Real Time Vision for Measuring Pipe Erosion

Donald Bailey, Matthieu Jones, Liqiong Tang
School of Engineering and Advanced Technology
Massey University
Palmerston North, New Zealand

Abstract—**The applicability of using structured lighting for measuring pipe erosion is discussed. Processing the high resolution images on an FPGA significantly reduces the data that is required to be transferred or recorded. The FPGA is able to implement the required image analysis directly on the images as they are streamed from the camera.**

Keywords–pipe inspection; FPGA; image processing

I. Introduction

Many wastewater pipes are made of concrete because of its relatively low cost and ease of installation. One of the main limitations of concrete, however, is its susceptibility to acid erosion. Erosion occurs when the pH of the wastewater is below 6.5, and the acidic water dissolves the cement compounds. Severe erosion can be a major problem because is damages the structural integrity of the pipe, and may lead to the contents of the pipe leaking into the surroundings. In some cases, the bottom of the pipe can be completely eroded away.

Unfortunately, inspecting pipes for erosion is often difficult as they are often underground, and not easily accessible. Repair or replacement are often expensive, and can require extensive excavation. Therefore, it is important to assess the extent of any erosion and severe damage before such drastic actions are taken. Inspection must be performed from within the pipe, and can be facilitated by a robotic platform manoeuvring through the pipe.

Most wastewater pipes are not run full. Therefore, it is the bottom surface of the pipes that is most frequently in contact with the water, and virtually all of the damage caused by acid erosion occurs on the bottom of the pipes. Unfortunately, this also complicates the design of the robotic platform as the bottom of the pipe is where the robot will run. While appropriate design of the robot will enable it to travel on the roughened eroded surface, this can make assessment of the extent of the erosion more difficult.

This problem was overcome by running the robot on the pipe walls, rather than along the bottom of the pipe. This provides a stable reference from which accurate measurements may be made of the extent of the damage. A detailed description of the robot platform is beyond the scope of this paper; the focus of this paper is on measuring the location and depth of any erosion present.

Simple visual inspection (for example by transmitting a video feed back to an operator) if frequently used to assess the pipe condition. However the subjective nature of such assessments has spurred much research into the automated processing of the resulting video data (see for example [1-3]). Many of these visual inspection techniques are looking for more than just erosion. This simplifies they complexity of the processing in this application.

When measuring erosion , two outputs are of particular importance. Firstly, it is necessary to measure the depth of any erosion. This should be to approximately one millimetre of accuracy, or better, to enable the severity of any damage to be assessed. Secondly, it is also desired to map the position and extent of erosion along the pipe. This can be of lower accuracy: one to two millimetres laterally, and 5-20 millimetres along the length of the pipe. (Note that crack detection requires a resolution of better than 0.5mm per pixel [1].)

While there are several methods of measuring the erosion, an image processing method was selected because of its speed and simplicity. Laser based profiling is commonly used for this purpose. In [1], the laser scanner consisted of a single laser beam scanned using a rotating mirror, which was synchronised with a video camera to capture the image as a ring. Images were transmitted via an umbilical to a processing unit which extracted the profile in real time. In [2], a laser ring projector was used, with the camera positioned beside the projector. The position and intensity of the reflected profile were analysed by a neural network classifier to detect defects. Again, the image data was transmitted via an umbilical for processing.

This paper describes the laser stripe imaging arrangement used, and the implementation of it on a robotic platform. The novel contributions of this paper are the analysis of different structured light stipe configurations in the context of pipe inspection, and the real-time hardware implementation of the analysis algorithm on an FPGA platform on the robot itself, which removes the need for the umbilical to transmit raw video data back to an operator.

The outline of the remainder of this paper is as follows. Section II describes the physical setup of the image system and compares the merits of a number of different structured light imaging topologies. Section III presents the software image processing algorithm for calculating the erosion at each point. Section IV develops a parallel implementation of this algorithm on an FPGA platform. Finally section V concludes the paper.

II. Measuring Erosion

Structured lighting is a well known technique for measuring the height of objects, dating from the early 1970s [4-6]. The basic principle is to use triangulation between an active illumination source and a sensor. The scene is illuminated by a

978-1-4577-0330-0/11/$26.00 ©2011 IEEE

laser stripe, which is imaged by a camera. Any deviation in height above or below a reference plane results in a displacement of the corresponding point in the image, as shown in Figure 1.

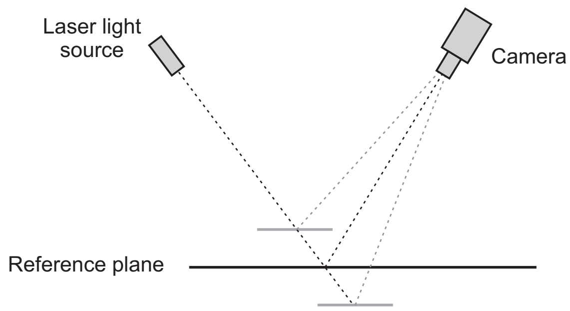

Figure 1. Structured light

The laser stripe is arranged perpendicular to the direction of motion, so that as the robot travels down the pipe, the region illuminated scans along the pipe, enabling the three-dimensional profile of the pipe to be captured.

Conventional structured light systems generally work with a planar reference (for example, a conveyor belt). In a pipe, however, the curvature of the surface adds a complicating factor. To determine the effect of offsets, it is first necessary to determine where the laser stripe will appear within the image. Let the coordinates be defined as in Figure 2, with x along the direction of the pipe, y across the width of the pipe, and z vertically. Also, let the pipe radius be r, and the angles of the light stripe and camera be φ and θ, respectively. Assume that the laser stripe along the bottom of the pipe (at the origin) is imaged to the centre of the field of view of the camera.

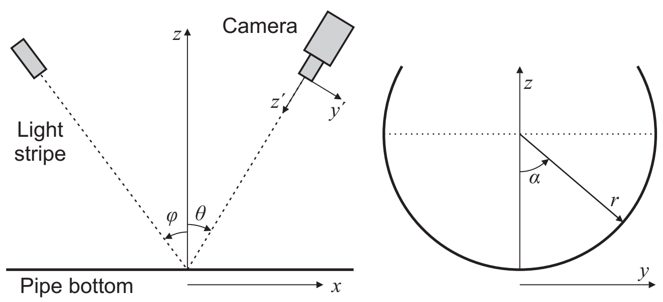

Figure 2. Coordinate system and definitions

The intersection of the laser stripe with the pipe wall will form an ellipse. If the position along this ellipse is parameterised by angle α then the points along the laser stripe will be

$$\begin{aligned} y &= r\sin\alpha \\ z &= r(1-\cos\alpha) - d \\ x &= -z\tan\varphi \end{aligned} \tag{1}$$

where d is the vertical depth of any erosion.

To determine where these points would lie within an image, it is convenient to transform these points into camera centred coordinates. For this, it is necessary to subtract off the camera position, and rotate in the x-z plane to align the axes with the optical axis of the camera. The required transformation gives

$$\begin{aligned} x' &= y \\ y' &= (x - x_C)\cos\theta - (z - z_C)\sin\theta \\ z' &= -(z - z_C)\cos\theta - (x - x_C)\sin\theta \end{aligned} \tag{2}$$

where the camera is located at $(x_C, 0, z_C)$ and

$$x_C = z_C \tan\theta \tag{3}$$

Substituting (1) and (3) into (2) gives

$$\begin{aligned} y' &= (r\cos\alpha - r + d)(\tan\varphi\cos\theta + \sin\theta) \\ z' &= (r - r\cos\alpha - d)(\tan\varphi\sin\theta - \cos\theta) + z_C / \cos\theta \end{aligned} \tag{4}$$

Finally, a pinhole model of the camera is used to project the light stripe onto the sensor in image coordinates.

$$\hat{x} = f\frac{x'}{z'}, \quad \hat{y} = f\frac{y'}{z'} \tag{5}$$

where f is the effective focal length of the lens (in pixels) and the origin for the image coordinates is defined to be in the centre of the image.

The lateral position within the pipe, y, is clearly mapped to the horizontal position within the image captured. However, since the pipe surface is curved, the mapping is distorted by the scaling of the factor z'. This scale factor is constant if

$$\tan\varphi\sin\theta - \cos\theta = 0 \tag{6}$$

which requires that the principle axis of the camera be perpendicular to the plane containing the light stripe, i.e.

$$\theta + \varphi = 90° \tag{7}$$

Unfortunately, such configurations are impractical because the large angle makes occlusion resulting from the erosion more difficult to handle.

There are two primary configurations of particular interest. One has the camera pointing directly down $(\theta = 0)$ and the other has the laser pointing directly down $(\varphi = 0)$. These two configurations will be compared in more detail.

A. Configuration 1: Camera pointing vertically

This is the typical arrangement within a structured lighting set-up because it minimises the perspective distortion when looking at horizontal flat surfaces. An example of an image captured in this configuration is shown in Figure 3.

Figure 3. Example image captured with camera vertical; $\varphi \approx 30°$.

Note that because the laser stripe is at an angle, the different positions within the image correspond to different positions along the pipe $(x \neq 0)$ because of the curved surface of the pipe bottom.

Substituting $\theta = 0$ and (4) into (5) and eliminating α simplifies to

$$\hat{x} = f \frac{y}{\sqrt{r^2 - y^2} - r + d + z_C}$$
$$\hat{y} = f \tan\varphi \frac{\sqrt{r^2 - y^2} - r + d}{\sqrt{r^2 - y^2} - r + d + z_C} \tag{8}$$

Assuming that the displacement in the denominator (d) is small relative to the other terms, the first line of (8) may be inverted to give the mapping between position in the image and lateral position within the pipe. The presence of the y term in the denominator means that for increasing $|y|$ the magnification increases, because the pipe surface is closer to the camera.

The offset in the image of the laser line resulting from the erosion may be found by

$$\Delta\hat{y} \approx \frac{df \tan\varphi}{\sqrt{r^2 - y^2} - r + z_C} \tag{9}$$

The shift of the line within the image is approximately proportional to the depth of the erosion, although the scale factor is again dependent on the lateral position within the image.

B. *Configuration 2: Laser pointing vertically*

The second configuration positions the laser vertically. This has the advantage that the laser stripe is on a fixed position along the length of the pipe $(x = 0)$. An example image captured in this configuration is shown in Figure 4.

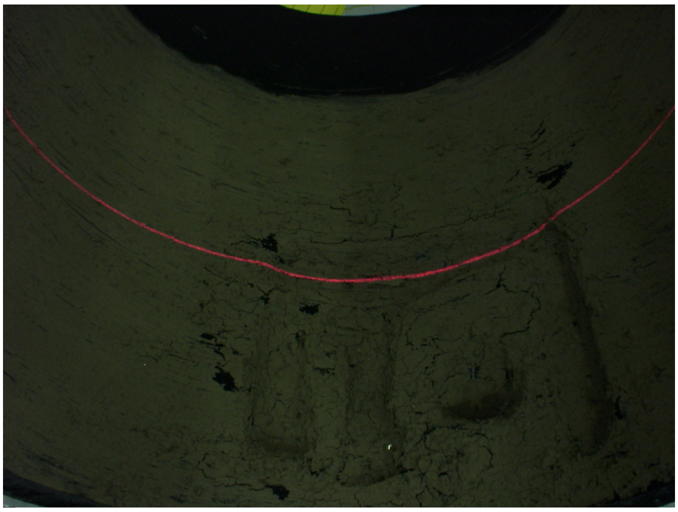

Figure 4. Example image captured with laser stripe vertical; $\theta \approx 40°$.

Substituting $\varphi = 0$ and (4) into (5) and eliminating α simplifies to

$$\hat{x} = f \frac{y\cos\theta}{\left(\sqrt{r^2 - y^2} - r + d\right)\cos^2\theta + z_C}$$
$$\hat{y} = f \frac{\left(\sqrt{r^2 - y^2} - r + d\right)\sin\theta\cos\theta}{\left(\sqrt{r^2 - y^2} - r + d\right)\cos^2\theta + z_C} \tag{10}$$

Again, there is a position dependent magnification of the image. This magnification factor is less than for the previous configuration because the camera is tilted. Within the confines of the pipe, tilting the camera allows the distance from the camera to the bottom of the pipe to be lengthened, which reduces the overall magnification, and increases the field of view.

As before, the shift of the line within the image is approximately proportional to the depth of the erosion:

$$\Delta\hat{y} \approx \frac{df \sin\theta\cos\theta}{\left(\sqrt{r^2 - y^2} - r\right)\cos^2\theta + z_C} \tag{11}$$

with the scale factor dependent on the lateral position within the image.

C. *Comparison of configurations*

In both configurations, the lateral position within the image is not linearly related with lateral position within the pipe. Although the variation is less with the laser pointing vertically, the distortion is such that it needs to be corrected in both cases.

One advantage of the camera pointing at an angle is that the optical path-length can be extended enabling a narrower angle of view to be used to cover the same width across the bottom of the pipe. This has two advantages: firstly, longer focal length lenses are less susceptible to lens distortion, and secondly, the more uniform depth to the laser stripe makes focussing less critical.

A significant disadvantage of the camera pointing vertically is that the position of the laser stripe along the length of the pipe depends on both the angle and any erosion present. This makes obtaining a regularly sampled erosion map more difficult.

Both configurations are prone to occlusion in situations of deep erosion, especially when the angle between the camera and laser stripe is large. However, with the camera looking straight down and the laser at an angle, the occluding surface will often still be visible enabling the line to still be detected. With the camera at an angle, the light stripe may not be visible within the image, and this must be accounted for in the processing.

On balance, the advantages of vertical laser configuration make it the preferred method in this application.

For testing the system (in the remainder of the paper) a Quarton VLM-635-27-LPA line laser diode module was mounted 270 mm from the bottom of a 310 mm internal diameter pipe, and positioned so that the laser stripe was perpendicular to the orientation of the pipe. A Sony DFW-SX900 firewire camera with a 6mm lens was used to capture the initial images (one example is shown in Figure 4). The camera was positioned 280 mm above the bottom of the pipe, and 210mm along the pipe from the laser ($\theta = 37°$). The initial images were captured on a desktop PC and processed using VIPS [7].

III. IMAGE PROCESSING

The image processing to detect the light stripe is relatively simple in this application. Basically, the position of the light stripe in each column must be determined. Then the reference profile (corresponding to no erosion) must be subtracted off, and the result scaled and resampled (to account for the varying magnification across the image).

With the red laser stripe, the maximum of the red component will generally give the position of the stripe. However, in situations where the reflectivity is low, the background may be brighter. The background may be removed effectively by subtracting the green component of the image from the red. This will have little effect on the laser stripe, but will eliminate much of the background, significantly improving the contrast.

The laser line is not a single pixel in thickness. Filtering the image with a 5x1 average filter smoothes the laser line, and results in a better defined peak which can be detected more reliably.

The row number of the maximum pixel value in each column then corresponds to the position of the laser line. Initially, this has simply been found to the nearest pixel, although the position may be found to sub-pixel accuracy if necessary [8]. If the maximum value is below a preset threshold and the detected maximum is to too far from its expected location then the sample is marked as occluded. The detected profile for the image in Figure 4 is shown in Figure 5, with the reference profile overlaid.

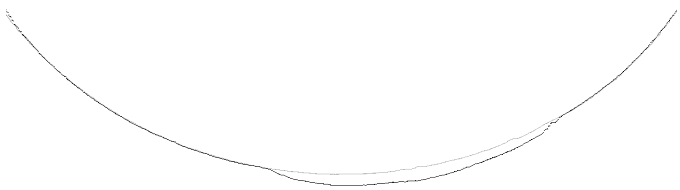

Figure 5. Detected profile from Figure 4, with the reference profile in grey

The reference level is may be subtracted from the detected profile, and the results scaled according to equation (11) to give the profile depth. An example profile is illustrated in Figure 6.

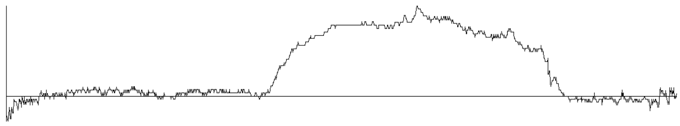

Figure 6. Erosion profile for the image in Figure 4.

IV. FPGA IMPLEMENTATION

While the processing is relatively simple, there is far too much data to process in real time using a low power embedded microcontroller. This may be overcome by processing the images in hardware, using an FPGA.

A. Implementation details

Rather than capture the images into a frame buffer, and process the images from there as was implemented in software, the images may be processed directly as they are streamed from the camera chip. This significantly reduces the memory requirements, and improves the latency since all of the data is available for each frame has been computed by the end of each frame. The processing logic for this is shown in Figure 7, and is described in more detail in the following paragraphs.

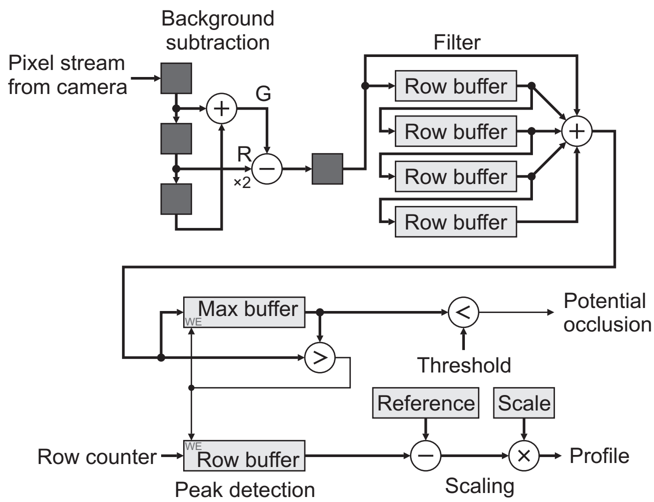

Figure 7. Processing logic for measuring pipe erosion profiles.

Using a CMOS camera, the whole image does not need to be captured. By selecting and reading out only the region of interest, the frame rate may be increased. The data coming from the camera consists of raw pixel values as captured by the camera with the Bayer colour filter array (see Figure 8). Consequently, each individual pixel will only have one colour component. The blue pixels are not required, so only every second line of the input image needs to be processed. The adjacent green pixels can be averaged and subtracted from the centre red pixel. This can be implemented as a simple horizontal filter, with the output register (and subsequent stages) clocked every second clock cycle.

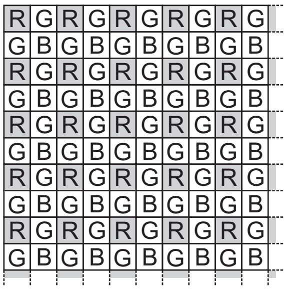

Figure 8. Bayer pattern with the red pixels highlighted.

For the vertical smoothing filter, rather than process each column individually (which would require loading the whole image first), all of the columns can be processed in parallel as the data is streamed in. A 5x1 window requires buffering the previous four rows of pixel values. Rather than average the pixels within the window, the pixels are simply summed. This removes the need to divide by 5.

Next, the current maximum value in each column is determined. Assume that the max buffer has been cleared during the blanking period. Each pixel in the filtered pixel stream is compared with the corresponding column maximum so far (read from the max buffer). If the new value is larger, then the value in the max buffer is replaced, and the corresponding row is recorded in the row buffer. At the end of the frame, the max buffer contains the maximum value in each column, and the row within the image that this occurred on is held in the row buffer.

The final stage, during the vertical blanking, is to convert the position of the laser stripe into an erosion profile. The laser stripe is read out from the row buffer, and the reference profile is subtracted from it. Since the reference profile is fixed for a given configuration, it may be precomputed (using (8) or (10) depending on the configuration) and stored in an on-chip memory. Finally, the difference is scaled to derive the depth in millimetres. Since this scale factor also depends on position, it is precomputed (using (9) or (11)) and stored in an on-chip memory. In parallel with this, the maximum value from the max buffer is compared with a predefined threshold and a flag is generated which indicates potential occlusions. At the same time that the maximum values are being read, the buffer is also cleared for the next frame.

The horizontal axis of the output erosion profile is distorted. If necessary, the output profile can be resampled to correct this. Such resampling can use a simple linear interpolation to give the uniformly sampled profile. At this stage, resampling was not considered to be a requirement for the prototype system.

B. Prototype

The prototype imaging system was implemented using a Terasic DE0 development board [9]. While this is a low cost, entry level board, it has more than sufficient logic resources for this application. This was coupled with a 5 megapixel D5M camera module [10].

The 5 Mpixel camera has a native resolution of 2592x1944 pixels with 12 bits per pixel. Since only the red pixels are of interest, the effective resolution reduces to 1296x972. Operating with column binning (combining the adjacent pairs of columns) halves the horizontal resolution. This is sufficient to give approximately 1 mm lateral resolution along the bottom of the pipe. Vertically, since the laser stripe passes through centre of the field of view, most of the bottom half of the image can be cropped off. Row binning is not used, to maintain the depth resolution at better than 1 mm per pixel offset.

Using windowing on the sensor to capture a 1296x1024 image gives an effective red resolution of 648x512 pixels. With a 96 MHz pixel clock (the maximum supported by the camera), a frame rate of 49 frames per second can be maintained.

The image processing algorithm was implemented using Handel-C (version 5.3), with the resulting net-list mapped to the FPGA using Altera's Quartus II tools (version 11.0 Web Edition).

V. Discussion and Conclusions

The description provided only gives a single profile. To build a map of the pipe erosion, it is necessary to determine the position in the pipe corresponding to the profile. This may be achieved by using a rotation based encoder on one (or more) wheels of the robot. As the robot travels down the pipe, this can trigger the capture and processing of images at regular intervals. The successive scans can simply be stacked to create a two dimensional map of the erosion.

Processing the pixel data directly as it is streamed significantly reduces the processing burden, and effectively reduces the two-dimensional image to a single line of data. The data rate is sufficiently reduced to be able to either stream the image data to flash memory (the DE0 board supports directly writing to an SD card) or to transmit the data wirelessly to an operator for immediate visualisation. (Wireless transmission removes the need for an umbilical cable to convey the data). For wireless transmission, a Zigbee module could be connected via the DE0's RS232 port, with the profile data transferred at up to 17.7 frames per second without compression. (With suitable data compression, this frame rate could easily be doubled.) Operating at this speed with a longitudinal resolution of 10 mm would take less than 6 seconds per metre of pipe inspected and measured.

In this paper we have analysed the different structured lighting configurations for pipe erosion inspection. The configuration with vertical laser and tilted camera is shown to give the best overall characteristics. The image processing algorithm is currently being implemented on an FPGA platform, enabling a map of erosion within the pipe to be built at relatively high processing speeds while maintaining low operating power for battery operation.

REFERENCES

[1] R. Kirkham, P.D. Kearney, K.J. Rogers, and J. Mashford, "PIRAT - A system for quantitative sewer pipe assessment," *International Journal of Robotics Research,* vol. 19, pp. 1033-1053, 2000.

[2] O. Duran, K. Althoefer, and L.D. Seneviratne, "Automated pipe defect detection and categorization using camera/laser-based profiler and artificial neural network," *IEEE Transactions on Automation Science and Engineering,* vol. 4, pp. 118-126, 2007.

[3] W. Guo, L. Soibelman, and J.H.G. Jr., "Automated defect detection for sewer pipeline inspection and condition assessment," *Automation in Construction,* vol. 18, pp. 587-596, 2009.

[4] Y. Shirai and M. Suva, "Recognition of polyhedrons with a range finder," in *International Joint Conference on Artificial Intelligence*, London, England, 1971, pp. 80-87.

[5] P.C. West, "Machine vision in practice," *IEEE Transactions on Industry Applications,* vol. 19, pp. 794-801, 1983.

[6] R.A. Jarvis, "A perspective on range finding techniques for computer vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 5, pp. 122-139, 1983.

[7] D.G. Bailey and R.M. Hodgson, "VIPS - a digital image processing algorithm development environment," *Image and Vision Computing,* vol. 6, pp. 176-184, 1988.

[8] D.G. Bailey, "Sub-pixel estimation of local extrema," in *Image and Vision Computing New Zealand (IVCNZ'03)*, Palmerston North, New Zealand, 2003, pp. 414-419.

[9] Terasic, *DE0 User Manual* vol. Version 1.4: Terasic Technologies, 2009.

[10] Terasic, *TRDB-D5M 5 Mega Pixel Digital Camera Development Kit* vol. Version 1.2: Terasic Techologies, 2010.