Error Analysis of a Real-Time Stereo System

Yalin Xiong
Apple Computer
Cupertino, CA 95014

Larry Matthies
Jet Propulsion Laboratory
Pasadena, CA 91109

Abstract

Correlation-based real-time stereo systems have been proven to be effective in ***applications such as robot navigation, elevation map building etc. This paper provides*** an ***in-depth*** analysis ***of the major error sources for such a real-time stereo system in the context of the cross-country navigation of an autonomous vehicle. Three major types of errors: foreshortening*** error, ***misalignment error and systematic error, are identified. The combined disparity errors can easily exceed three-tenth's*** of ***a pixel, which translates to*** a ***significant amount of range errors. Upon understanding these error sources, we demonstrate different*** ap***proaches to either correct them or model their*** magni***tudes without excessive additional computations. By correcting those errors, we show that the precision of the stereo algorithm can be improved by 50%.***

1 Introduction

Unmanned ground vehicles (UGV) have been under development for various missions ranging from interplanetary exploration, volcanic exploration to hazardous waste disposal. One critical sensing capability of such vehicles is to avoid obstacle for autonomous or semi-autonomous navigation, During the past decade, real-time stereo systems have emerged as the major alternative to achieve such a task other than the laser range finder, which is difficult to use in many missions for various reasons. Two main requirements for such a stereo system are its speed and precision. It must be fast enough to guide the vehicle in real-time, and precise enough to detect any far-away obstacle so that the vehicle has enough time to steer around it. The correlation-based stereo method [3, 4] can be as fast as several pairs per second, and it delivers reasonably precise results.

With the current state-of-art system as described in the next section, we believe that, in order to achieve even higher percentage of correct obstacle detections in the context of autonomous navigation, we need a better understanding of the errors in stereo than regarding them uniformly as consequences of image noises, and to correct those errors if possible without excessive additional computations. This paper presents our analysis of three major error sources in such correlation-based real-time stereo systems. All these three types of error are significant in magnitude and yet none of them is caused by image noise.

The foreshortening errors result from the fact that the 3D scene is not fronto-parallel. Therefore, for any small patch in the left image, its corresponding patch in the right image is not only translated but also distorted. The correlation-based stereo method models the translation but not the distortion. To measure the distortion effects on disparity values, we compute a 'Foreshortening sensitivity map", which represents how sensitive the disparity value at any pixel is with respect to the magnitude of the foreshortening. Such a measure is very useful in determining the reliability of disparity values. In order to correct the errors caused by the foreshortening effects without changing the stereo algorithm, we can pre-warp the right image such that the foreshortening is zero for an arbitrary 3D plane instead of the fronto-parallel planes. For the task of vehicle navigation, we can set such a plane to be the ideal ground plane. In practice, such a simple method usually results in an 80% reduction of the foreshortening errors.

Precise stereo calibration has been extensively researched during the past decade [2]. The problem with the outdoor and cross-country navigation is that extensive mechanical vibrations and rough terrain can often perturb camera parameters, therefore, render the pre-calibrated stereo rig either partially or even completely out of calibration. The second major error source is the misalignment errors caused by stereo jig being slightly out of calibration. We will not explore the self-calibration approach since self-calibration for lens with significant distortion is beyond the scope of this paper. Instead, we first try to model the effects of misalignment by computing a "misalignment sensitivity map", which measures the sensitivity of the dis-

parity values with respect to the misalignment. Such a measure is also very useful in determining the reliability of disparity values. In order to correct the misalignment errors, we model the misalignment field by a low-order bivariate polynomial, and then pre-shift the right image vertically in the opposite direction. We show that such a simple approach usually results in 50% reduction of misalignment errors.

The third error source for the correlation-based stereo is the systematic error which is the consequence of window effects and quadratic approximation in the subpixel registration. For typical images used in cross-country navigation, the magnitude of the systematic error is between 0.05 and 0.10 pixels. We test a quartic fitting algorithm to replace the quadratic interpolation. The quartic approach appears to reduce the systematic errors by about 0.03 pixels. Without drastic change to the simple correlation algorithm and the amount of computation, we believe the systematic error together with errors from image noises represents the upper limit of the precision that the real-time stereo system can achieve. We also outline several approaches to alleviate the systematic errors if a significant amount of additional computation is feasible.

2 Background

Figure 1 shows the block diagram of a real-time stereo system for obstacle detection. The image rectification process utilizes the stereo calibration information to correct radial lens distortion and align two images along scanlines. Image pyramids for both left and right images are built. The integer disparity values are computed by search along the scanlines using a 7 x 7 window. The objective function is the Sum-of-Squared-Difference (SSD):

$$M(d) = \sum_{x=-3}^{3}\sum_{y=-3}^{3}\left(I_l(x,y) - I_r(x-d,y)\right)^2 , \quad (1)$$

where I_l is the left image, and I_r is the right image.

The SSD values are evaluated discretely along scanlines of the right image, and the integer disparity values are where the SSD values are the smallest. Once the smallest SSD value So is identified, the two adjacent SSD values S-1 and S_{+1} on the left and right side of So are used together with So to approximate the SSD curve by a second-order polynomial. The subpixel adjustment on the integer disparity value is the minimum of the quadratic curve:

$$\text{A d} = \frac{S_{-1} - S_{+1}}{2(S_{-1} + S_{+1} - 2S_0)}. \quad (2)$$

To achieve robustness, we also employ many additional techniques such as smoothing, statistical outlier detection, blob filtering, consistency checking and other heuristic methods. Since most of these ad hoc techniques are irrelevant to this paper, we will not specify the details of these technique. Interested readers may refer to [3, 4].

In the latest demo in Fort Hood of the stereo-guided navigation, the vehicle successfully traversed 2 kilometers autonomously. Despite the success, we believe that the system needs to be more reliable in detecting even smaller obstacles. And one of the major options to achieve that goal is to improve the precision of the range image from the stereo system without excessive additional computations. Given the current configuration of the stereo system, a disparity error of 0.1 pixels results in a range error of 14 centimeters at 10 meters. And we need to detect obstacles of 41 centimeters high (i.e. the axle clearance of the HMMWV) at that distance. The three major error sources presented in this paper usually generate a combined error in a magnitude of around 0.3 pixels, which is substantial for the obstacle detection purpose. In the next three sections, we will look into each of the three major error sources in the real-time stereo system, and propose various approaches to either correct or model them.

3 Foreshortening Error

The foreshortening error results from image deformations unmodelled in the original SSD formulation in Eq. 1. When the right image I_r is not only translated but also deformed with respect to the left image I_l, minimizing the SSD function $M(d)$ will not yield the correct disparity. Such a deformation is caused by non-zero disparity gradients, which, for rectified stereo pairs, is caused by the 3D surface not being fronto-parallel. Depending on the image texture and the amount of deformation, the estimated disparity d will contain errors of various magnitudes. In the UGV navigation task, since the surface of the terrain is rarely fronto-parallel to the camera pointing direction, there is always a similar deformation. In this section, we will first model quantitatively the foreshortening errors for terrain imagery, then reduce those errors by pre-warping the right image so that the equi-disparity plane coincides with the ideal ground plane. Similar pre-warp method was also proposed in [5, 1].

In the SSD minimization in Eq. 1, if the disparity values have significant gradient in the column (vertical) direction, which is usually true for stereo images shot from the top of a vehicle, the deformation between the left and right image can be modeled as:

$$I_r(x,y) = I_l(x + d_0 + ay, y), \quad (3)$$

where a is the non-zero disparity gradient in the col-

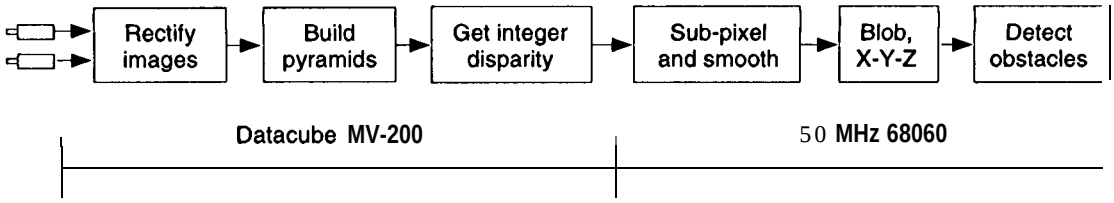

Figure 1: Main Blocks of a Real-Time Stereo System for Obstacle Avoidance

umn direction.

Taking the Taylor expansion of the above equation, we have

$$I_l(x + d_0 + ay, y) \approx I_l(x, y) + (d_0 + ay)\frac{\partial I_l}{\partial x}. \quad (4)$$

Replacing it into Eq. 1 and simplifying, we then have

$$M(d) \approx \sum_{x=-3}^{3}\sum_{y=-3}^{3}\left((d - d_0 - ay)\frac{\partial I_l}{\partial x}\right)^2, \quad (5)$$

the minimum of which is

$$\hat{d} = d_0 + ae_f, \quad (6)$$

where the foreshortening sensitivity e_f is defined as:

$$e_f = \frac{\sum_{x=-3}^{3}\sum_{y=-3}^{3} y\left(\frac{\partial I_l}{\partial x}\right)^2}{\sum_{x=-3}^{3}\sum_{y=-3}^{3}\left(\frac{\partial I_l}{\partial x}\right)^2}. \quad (7)$$

The foreshortening sensitivity e_f measures at every pixel location the ratio between the magnitude of disparity error and the magnitude of the disparity gradient which causes the error. We call the image of these sensitivity values "Foreshortening Sensitivity Map". And it can be computed very efficiently.

Figure 2 shows a typical image of a nearly flat terrain. Figure 3 shows the foreshortening sensitivity map of the terrain image. The sensitivity values are quantized uniformly bet ween [-3, 3]. Based on the stereo jig configuration and the image resolution at 240 x 256, the magnitude of vertical foreshortening of the rectified stereo images is always around 1/8, i.e., the disparity increases from about O to 30 pixels in 240 rows. Figure 4 shows the image resulting from horizontally shifting the image in Figure 2 by an increasing disparity from top to bottom. The corresponding true disparity map has a vertical gradient of 1/8. We then feed this synthetic stereo pair to the real-time stereo system to measure the disparity errors caused by this foreshortening.

Figure 5 shows the errors caused by the foreshortening at every pixel and its correlation with the foreshortening sensitivity values. The line in the correlation plot is the ideal linear relation predicted in

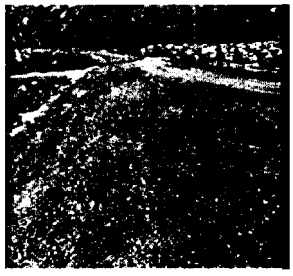

Figure 2: A Terrain Image

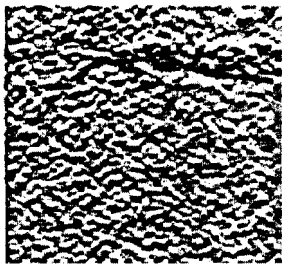

Figure 3: Foreshortening Sensitivity Map

Eq. 6. The errors are quantized uniformly between [–3/8, 3/8]. We conclude that the foreshortening sensitivity map formulated in Eq. 7 is an accurate model of sensitivity to disparity gradient at every pixel.

We also measure the statistics of the all the sensitivity values in order to have an idea about the average magnitude of the errors caused by the foreshortening. Figure 6 shows the histogram of the foreshortening sensitivity map. The mean is about zero, and the

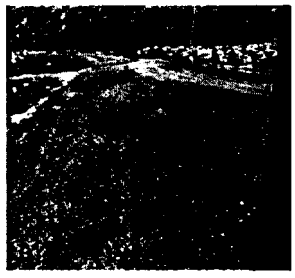

Figure 4: Synthetically Warped Image

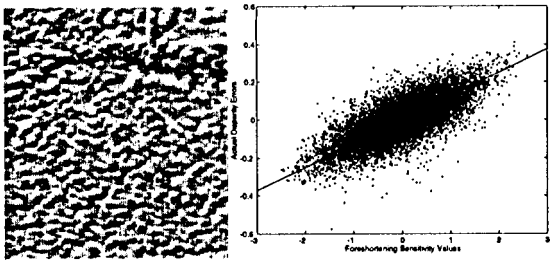

Errors Correlation

Figure5: Errors caused by Foreshortening and Correlation with Foreshortening SensitivityValues

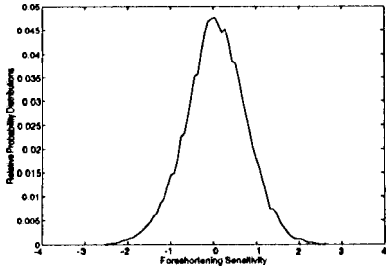

Figure 6: Distribution of Foreshortening Sensitivity Values

standard deviation is 0.71. If the expected disparity gradient is around 1/8, these statistical measures predict that the foreshortening error has a zero mean, and a standard deviation of about 0.09 pixel.

Given such a substantial magnitude of the average errors caused by foreshortening, we attempt to correct them. Instead of having fronto-parallel planes be equi-disparity, we can pre-warp the image such that the ideal ground plane is equi-disparity as we did in Figure 5. Figure 7 is the real right image taken together with the left image in Figure 2. Figure 8 shows side-by-side the disparity maps resulting from the original stereo algorithm and the pre-warp method. Note that the disparity gradient in both maps has been factored out, and the disparity maps are quantized uniformly between [1, 4] pixels. Though we do not have ground truth to quantify how much improvement the prewarp method has introduced, the reduction of small ripples in the disparity maps is obvious. The residual disparity gradient is about 0.01 to 0.02 in both the vertical and horizontal directions. Such a reduction of the disparity gradient results in more than 80% of reduction in the foreshortening and consequent foreshortening errors.

As a summary, we showed in this section that the foreshortening causes significant errors, and the foreshortening sensitivity map accurately models how sensitive a disparity value at a particular location is with respect to the foreshortening. For the vehicle navigation purpose, we can use the pre-warp method in order to minimize the foreshortening in stereo images. The results show 80% improvements in precision.

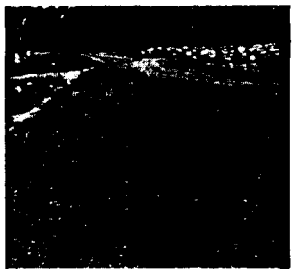

Figure 7: Right Image of the Same Terrain

Original algorithm Foreshortening Correction

Figure 8: Disparity Maps

4 Misalignment Error

When there is a certain amount of misalignment after the rectification, it will introduce additional errors. In practice, calibrating a stereo jig so that the misalignment is less than a tenth of a pixel is challenging. Even if it can be calibrated to such a high precision, it can rarely survive the severe vibrations of the vehicle and the rough terrain without being slightly out of calibration. In this section, we will look into this misalignment problem and try to mode] and correct the disparity errors caused by misalignment.

In the SSD minimization in Eq. 1, if there exists a small vertical misalignment m, the relation between the left and right images can be modeled locally as:

$$I_r(x, y) = I_l(x + d_0, y + m), \qquad (8)$$

where, again, d_0 is the true disparity value.

Taking the Taylor expansion of the above equation, we have

$$I_l(x + d_0, y + m) \approx I_l(x, y) + d_0 \frac{\partial I_l}{\partial x} + m \frac{\partial I_l}{\partial y} \qquad (9)$$

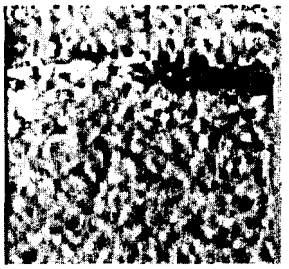

Figure 9: Misalignment Sensitivity Map

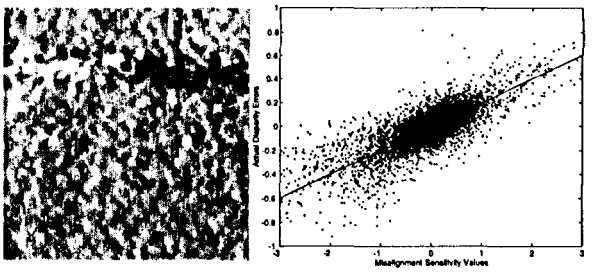

Errors Correlation

Figure 10: Disparity Errors Caused by Misalignment and Correlation with Sensitivity Values

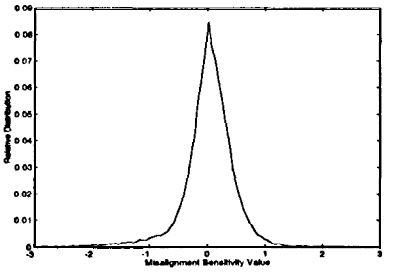

Figure 11: Distribution of Misalignment Sensitivity Values

Replacing it into Eq. 1, we have

$$M(d) \approx \sum_{x=-3}^{3}\sum_{y=-3}^{3}\left((d-d_0)\frac{\partial I_l}{\partial x} - m\frac{\partial I_l}{\partial y}\right)^2, \quad \text{(lo)}$$

the minimumof which is

$$\mathbf{d} = d_0 + m e_m, \quad (11)$$

where the misalignment sensitivity value e_m is defined as:

$$e_m = -\frac{\sum_{x=-3}^{3}\sum_{y=-3}^{3}\frac{\partial I_l}{\partial x}\frac{\partial I_l}{\partial y}}{\sum_{x=-3}^{3}\sum_{y=-3}^{3}\left(\frac{\partial I_l}{\partial x}\right)^2}, \quad (12)$$

The misalignment sensitivity e_m measures at every pixel the ratio between the magnitude of the disparity error and the magnitude of the misalignment which causes the error. We call the image of these sensitivity values "Misalignment Sensitivity Map", which again can be computed very efficiently.

For the left image shown in Figure 2, its misalignment sensitivity map is shown in Figure 9. The sensitivity values are quantized uniformly in [-3, 3]. We synthetically shifted the same image vertically by 1/5 pixel, and feed the pair to the real-time stereo system to generate the disparity map. Figure 10 shows the disparity error caused by the misalignment and the correlation between the errors and the misalignment sensitivity values. The line in the correlation plot is the ideal linear relation predicted in Eq. 11. The errors are quantized uniformly in [-3/5, 3/5]. Overall, the misalignment sensitivity map predicts fairly accurately the disparity errors caused by the misalignment.

Figure 11 shows the histogram of the misalignment sensitivity map. The mean value is zero, and the standard deviation is 0.51 pixel. Therefore, for a misalignment ranging from 0.1 to 0.2 pixel, the average error caused by the misalignment for road imagery ranges approximately from 0.05 to 0.10 pixel.

[1] **The errors also include the systematic errors presented in the next section, which may explain the differences between what is predicted by the misalignment sensitivity map and the actual disparity errors.**

Since the errors caused by misalignment are significant, it is desirable to minimize the magnitude of the misalignment. The traditional way to achieve this goal is to calibrate the stereo jig more precisely. Unfortunately, for reasons cited before, achieving even higher precision in calibration is both difficult and ineffective. Self-calibration may be the only option when the stereo calibration is completely out of calibration, i.e. the magnitude of the misalignment is more than a pixel. But for most cases, the magnitude is well within half a pixel. Therefore, we can directly correct the misalignment in images instead of indirectly compensating it through modifying the camera parameters.

Before we adopt the approach to directly correct misalignment in images, we need to experimentally verify that the misalignment field is consistent, i.e., it does not change from one frame to next. If it is approximately constant, we then only need to calibrate the misalignment field occasionally to keep our mode] up-to-date. Otherwise, calibrating the misalignment field for every stereo pair will be too much computation, and further analysis is needed to determine the cause of the dynamic misalignment.

In order to calibrate the misalignment field, we augmented the original SSD search in Eq. 1 to include a vertical search in the misalignment calibration mode. In fact, the vertical search only needs to be ± 1 pixel.

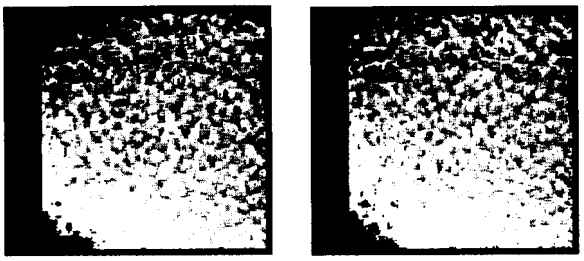

Figure 12: Misalignment Fields

In the subpixel registration in Eq. 2, we fit a bivariate quadratic equation to a nine-point neighborhood to compute the subpixel shift in both the horizontal and vertical directions. Note that in the experiments shown in this section, we have already applied the foreshortening correction technique presented in the previous section. Figure 12 shows the misalignment fields for the stereo pair (Figure 2 and Figure 7) and another stereo pair shot in the same trip. It shows that the misalignment fields are fairly consistent. For both pairs, we have significant misalignments on the left side of the image. The maximal misalignments in lower-left and upper-left corners are roughly ± 0.4 pixels respectively. And the overall standard deviation of the vertical shift is 0.15 pixels[2].

Given that the misalignment fields are consistent, we can model them by a bivariate polynomial. In practice, we find that a 4th order bivariate polynomial can model the misalignment fields very well:

$$a_0 + a_1 x + a_2 y + a_3 x^2 + a_4 xy + a_5 y^2 + a_6 x^3 + a_7 x^2 y + a_8 x y^2 + a_9 y^3 + a_{10} x^4 + a_{11} x^3 y + a_{12} x^2 y^2 + a_{13} x y^3 + a_{14} y^4.$$

The modelling of the misalignment field by using a 2D SSD search and fitting a bivariate polynomial is computationally expensive. But we do not need to do it for every pair. And it is much easier and faster than a full calibration.

After the misalignment field is modeled, we then pre-warp the stereo pair such that the misalignment is canceled. The warp is simply a vertical shift. Figure 13 shows the disparity map for the same stereo pair we used in the previous section. We compensated the misalignment. And the improvement in the lower-left area is very obvious since the misalignment was large in that area. In fact, Figure 14 shows the misalignment field after we have compensated for the misalignment. Statistically, the standard deviation of the misalignment in Figure 14 is 0.10 pixel for the whole image. Given that the systematic error is at least 0.05 pixel, we conclude that the misalignment compensation method reduces the amount of misalignment by at least 50%.

[2] **Again, it also includes the systematic error.**

Figure 13: Disparity Map after Misalignment Compensation

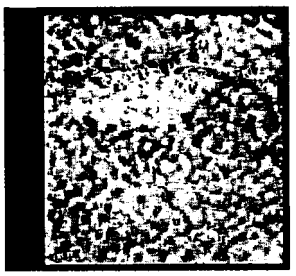

Figure 14: Misalignment Field after Misalignment Compensation

5 Systematic Error

The third type of error source in the real-time stereo system is what we called the "systematic errors", which are caused by the implicit assumptions in the algorithm. There are two kind of systematic errors, which are significant:

- ***Window Effect Errors:*** When the window mask slides on the right image, every time we move the window to the right by one pixel, the leftmost column is moved out and the rightmost column is moved in. With a finite window, the difference in image content by replacing the leftmost column by the new rightmost column is random. Therefore, there is certain amount of randomness in the SSD values. When the SSD values are interpolated, such a randomness causes errors in disparity.

- ***Linearization Errors:*** The implicit assumption in the quadratic interpolation is that the right image can be linearized using Taylor expansion around the true disparity. Unfortunately such an approx-

imation is poor when there is a significant amount of high frequency information.

Since the window effect errors are caused by the discrepancy between the leftmost column which is moved out and the rightmost column which is moved in, we have the following two observations:

1. The errors have zero-mean because the discrepancy should have a zero-mean distribution.

2. When the image is composed of information only at the harmonic frequencies of the window width, the window effect errors should be zero. In other words, if the frequency of the information is always an integer times the inverse of the window width, the moved-out leftmost column is exactly the same as the moved-in rightmost column.

Let us suppose that an image is composed of one single sinusoid at frequency f which is an integer times the inverse of the window width, and the true disparity is d. Therefore, using the Parsavel's theorem, we have SSD values at $x = -1, 0, 1$ as

$$S_{-1} = \| e^{-jf} - e^{jfd} \|^2, \quad (13)$$

$$S_0 = \| 1 - e^{jfd} \|^2, \quad (14)$$

$$S_1 = \| e^{jf} - e^{jfd} \|^2, \quad (15)$$

where the magnitude of the sinusoid is omitted since it is the same for all three SSD values. In this case, the only systematic error is the linearization error.

After fitting a parabola to the three SSD values above, the estimated disparity $\hat{d}$ in Eq. 2 is:

$$\hat{d} = 2 \frac{\tan(fd)}{\tan(f/2)}. \quad (16)$$

Figure 15 shows the true disparity d versus the estimation bias $\hat{d} - d$. The different curves are for different frequency $f = +2.0, +1.5, \pm 1.0, \pm 0.5$, with the outmost corresponding to $f = \pm 2.0$, and the flattest corresponding to $f = \pm 0.5$.

Of course, an image contains information at frequencies ranging from $-\pi$ to π. But the different biases generated by information at different frequencies do *not* cancel each other since they all bias toward the same direction. As a result, the linearization error is in fact a bias toward integer disparity values. The more high frequency information the image contains, the larger the magnitude of the bias is.

We test the real-time stereo system on stereo pairs with uniform disparity and zero misalignment, i.e. both the foreshortening and misalignment errors are

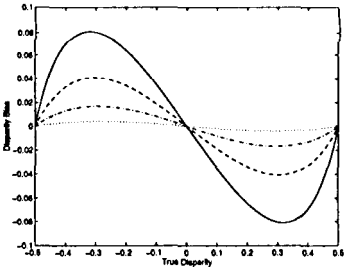

Figure 15: Bias for different Frequency

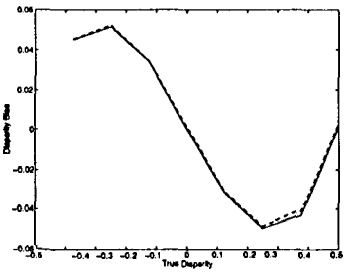

Figure 16: Systematic Bias from the Experiments

zero. We translate the whole terrain image horizontally by the following amount: –3/8, -1/4, $-1/8$, O, 1/8, 1/4, 3/8, 1/2 pixel. For every pair, we run the real-time stereo system on it, and compute the statistical distribution of the disparity values. The second set of experiments is the same as the first one except that we also add in the right image a white noise whose standard deviation is 3.0 out of 256 greylevels.

Figure 16 shows the disparity bias versus true disparity for the two sets of the experiments. And Figure 17 shows the standard RMS errors, i.e. the zero-mean errors, for the same sets of experiments. The solid curves are from the first set, and the dashed curves from the second set.

From these synthetic examples, we draw the following conclusions:

1. The linearization error is a bias modeled by Eq. 16. The maximal bias for terrain imagery

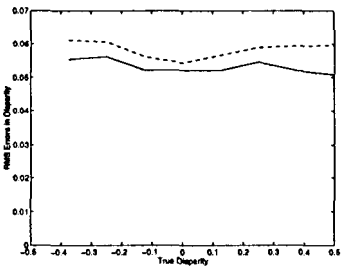

Figure 17: RMS Errors from the Experiments

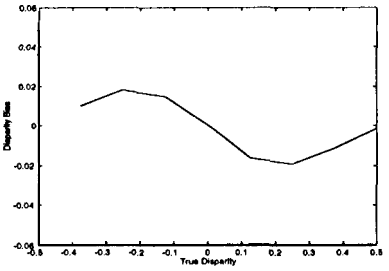

Figure 18: Bias using Quartic Curve Fitting

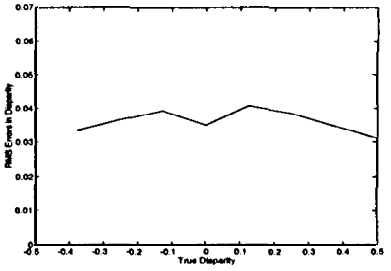

Figure 19: RMS Errors using Quartic Curve Fitting

is about 0.04 to 0.05 pixel, and is reached around $\pm 1/4$ from an integer disparity value.

2. The window effect error is zero-mean with its standard deviation for terrain imagery about 0.05 pixel.

3. The effect of image noises is relatively insignificant. The RMS errors show only a very small amount of increase when image noises are present.

One way to reduce the linearization errors is to extend the Taylor expansion to quadratic terms. In other words, the SSD curve ***M(d)*** should be approximated aa a quartic instead of a quadratic curve. We represent the quartic curve as

$$Q(x) = c_0 + c_1 x + c_2 x^2 + c_3 x^3 + c_4 x^4. \qquad (17)$$

In order to fit the SSD values against the quartic curve, we need five SSD sample values: S_{-2}, S_{-1}, S_0, S_1, and S_2. The minimum of the quartic curve is one root of its derivative function. Since we already have a good subpixel solution from the original quadratic approximation, we can use it as an initial estimate of the minimum. Usually the minimum is reached in a couple of iterations using the Newton-Raphson method.

Figure 18 and Figure 19 show the bias and RMS errors for the same stereo pair by using the quartic curve fitting. We can see that the quartic curve method reduces both the linearization error (bias) and window effect errors (RMS) by about 0.03 pixels.

There are a number of other approaches which can reduce the systematic errors significantly, though they usually require a significant amount of additional computations:

- ***Sub-pixel Sampling*** The magnitude of the bias is proportional to the sampling distance in the SSD function. If we can sample the SSD values by half-pixel distance instead of one pixel distance, the magnitude of the bias can be cut by half.

- ***Larger Window:*** Since the window effect errors are the direct consequence of the discrepancy of the moved-out leftmost column and moved-in rightmost column, increasing the window size will decrease the effects of the discrepancy simply because there are more columns.

6 Summary

In this paper, we have analytically identified the three major sources of disparity errors in a real-time stereo system. For each error sources, we first quantify its magnitude for a typical terrain image pair and create a mathematical model. Based on the models of the error sources, we propose various simple and effective approaches to reduce their magnitudes without excessive additional computations. By applying the techniques presented in this paper, we demonstrated that the combined disparity errors can be reduced from about 0.3 pixels to 0.10 to 0.15 pixels.

Acknowledgments

The authors would like to thank Todd Litwin in JPL for helping with the datasets and programs.

References

[1] **P.** Burt, **L. Wixson and G.** Salgian, **Electronically directed "focal" stereo, in** *ICCV'95*; **pp. 94-101.**

[2] **D. Gennery, Least-squares camera calibration including lens distortion and automatic editing og calibration points, in Calibration** *and* **Orientation of Cameras in** *Computer* **Vision, Springer-Verlag, 1993.**

[3] **L. H. Matthies and P. Grandjean, Stochastic performance modeling and evaluation of obstacle** detectability **with imaging range sensors, in** *IEEE Transactions on* Robotics **and Automation, 10(6), December 1994, pp. 783-791.**

[4] **L. Matthies, A. Kelly, T. Litwin, and G. Tharp, Obstacle detection for unmanned ground vehicles: a progress report,** *in Robotics Research: The Seventh International Symposium*, **Springer, 1996, pp. 475-486.**

[5] **Lynn H. Quam, Hierarchical warp stereo, in** *DARPA Image Understanding Workshop, 1984*, **pp. 149-155**