(a) where a tall dark obstacle is present in the left. Number of pixels in each of the five regions is summed. If this sum exceeds the threshold N_0 for this location of the robot, the directions corresponding to that vertical region in the camera image are considered to be unsafe.[3] Recall from our previous discussion, N_0 is the smallest number of N_m which is the maximum of the pixel numbers N_1 through N_5 counted in the five vertical regions of the difference edge image, and N_0 results from the application of the optimum threshold T_0. Of course, as with the application of the edge detection thresholds T_0, the location for which N_0 is available will often not coincide exactly with the current location of the robot. So, as was the case for edge detection, N_0 for the nearest entry in the table is used. Shown in Fig. 10(b) are the summed values for each of the five regions marked in Fig. 10(a).

The view-space in front of the robot is given one of three labels:

1) obstacles;
2) unknown;
3) open space.

---

[3] The threshold that is actually used for detecting obstacles is $(1 + \epsilon) \times N_0$ where a nonnegative ϵ accounts for the nearly always-present discrepancy between the nearest location of the robot where N_0 is available and the location of the robot where it is needed during navigation. A value of $\epsilon = 0.5$ has worked for us in almost all our experiments.

Initially, the directions corresponding to "obstacles" are made to correspond to those of the five vertical regions in the camera image whose difference pixel sum exceeds N_0, as explained above. By the same token, initially the directions corresponding to "unknown" are all those that are outside the view cone corresponding to the camera image. Both the "obstacles" and the "unknown" are expanded by half the width of the robot to ensure safe clearance. The rest of the directions then yield us the directions of the safe passage for the robot. Since the width of the YAMABICO robot is assumed to be twice as large as one of the vertical regions in Fig. 10 at a distance of one meter in front of the camera, free passage space is narrowed to the size of one vertical region on the right of the obstacle in Fig. 10(c). Distance to obstacles cannot be computed from the visual information, since only a single camera is used in this method. It is for this reason we make the conservative assumption that all obstacles and unknown regions are situated 1 m in front of the robot. After obstacle avoidance, the robot seeks to approach its originally planned path to the destination.

B. Moving Obstacle Detection

As mentioned before, the currently implemented vision system is effective only for detecting stationary obstacles because of limitations on the computing power available to the robot. The ultrasonic sensors, which can measure the distance to the object in almost real time by the pulse-echo technique, are therefore used for detecting moving obstacles. Of course, since the ultrasonic sensors have no intrinsic way of distinguishing between moving and stationary obstacles, the interaction between the two sensor systems—the vision sensor and the ultrasonic sensors—becomes important. How the robot is controlled by the two sensor systems is predicated on the following attributes of the sensors:

1) view angles of the two sensor systems are nearly identical ($60°$);
2) by using appropriate time gating of the received echoes, the ultrasonic sensors do not detect an obstacle if it is farther than 50 cm from the robot;
3) vision sensor is capable of detecting stationary obstacles at ranges far exceeding 50 cm.

Shown in Fig. 11 are the view angles for the two sensor systems. Note that the ability of an ultrasonic sensor to detect an obstacle depends on many factors, such as the orientation, reflectively, curvature, etc., of the surface of the obstacle toward the sensor and on the threshold used for the detection of the received echoes. The view angle for the ultrasonic sensor system shown in Fig. 11 corresponds to the case of a unit area metallic surface (of nearly unit reflectivity) perpendicular to the line of the sight from the center of the sensor system and the same detection threshold used for all directions.

IV. SYSTEM ARCHITECTURE

The navigational system described in this paper was implemented on the YAMABICO robot [20] shown in Fig. 12. Because this robot doesn't have a powerful image processing module, a workstation is used as an off-board image processing unit at this time. The robot communicates with the workstation over two links, a video link for the transmission of images and a two-way RF link for data. YAMABICO has an ultrasonic sensing module and can use several pairs of ultrasonic transmitters and receivers that can independently measure distances up to 5 m to an obstacle. The directivity of each sensor pair spans $30°$. Three pairs of transducers can monitor about $60°$ of the view-space in front of the robot, in accordance with Fig. 11. Since we wanted to predominantly use vision for collision

Fig. 9.   Sample images for obstacle detection by vision: (a) model vertical edges overlaid on the camera image, (b) detected vertical edges from the camera image, (c) the difference edge image, and (d) robot's position, obstacle information and passage space in the hallway.

avoidance, time gating was used to limit the range of ultrasonic collision avoidance to 50 cm.

Fig. 13 shows the system hardware configuration. An image is taken by the camera mounted on the robot and it's sent to the workstation through a wireless video transmitter and receiver. The received image is processed on the workstation and the resulting data sent to the robot over the RF link. The camera used here is XC-999 (Sony). As was shown previously in Fig. 11, the horizontal angle of viewing field is 60° approximately. The image is digitized into 256 × 240 pixels.

Fig. 14 shows the system software configuration. On the mobile robot side, the robot is basically navigated on the planned path by a "route-runner" process. For self-localization, the "route-runner" process gets the robot's current position from the "retroactive position correction" process and sends it to the "self-localization" process on the workstation. The "self-localization" process analyzes the image received from the camera on board the robot and calculates the robot's position. The calculated position, corresponding to the instant when the camera image was taken, is sent to the "retroactive position correction" process that updates and corrects the odometry information for the latest position of the robot. For obstacle detection by vision, "obstacle detection" process on the workstation generates the obstacle information from the camera image and the corrected robot's position calculated by the "self-localization" process. The obstacle information is sent to the "route-runner" process and is used for the calculation of the direction of safe passage. For obstacle detection by ultrasonic sensing, the "US-checker" process checks the distance measured by the ultrasonic sensors and sends obstacle information to the "route-runner" process in order to stop the robot when it detects an obstacle.

We will now explain how the two sensor systems interact for collision avoidance. When the robot faces a stationary obstacle, the obstacle is first detected by vision since the range of ultrasonic sensors is limited to 50 cm. In the manner explained previously, the vision data is used for the calculation of the direction of safe passage and the robot turns toward this direction. For those obstacles that are avoided on the basis of vision data, the ultrasonic sensors will usually not record any echoes for the simple reason that the robot has already steered away from them. If per chance one or more of the ultrasonic sensors do receive echoes from such an obstacle, the robot comes to
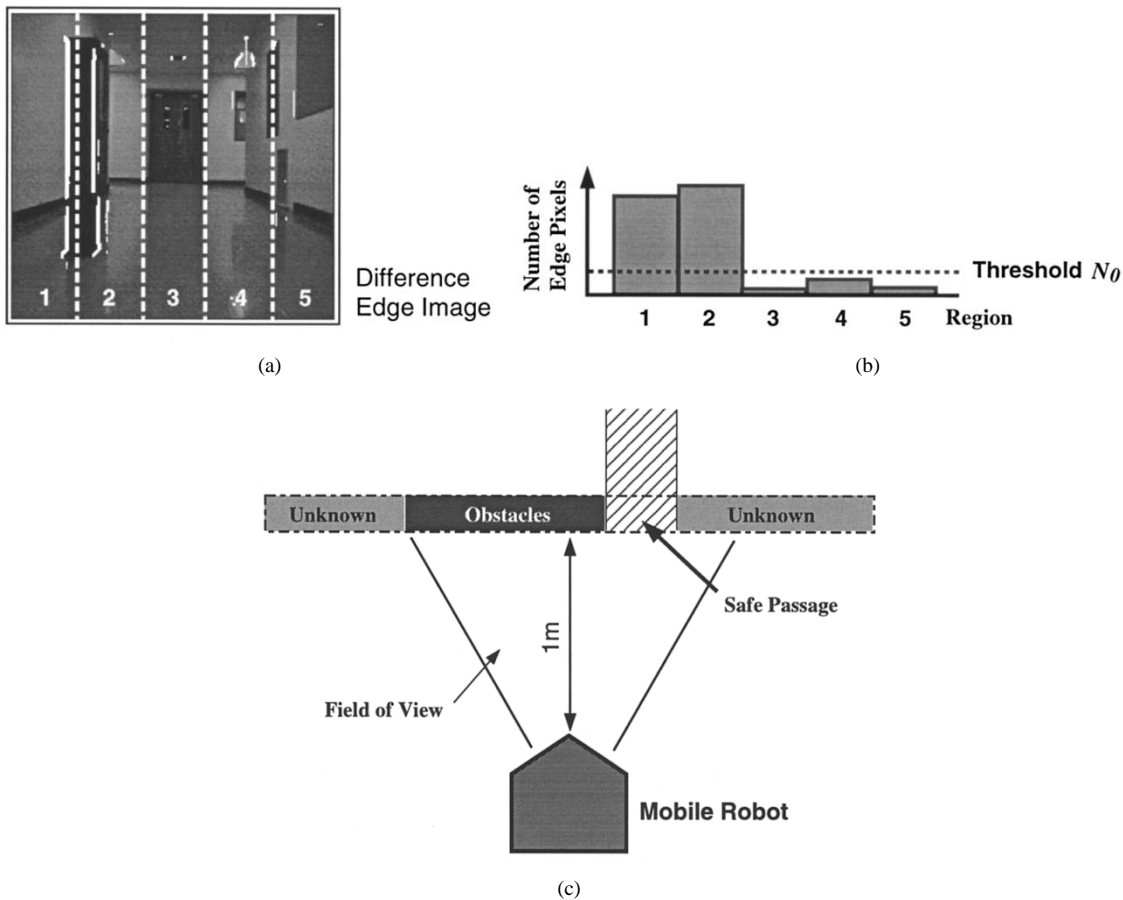
Fig. 10. Determining the direction of safe passage in the presence of obstacles: (a) division of the difference edge image into five vertical segments, (b) total number of pixels in each of the five vertical segments of the difference edge image, and (c) direction of the safe passage obtained by excluding "obstacles" and "unknown" regions.

a halt and uses its vision again for calculating afresh the direction of safe passage. In case of a moving obstacle, of course, it could be detected by vision if the obstacle makes its appearance just at the moment the camera image is recorded. Such an obstacle would be treated in the same manner as a stationary obstacle. Of course, the obstacle such as this would have moved to a different location by the time the robot arrives in its vicinity. In such cases, the robot would seem to be avoiding phantom obstacles, but not to any great detriment of the overall navigational performance. Note that after each exercise in collision avoidance, the robot seeks to go back to its planned path to the destination. If the moving obstacle gets too close to the robot, the ultrasonic sensors take over and bring the robot to a momentary halt. Subsequently, the vision based processing is reinitiated for further navigation.

V. EXPERIMENTAL RESULTS IN AUTONOMOUS NAVIGATION

Using the system presented in the previous section, several experiments were performed in the hallway shown in Fig. 15. The size of the grid shown by the dotted lines is 1 m on each side.

An example of autonomous navigation on the planned path, as displayed by the black line in Fig. 15, is shown in Fig. 16. In this figure, the robot's positions corrected by each self-localization exercise are shown as x's. The total length of the path is approximately 12.5 m. During this experiment, a moving obstacle "A" (a human) suddenly crossed the path just in front of the robot. The ultrasonic sensors brought the robot to a halt, followed by the invocation of vision sensors for determining the direction of safe passage. The robot also


Fig. 11. View angle of the camera and detection area of the ultrasonic sensor. In the next section, we will explain how the two sensor systems interact.

successfully avoided a stationary obstacle "B" (a wooden box) by using vision.

The speed of the robot was set at 10 cm/s for this experiment. The processing time for one image was approximately 10 s on the workstation (a SUN SPARCstation 20).

VI. CONCLUSION

We presented in this paper a vision-based navigational system for mobile robots that is also capable of avoiding at least the stationary
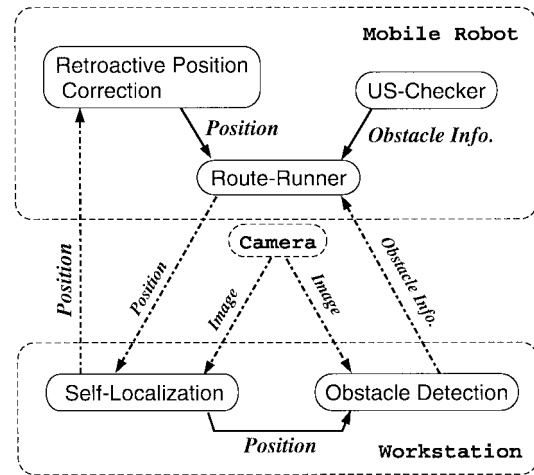
Fig. 12. Autonomous mobile robot YAMABICO.

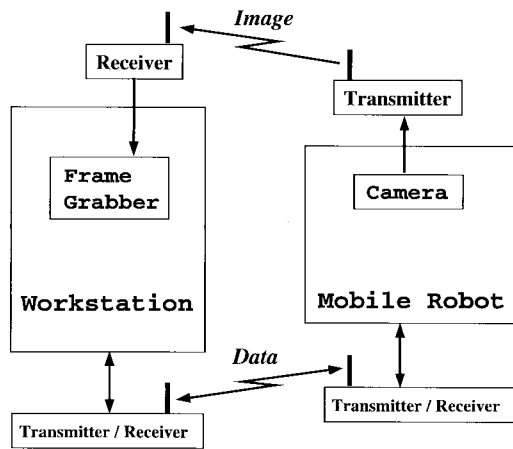Fig. 14. System software configuration.


Fig. 13. System hardware configuration.

Fig. 15. Map of the hallway. The dark line denotes the planned path.


Fig. 16. An example of autonomous navigation. The robot's actual trajectory is shown as x's. A moving obstacle "A" crossed the path and there was a stationary obstacle "B" to deal with.

obstacles using vision data. By using a combination of model-based vision for self-localization; retroactive position updating to cope with the time delays associated with image processing; using vision data for not only self-localization but also for the calculation of directions of safe passage in the presence of obstacles; and ultrasonic sensors for the detection of close-range moving obstacles; we have created a navigational system that makes optimum use of all the sensors for smooth and continuous navigation in indoor environments.

As with all such systems dealing with higher-level robotic intelligence, the performance can never be expected to be completely foolproof. The best that one can do is to devise appropriate automatic error correction and detection strategies. To briefly discuss the various failure modes of our system, the vision-based collision avoidance capability depends obviously on the visual contrast between the obstacle and the interior of the hallway. The size of an obstacle will also play an important role in its detectability by vision. To gain an

understanding of these limitations, we performed experiments using two small cardboard boxes of two different colors, brown and white, each of height 35 cm and width 24 cm, as test obstacles. As for the results, the robot was able to detect the white box in all cases, but in 50% of the cases failed to detect the brown box. When the robot did fail to detect the brown box, in some cases it was because it could

extract edges for only one side of the box. It is entirely possible that superior image processing strategies would enhance the performance of vision-based collision avoidance. Our future research will address this issue.

Another interesting limitation of our system arises if the number of obstacles in the environment is such that they fill the camera image and make it difficult for the robot to compute a direction for safe passage. This can happen even when the obstacles are situated in such a manner that the robot could meander through the obstacles and get to its destination. Since at this time we are using only a single camera, the range to the obstacles cannot be computed using just vision data. So, too many obstacles, even when situated sufficiently apart to permit navigation, can make it impossible for the robot to proceed.

ACKNOWLEDGMENT

REFERENCES

[1] J. Borenstein, "Real-time obstacle avoidance for fast mobile robots," *IEEE Trans. Syst. Man Cybern.,* vol. 19, pp. 1179–1187, Sept./Oct. 1989.

[2] J. Borenstein and Y. Koren, "Histogramic in-motion mapping for mobile robot obstacle avoidance," *IEEE Trans. Robot. Automat.,* vol. 7, pp. 535–539, Aug. 1991.

[3] T. Camus, D. Coombs, M. Herman, and T. Hong, "Real-time single-workstation obstacle avoidance using only wide-field flow divergence," in *Proc. 13th Int. Conf. Pattern Recog.,* Aug. 1996.

[4] J. L. Crowley and P. Reignier, "Asynchronous control of rotation and translation for a robot vehicle," *Robot. Auton. Syst.,* vol. 10, pp. 243–251, 1992.

[5] M. Devy and H. Bulata, "Landmark-based vs. feature-based localization of a mobile robot in a structured environment," in *Proc. Int. Conf. Adv. Robot.,* 1995.

[6] A. Kosaka and A. Kak, "Fast vision-guided mobile robot navigation using model-based reasoning and prediction of uncertainties," *CVGIP—Image Understanding,* vol. 56, no. 3, pp. 271–329, Nov. 1992.

[7] A. Kosaka, M. Meng, and A. Kak, "Vision-guided mobile robot navigation using retroactive updating of position uncertainty," in *Proc. IEEE Int. Conf. Robot. Automat.,* May. 1993, pp. 1–7.

[8] J. J. Leonard and H. F. Durrant-Whyte, "Mobile robot localization by tracking geometric beacons," *IEEE Trans. Robot. Automat.,* vol. 7, pp. 376–382, June 1991.

[9] J. J. Leonard, H. F. Durrant-Whyte, and I. J. Cox, "Dynamic map building for an autonomous mobile robot," *Int. J. Robot. Res.,* vol. 11, no. 4, pp. 286–298, Aug. 1992.

[10] S. Maeyama, A. Ohya, and S. Yuta, "Non-stop outdoor navigation of a mobile robot—Retroactive positioning data fusion with a time consuming sensor system—," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.,* Aug. 1995, pp. 130–135.

[11] Y. Matsumoto, M. Inaba, and H. Inoue, "Visual navigation using view-sequenced route representation," in *Proc. IEEE Int. Conf. Robot. Automat.,* Apr. 1996, pp. 83–88.

[12] P. Moutarlier and R. Chatila, "Incremental free-space modeling from uncertain data by an autonomous mobile robot," in *Proc. IEEE/RSJ Int. Workshop Intell. Robots Syst.,* Nov. 1991, pp. 1052–1058.

[13] A. J. Muñoz and J. González, "Localizing mobile robots with a single camera in structured environments," in *Proc. World Automat. Congress—Robot. Manufact. Syst.,* May 1996, pp. 539–544.

[14] J. Neira, J. Horn, J. D. Tardós, and G. Schmidt, "Multisensor mobile robot localization," in *Proc. IEEE Int. Conf. Robot. Automat.,* Apr. 1996, pp. 673–679.

[15] T. Ohno, A. Ohya, and S. Yuta, "Autonomous navigation for mobile robots referring pre-recorded image sequence," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.,* Nov. 1996, pp. 672–679.

[16] K. Onoguchi, M. Watanabe, Y. Okamoto, Y. Kuno, and H. Asada, "A visual navigation system using a multiinformation local map," in *Proc. IEEE Int. Conf. Robot. Automat.,* May 1990, pp. 767–774.

[17] B. Schiele and J. L. Crowley, "A comparison of position estimation techniques using occupancy grids," *Robot. Auton. Syst.,* vol. 12, pp. 163–171, 1994.

[18] K. Sugihara, "Some location problems for robot navigation using a single camera," *Comput. Vision, Graph. Image Process.,* vol. 42, no. 1, pp. 112–129, Apr. 1988.

[19] T. Tsubouchi and S. Yuta, "Map assisted vision system of mobile robots for reckoning in a building environment," in *Proc. IEEE Int. Conf. Robot. Automat.,* Mar.–Apr. 1987, pp. 1978–1984.

[20] S. Yuta, S. Suzuki, and S. Iida, "Implementation of a small size experimental self-contained autonomous robot—sensors, vehicle control, and description of sensor based on behavior—," in *Experimental Robotics II,* R. Chatila *et al.*, Eds. Berlin, Germany: Springer-Verlag, 1993, pp. 344–358.

A Phase Management Framework for Event-Driven Dextrous Manipulation

James M. Hyde and Mark R. Cutkosky

Abstract—Multifingered robotic hands have not yet made significant inroads into practical applications, partly due to the complexity of dextrous manipulation tasks, and also due to control software shortcomings. High level task controllers exist, as do low level grasp controllers, but neither of these fully address the problems of changing kinematic and dynamic constraints that arise during a grasping and manipulation task. This paper develops a framework, utilizing phases, events, and transitions, that bridges the gap between high- and low-level control. Results from constraint handling and transition experiments conducted with a two-fingered hand are included.

Index Terms— Dextrous manipulation, events, phases, task planning, transitions.

I. INTRODUCTION

Although multifingered robotic hands have been available for over a decade, their use has been largely confined to a few research laboratories. One reason for the slow progress in putting such hands to work is the overall complexity of dextrous manipulation, as viewed from the standpoints of kinematics, dynamics, sensing, control, and planning.

It has been observed by a number of robotics researchers (e.g., Allen [1], Brock [2], Howe [8], Ricker [23], and Speeter [28]), as well as by physiologists studying human manipulation (Johansson

J. M. Hyde is TenFold Corporation, Draper, UT 84020 USA (e-mail: jhyde@10fold.com).

M. R. Cutkosky is with the Center for Design Research, Stanford University, Stanford, CA 94305-2232 USA (e-mail: cutkosky@cdr.stanford.edu).