THE UNIVERSITY OF
WESTERN AUSTRALIA

Faculty of Engineering, Computing and Mathematics

TERRAIN MAPPING WITH A TRACKED ROBOT

David Wells (10119598)

Supervisor: A/Prof Thomas Bräunl

June 2006

Letter to the Dean

David Summerfield Wells
33 Shann Street
Floreat, Western Australia, 6014

5th June 2006

The Dean
Faculty of Engineering, Computing and Mathematics
The University of Western Australia
35 Stirling Highway
Crawley, Western Australia, 6009

Dear Sir

I submit to you this dissertation entitled “Terrain Mapping with a Tracked Robot” in partial fulfilment of the requirement of the award of Bachelor of Engineering.

Yours faithfully

David Wells

Acknowledgements

I would like to thank the University of Western Australia, specifically the Faculty of Engineering, Computing and Mathematics for the opportunity to undertake an interesting and fulfilling project and for providing an exciting university course.

This project would not have been possible without the supervision and organisation provided by Associate Professor Thomas Bräunl and the associated administrators of the Mobile Robot Lab. I would like to express my gratitude for the staff involved in making my project possible.

Many thanks to my parents, Dr Fred Wells and Ms Patricia Summerfield, for supporting me throughout my university career and providing essential feedback on my dissertation.

Finally, I would like to thank all my friends who helped me develop ideas and critically analyse my project. This was essential in increasing the depth and quality of my tracked robot project.

Abbreviations

GPS	Global Positioning System
LCD	Liquid Crystal Display
NASA	National Aeronautics and Space Administration
PSD	Position Sensitive Device
RAM	Random Access Memory
ROM	Read Only Memory
RS232	Recommended Standard - 232
TPU	Time Processor Unit

Abstract

The objective of this thesis was to perform and analyse a method of terrain mapping with a tracked robot. Many robots already have this functionality using different sensors and techniques for obtaining altitude measurements. The tracked robot achieves this using an inclinometer, for measuring the pitch of the robot and an optical encoder, for measuring the distance travelled.

Advanced robots have a vast array of very expensive sensors which provide accurate altitude readings relative to their application. These sensors include altimeters and GPS receivers. The tracked robot project investigated the feasibility and methodology for performing terrain mapping using an inclinometer and an encoder.

Using these sensors, altitude calculations proved to be sufficiently accurate and actually produced a better terrain map than what would be produced using more expensive devices mentioned above.

With the inclusion of altimeters or GPS receivers in larger scale robots, the applicability and area of operation for the robot were brought into question. Traditionally, these robots operate in a different environment to the small, enclosed, area in which the tracked robot was tested. The tracked robot's method proved to be more practical and accurate for this application.

The outcomes of this project demonstrated the worthiness of the finalised method for obtaining a terrain map. The result of this project produced a better way of performing terrain mapping, in small areas, compared to that of the methods used in more expensive robots. The cost effectiveness and simplicity of the tracked robot were key factors in determining the practicality of the mapping method.

Table of Contents

List of Figures

List of Tables

List of Equations

1. Introduction

Mobile robots have attracted huge amounts of interest over the past decade and will continue to do so well into the future. One of the main reasons for their growing popularity is the powerful microprocessors which are now available. These were previously considered impossible to produce [12]. The potential benefits of these processors have barely been investigated but are likely to be very significant. The American National Aeronautics and Space Administration (NASA) is at the forefront of designing and producing robots that utilise many of the capabilities of mobile robots. These robots are extremely well-publicised and are very exciting to see in operation. The objective of these planetary robots is to explore a designated planet and obtain as much environmental information as possible [11]. One of these functions is to obtain a terrain map of the designated area. This functionality exists in many types of robots and it is critical that this operation is performed as efficiently as possible, due to the high expense of development and production. This project will investigate an approach to obtaining a terrain map using a tracked robot fitted with an inclinometer and optical encoder and determine where this method may be applicable, if at all.

The key objective of this project is to obtain a meaningful representation of a physical area by mapping its terrain. This representation will describe the area using a surface plot which can be easily interpreted by humans. This surface plot should resemble the actual area itself with features being easily recognisable whilst providing as much accuracy as possible. The terrain map must also be available in a format that can be used by various robots, enabling them to determine a path for traversing the terrain. To produce the terrain map, the robot must ascertain the altitude of each cell within the area.

In order to obtain the terrain map, the robot must traverse each segment within the area and calculate the altitude as the robots passes these segments. To reach every segment possible, the robot will drive in a lawnmower pattern. The robot will start in a corner and then drive forward until a wall is reached. At this point, the robot will turn around and drive parallel to the previous lap, shifted over by the width of the robot. This will continue until the entire terrain is traversed.

Many aspects of the robot, whilst in motion, will affect the accuracy of the terrain map produced. Ensuring the robot drives in straight lines and performing turns that leave the robot perpendicular to the wall, are essential in guaranteeing that the calculations performed by the robot are done with a high degree of accuracy. This will allow the focus of the errors to be on the inclinometer and its characteristics whilst traversing over different features of the terrain.

The performance of the terrain mapping method will allow a judgement to be made on whether or not the method is viable. This decision will be based on the accuracy of the terrain map and the effectiveness of the method itself. Driving the robot in a lawnmower pattern may not produce the best results for terrain mapping. Controlled factors such as this must be scrutinised in order to increase accuracy. Although certain components of the robot's method and design may be undesirable, the method could be very useful in the areas for which the robot is designed. Factors such as simplicity and cost would also influence a decision on the applicability of this terrain mapping method.

2. Background

This section is a brief overview of mobile robots that perform terrain mapping and navigation through an unknown terrain. Conducting this review before designing or developing the tracked robot is essential, as it provides a strong knowledge base against which decisions and designs can be compared. It also assists in knowing what to expect throughout the course of the tracked robot research. Generally, the focus will be on planetary robots as they perform much the same function as that required of the tracked robot. Other forms of robots exist, such as legged robots, but these robots are of lesser relevance to this project.

As with planetary robots, the tracked robot must be able to traverse many types of surfaces without much difficulty [1]. The reason that tracks are used is because off-road performance is crucial in being able to perform terrain mapping. The performance of the robot is based on a number of factors. The first of these is its ability to traverse the terrain with a minimal amount of slippage on unknown surfaces [1]. Slippage can cause the robot to lose track of its location because the rate at which wheels or tracks that are turning is no longer proportional to the robot's actual speed [2]. This means that readings taken from the rotational velocity of the drive shaft to calculate the speed are misleading. Slippage can also affect the robot's ability to traverse the terrain. For example, if there is not enough traction the robot may be unable to climb slopes. The robot would then have to find another path to enable it to reach the desired location. The robot may even become bogged which for planetary exercises is extremely undesirable. Tracked robots distribute their weight over a greater surface and tend to be positioned higher on soft surfaces, such as sand, compared to wheeled robots that sink into the sand [2]. Greater traction and reduced slippage make tracked robots superior in negotiating rugged, hostile and unknown areas.

The increased surface area contact of tracked robots over wheeled robots provides greater stability over various surfaces [5]. When traversing corrugated surfaces, the tracks of the robot will tend to locate themselves on top of the ridges and troughs.

This reduces the amount of vertical, oscillating motion of the robot and its suspension components, leading to smoother navigation of the terrain. The advantage of this is reduced probability of becoming immovable in troughs, and reduction in the amount of shock and vibration on the components of the robot. These reductions increase the lifetime of the robot and its individual components.

The manoeuvrability of tracked robots is also greatly beneficial when negotiating unknown terrain. These robots are able to turn 360° on the spot simply by driving the tracks on either side of the robot in opposing directions. This enhances the manoeuvrability of the robot greatly, as the robot can avoid obstacles and modify its path with ease [1]. Wheeled robots can also perform this function but they may require more motors, increasing the amount of control required in order to complete the operation [7]. This manoeuvre gives tracked robots superiority over wheeled robots due to its simplicity.

For the purposes of terrain mapping, both planetary and terrestrial robots store the information in similar ways. These robots will often have a limited amount of information of the area they are mapping before entering. For planetary robots, this is done using satellites or other space vehicles orbiting the planet before landing a robot to perform terrain mapping, along with many other functions [11]. This allows the robot to function in a certain area for which the robot is designed. This area can be divided into cells and topographic information, then recorded into each of these cells [4]. This is usually done using a simple two dimensional array representing the area, with each array element signifying a cell. For each element, data relating to the terrain can be stored and for highly advanced robots, soil samples, wind speed and other measurements can be taken and stored.

The tracked robot project will be focusing on the altitude of the cells in relation to the starting point. With these ideas in mind, design and implementation of the tracked robot may commence.

3. Robot Design

The tracked robot, as the name suggests, employs the use of tracks to traverse the terrain in a similar fashion to a military tank. The manoeuvres are performed in conjunction with an EyeBot controller with a range of sensors and actuators.

3.1 EyeBot Controller

Figure 3.1.1: EyeBot Controller

The EyeBot controller is the backbone of the robot, as it provides an interface to all the associated devices with a reasonably easy-to-use operating system known as RoBIOS (version 6.4). The controller contains a 32-bit microprocessor with a clock rate of 33MHz along with 512KB ROM and 2048KB RAM. The controller accepts many input types, often for specific sensors and actuators, but also includes digital and analogue inputs. The outputs are controlled in a similar manner. There is a Liquid Crystal Display (LCD) connected for displaying the controller's status which also functions as a form of output for the programs developed for it. Assembly or C programming languages may be used.

3.2 Sensors

Position Sensitive Devices

Position sensitive devices (PSDs) are commonly used on mobile robots as they provide the distance from the device to the object directly in front of them. These

Figure 3.2.1: PSD

devices produce a raw reading which is used for calibration. The calibrated values are then stored in an array in millimetres, so that the calibrated distance measurement matches the corresponding raw value. These devices are quite prone to error and care must be taken when using them.

There is an effective range when using these types of PSDs of approximately 10cm to 35cm, from the PSD to the object it is detecting. At distances less than 10cm, the reading from the PSD is unpredictable, but increases in reliability up to 30cm. The accuracy of the PSD degrades as the distances increases from around 30cm, with a very dramatic increase in error around 40cm and above. The effective range of the PSD is rather dependent on the specific PSD because they all produce their own performance characteristics. This is why it is important for each PSD to be calibrated with prior knowledge of the device, so that caution can be exerciesed when using the information they provide.

Inclinometer

The inclinometer is a small device for measuring the inclination relative to the horizontal plane. On the tracked robot, its function is to measure the pitch of the robot so that it can determine the angle of the slopes that the robot drives on. The output of the inclinometer is simply read from an analogue input on the EyeBot. These values must then be calibrated to the actual inclination or declination angle of the robot.

Figure 3.2.2: Inclinometer

The accuracy of the inclinometer itself, for the purposes of this robot, is satisfactory but it is prone to error from other sources. When the robot is in motion, the reading from the inclinometer can be misleading because of the turbulence caused by the robot's motion. This causes the inclinometer to produce erroneous inclination values even when the robot is traversing over a flat, horizontal surface. These types of errors are mainly due to the robot's movements and not the inclinometer. Measures can be taken to decrease the impact with knowledge of where these errors arise.

Optical Encoder

Figure 3.2.3: Optical Encoder

The optical encoder is a device for measuring the number of revolutions of a rotating shaft or axle. In this case, it measures the number pulses produced as the motor's drive shaft rotates. The number of pulses is proportional to the distance travelled by the robot. A simple calculation, using the number of pulses, can then be made after a calibration factor has been obtained. This encoder produces 128 pulses per revolution which is an extremely high degree of resolution. For this robot, the number of pulses per centimetre is 187 ±1 pulses, meaning that the calculation of the distance travelled by the robot is excellent if there is no slippage between the tracks and the surface it is travelling over.

3.3 Actuators

Motor

Figure 3.3.1: Motor

The tracked robot has a single motor for powering the tracks on either side through a differential. The speed of the motor is controlled by a regulator that produces a voltage level relative to the input into the regulator. Control of the regulator is similar to that of a servo where there is a specified input range from 0 to 255. Values above 128 will drive the robot forward and values below 128 will drive the robot in reverse. The speed increases as this value diverges farther away from 128, hence increasing the voltage output of the regulator in either a positive or negative fashion.

Brakes

To enable the tracked robot to turn, the brake is applied to the axle after the differential on the side in which direction the turn is to be made. To turn left pressure is applied to the left-hand brake. Pressure is applied by pulling the brake discs together which is done using a servo. Rotating the servo in one direction will apply braking on one side. Rotating the servo in the other direction will apply pressure to the brake on the opposite side. This means that braking can only be done on one side at a time.

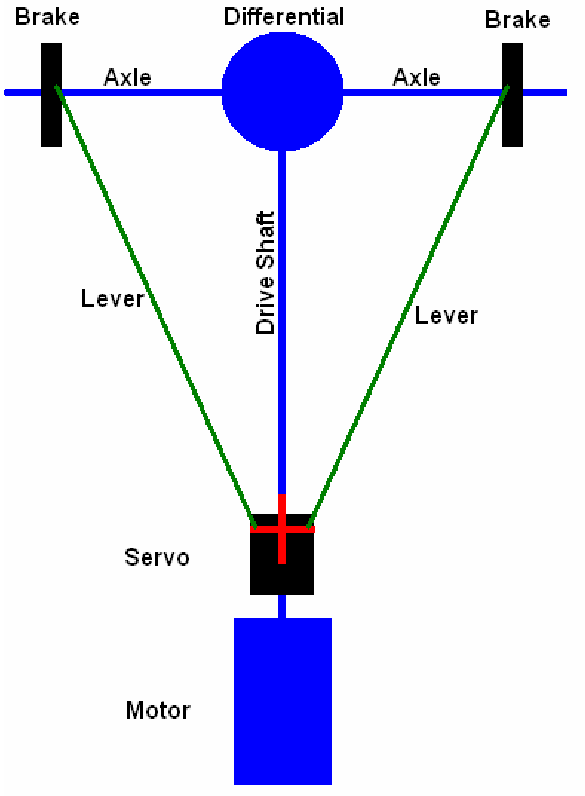

Figure 3.3.2: Braking System

3.4 Other devices

Remote Control

A universal remote control is used for simple operations on the tracked robot. These operations include starting the terrain mapping process, stopping the robot in case of emergency, and selecting options once the mapping has been completed. The remote control is only used for emergencies and selecting options, not controlling the movements on the robot itself as that is fully autonomous.

Figure 3.4.1: Remote Control

4. Terrain Mapping

The objective of the terrain mapping that the tracked robot is designed to perform, is to produce a representation of the area which it is traversing. This representation shows the altitude of each segment of the enclosed area by means of a graphical picture or simply storing the information in memory for other functions.

To perform the mapping of the terrain, the robot must manoeuvre over the entire area, enabling it to calculate the inclination of the slopes within the area. The inclination angles are then used to compute the height at points along the slopes, as well as on the flat surface, producing a two-dimensional array with each array element storing the height at these points. The height calculated is relative to the starting position.

The area used for testing is approximately a 3x3 metre area with walls enclosing it. Within the walls are various slopes and platforms that the robot negotiates in order to perform that terrain mapping.

4.1 Traversing the Terrain

As previously discussed, the robot must traverse the entire area enclosed by walls to be able to perform terrain mapping. The tracked robot manoeuvres in a lawnmower pattern as shown in Figure 4.1.1. This pattern enables the robot to drive over every segment within the area given it can reach any segment from any other segment in a straight line. The process used to drive the robot in this pattern is shown in Figure 4.1.2. The turning process is an extremely important part of the robot's function because the turns must be made so that the robot is perpendicular to the wall once the turn is complete. This means that the robot can then drive forward after the turn, in a straight line parallel to the all the other lines. This stops the robot from drifting over other segments and reduces the amount of errors in the production of the array.

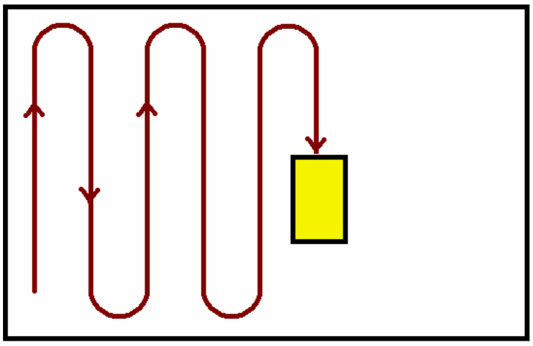

Figure 4.1.1: Lawnmower Pattern

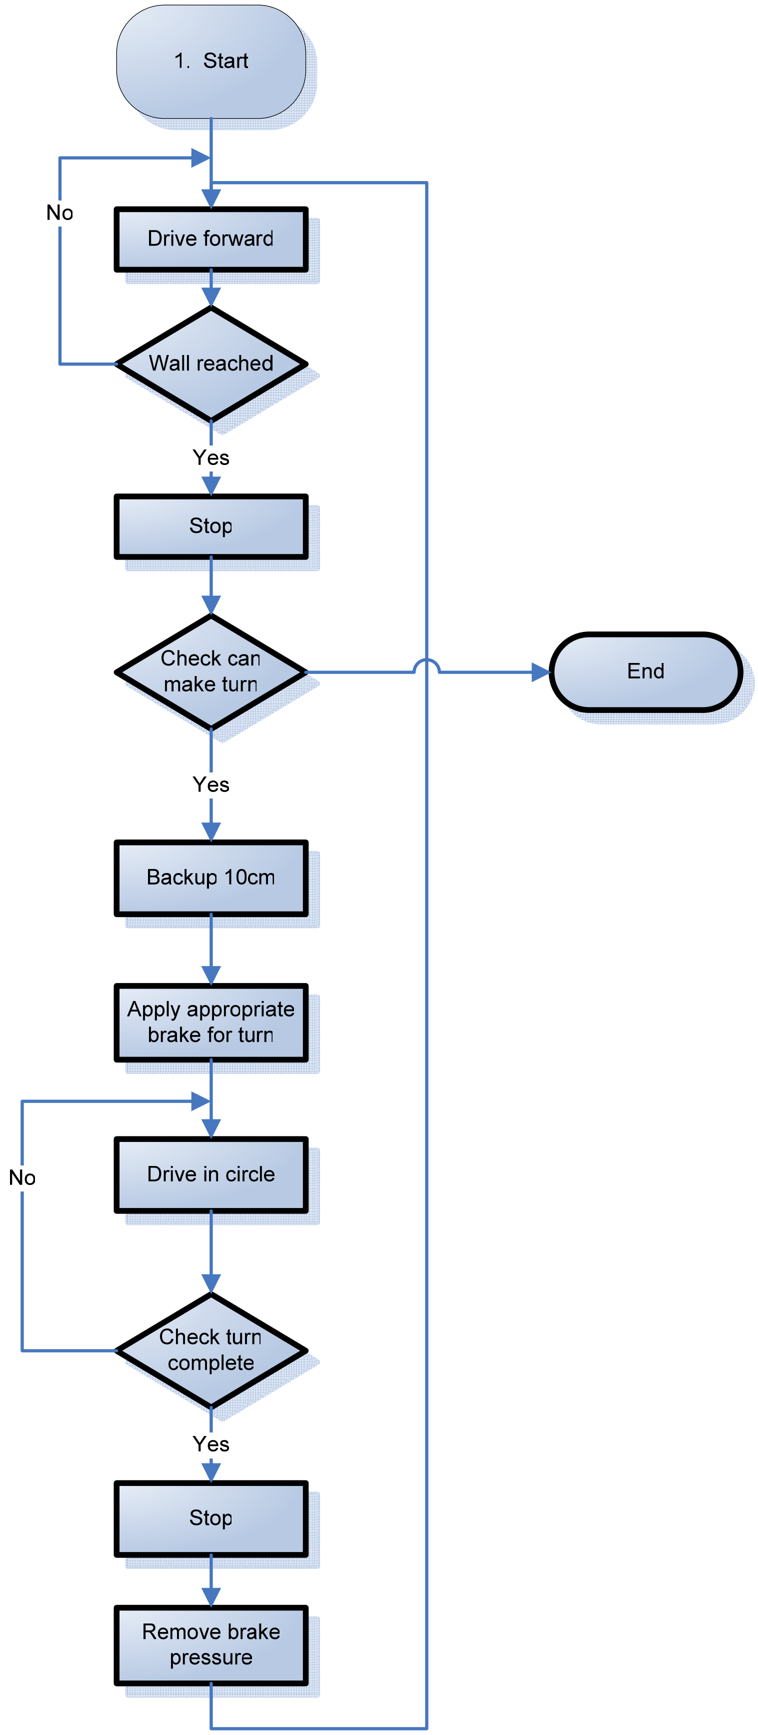

Figure 4.1.2: Turning Algorithm

To begin with, the robot starts in a corner of the enclosed area and drives forward, as shown in Figure 4.1.1. As the robot moves forward, the front position sensitive device (PSD) is used to detect the distance to the wall in front of the robot. When that distance reaches a threshold value, the robot stops.

Once the robot has stopped, it reads the PSD on the side of the robot that relates to the direction that the robot will turn. For example, if the robot is going to turn right, the PSD located on the right hand side of the robot is read. This simply checks for objects or walls that may stop the robot from being able to perform the turn. If a wall is detected, then the robot must have traversed the entire area. If nothing is detected, then the robot reverses 10cm, allowing enough room for the turn to be completed.

The robot now has enough room to make the turn, so the brake is applied on the track that is closest to the apex of the turn. That is, if the robot were to turn right, then the right hand brake is applied. Power to the motor is supplied so that the turn can be performed.

As the robot is turning, the PSD located on the back of the robot is determining the distance from the back of the robot to the wall behind it. As the robot turns, the value that this PSD reads decreases until the robot is perpendicular to the wall, at which point the value would start increasing. When this point of decreasing values turning to increasing values is detected, the robot will stop and the turn will be complete. Since the PSDs are subject to many errors, such as noise, the encoder measuring the distance travelled is used to ensure that the robot makes a significant portion of the turn before relying on the PSD for accurate measurements. When the robot arrives at a wall, it will be almost perpendicular to the wall meaning that the encoder can be used to determine how far through the turn the robot is because each turn is the same and the start and end points of the turn are similar for each turn. This ensures the PSD is only being used when the distance is within the effective range of the PSD.

Once the turn is complete, the brake pressure is removed and then robot drives forward to the next wall and repeats the process until it detects a wall before performing the final turn.

4.2 Calibration of the Inclinometer

The inclinometer must be calibrated in order to produce the most accurate measurements possible. To accurately calibrate the inclinometer, various tests were performed. Firstly, the robot was tested in a stationary state on horizontal, inclined and declined surfaces. The average reading from each of these tests was then mapped to the actual angle of the slope in degrees. This was simply done using a standard linear equation to find the relationship between the inclinometer's raw reading and the actual angle in degrees. This assumes the zero value, the raw value of the inclinometer on an incline of 0°, to be the raw reading of the inclinometer on a horizontal surface.

This zero value was then adjusted in realisation that the robot's front end lifts as the robot drives. Various tests of the robot in the dynamic state were performed by examining the inclinometer's readings as the robot accelerates and reaches cruising speed. The robot was then driven over a mound in order to methodically test various zero values based on the results of the dynamic tests. Analysing the performance of the altitude calculations obtained in the mound tests would then provide the best zero value for maintaining accuracy and reducing error accumulation.

4.3 Recording Data

As the tracked robot is traversing the enclosed area, it must store the calculated information for uploading to a computer after completion of the mapping or store it on the EyeBot allowing the robot to use this information to perform other tasks. The data is stored in a two dimensional array with the columns indicating the lap number and the rows indicating the segment along that lap. A lap is the movement from one wall to the opposite wall in a straight line. The array elements store the altitude computation for the each segment on each lap.

Each lap is divided into equal-sized segments so that calculation of the altitude can be made at each junction between segments along the lap. The length of each segment in this case is 4cm, but this can easily be changed depending on how much detail about the terrain is required and whether or not the processor on the EyeBot can handle the additional computations.

By reducing the size of each segment, the frequency in which the altitude calculations are made also increases. This significantly increases the number of computations and operations that must be performed within the same timeframe, putting extra pressure on the microprocessor, which may lead to errors becoming apparent. Since the robot is still in motion while the computation is being made, if the computation of the previous altitude has not finished when the robot arrives at the next point, then the calculation for the new point may be missed leaving gaps within the lap. This can be combated by reducing the speed of the robot but this means that it takes much more time for the robot to complete the terrain mapping exercise. This trade-off between the resolution of the map and the time to complete the mapping must be examined depending on the overall purpose of the robot. For the tracked robot working in the testing environment, using 4cm segments produces a good mapping of the terrain.

Another problem with added pressure on the EyeBot controller, with the increase of segments per metre, is that other devices also struggle. The readings from the PSDs can become erratic and misleading, causing the robot to perform unpredictably. Other operations on the robot can also perform undesirably but in most cases the robot will still carry out the required task.

Aside from the computation problems, each segment must represent the distance from the start of the segment to the end of the segment along the horizontal plane. To be able to divide up the enclosed area for storage into a meaningful array, each array element must represent the height for each segment over the horizontal plane. When the robot is travelling up or down slopes, the distance the robot travels is greater than the distance it travels on the horizontal plane. This means that determination of the segments must be calculated from the distance travelled by the robot over the horizontal plane. This is a simple calculation as shown in Equation 4.3.1.

$$d_h = d_t \cos\theta$$

Equation 4.3.1

Where d_h is the distance travelled over the horizontal plane

d_t is the distance actually travelled by the robot from start of the segment

θ is the pitch angle of the robot

This means that each time d_h is equal to 4cm, a calculation is made, not when d_t is equal to 4cm. Obviously, when the robot is driving over a horizontal surface, the pitch angle of the robot is zero causing the cosine of this angle to equal one, so that the distance travelled and the distance over the horizontal plane become equal.

Now that the laps are divided into small segments, calculation of the altitude for each of these segments must be carried out. This process is similar to the division of the lap into segments. Each time the distance travelled over the horizontal plane equals 4cm, an altitude calculation is performed for the corresponding segment. Calculation of the altitude is done by computing the vertical distance gained or lost for the current segment and summing it with the previous computed altitude. When the robot begins the terrain mapping, the initial altitude is set to zero and all altitudes calculated are relative to this initial value. The calculation for the height as the robot is traversing the terrain is shown in Equation 4.3.2

$$a_c = a_p + d_t \sin\theta$$

Equation 4.3.2

Where a_c is the altitude of the segment

a_p is the previous segment’s altitude

d_t is the distance actually travelled by the robot from start of the segment

θ is the pitch angle of the robot

Essentially, the robot will calculate the altitude every 4cm along each lap based on the previous altitude computation. Due to the restriction of available sensors, this method of adding the previous altitude to the vertical distance gained or lost is the only method possible. Results of this method will be discussed later.

4.4 Producing the Map

Once the robot has performed the terrain mapping, the two dimensional array storing the altitude of each segment is uploaded to a computer and plotted in MATLAB, using the surface function. This produces a representation of the physical area and looks just like the real area with slopes and horizontal surfaces. An ideal surface plot depicting the testing area is shown in Figure 4.4.1 below.

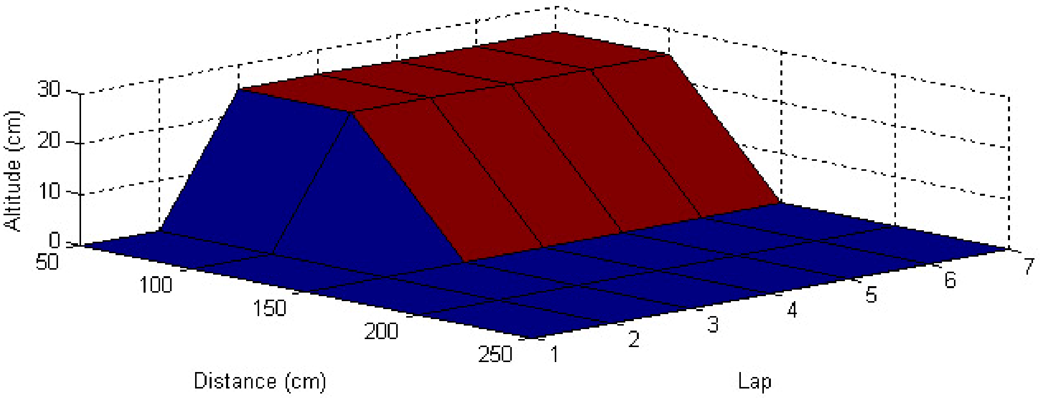

Figure 4.4.1: Ideal Terrain Map

5. Performance Analysis

The readings from the inclinometer are the most important readings with regards to producing accurate terrain representations. For this reason, tests were performed to see how well the inclinometer operates in certain situations. Some of these tests were to see the variance of the inclinometer's raw measurement in both static and dynamic situations. The static tests were performed on an incline, a decline and also on a horizontal surface. The results of these tests show some interesting characteristics of the inclinometer especially in the dynamic tests.

As described in Section 4.2, the results of these tests will be used to calibrate the inclinometer. The static tests will provide points at which the linear calibration of the inclinometer will be made. That is, to determine what is the relationship between the raw inclinometer and the actual angle in degrees. The results from the dynamic tests will then offset the raw inclinometer value that corresponds to zero degrees.

5.1 Stationary Performance

The static tests performed on the robot were used to calibrate the inclinometer. The technical specifications of the inclinometer specify that it is a linear device, so these tests will determine the relationship between the raw inclinometer value and the angle in degrees.

To begin with, the inclinometer was tested on a horizontal surface to see how much the raw reading actually changes over 100 seconds. Readings were taken every tenth of a second and produce the results shown in Figure 5.1.1 shown below. Note that the figure produced includes the angle in degrees for each corresponding raw value which was calculated after performing these tests and was inserted to produce a more meaningful graph.

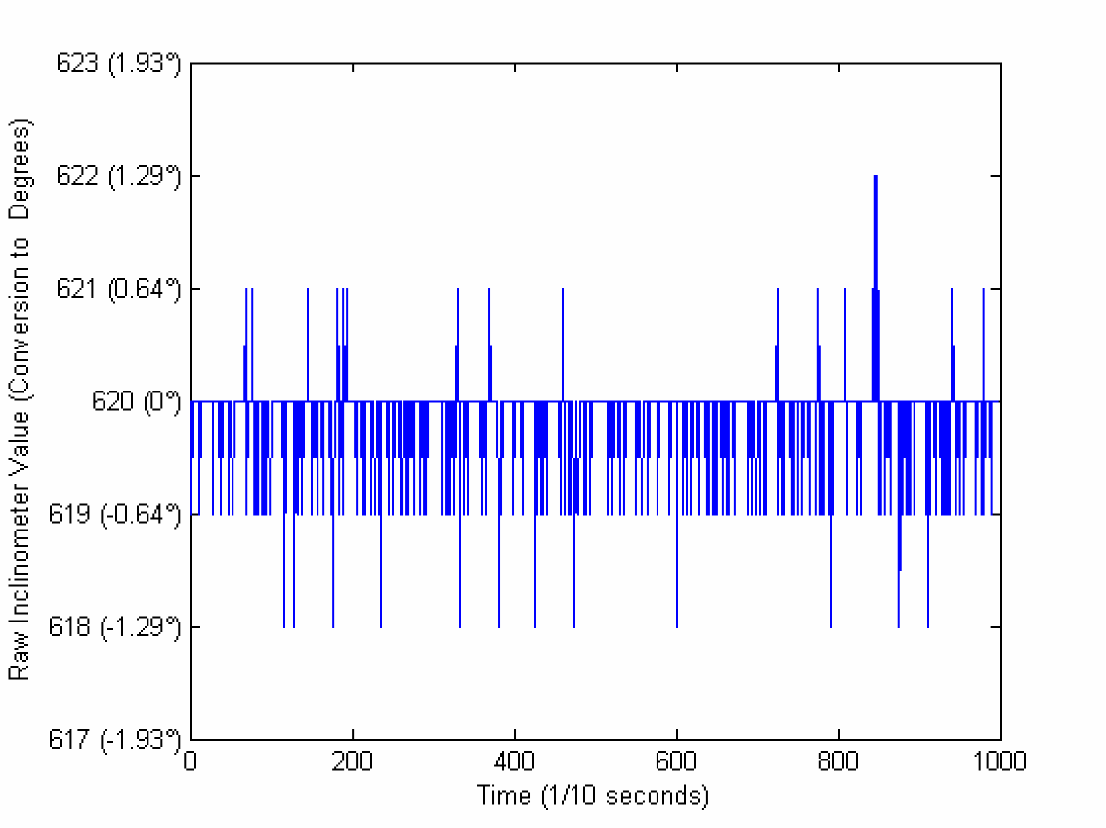

Figure 5.1.1: Stationary Horizontal Inclinometer Test

Figure 5.1.1 shows that when the robot is stationary on a horizontally flat surface, the inclinometer's raw readings do not vary much at all. In fact, the value does not exceed 4 units either side of the average value of 620. It also shows that the reading is equal to the average value 82% of the time and is between 619 and 620 for 97% of the time.

Translating these maximum errors into degrees shows that the inclination does not vary by more than ±1.29°. Since the inclinometer's reading is equal to the average value 82% of the time, -0.64° for 15% of the time and ±1.29° for the remaining 3% of the testing time, then the inclinometer's reading on a horizontal surface is resilient to error. This means that when the robot is on a horizontal surface and is stationary, the reading from the inclinometer is very reliable and does not contribute much error.

Testing of the inclinometer on an incline and a decline produced similar results to those shown for a horizontal surface. The results of the tests on an incline and decline are shown in Figure 5.1.2 and Figure 5.1.3 respectively. Note again that the figure produced includes the angle in degrees for each corresponding raw value which was calculated after performing these tests and was inserted to produce a more meaningful graph.

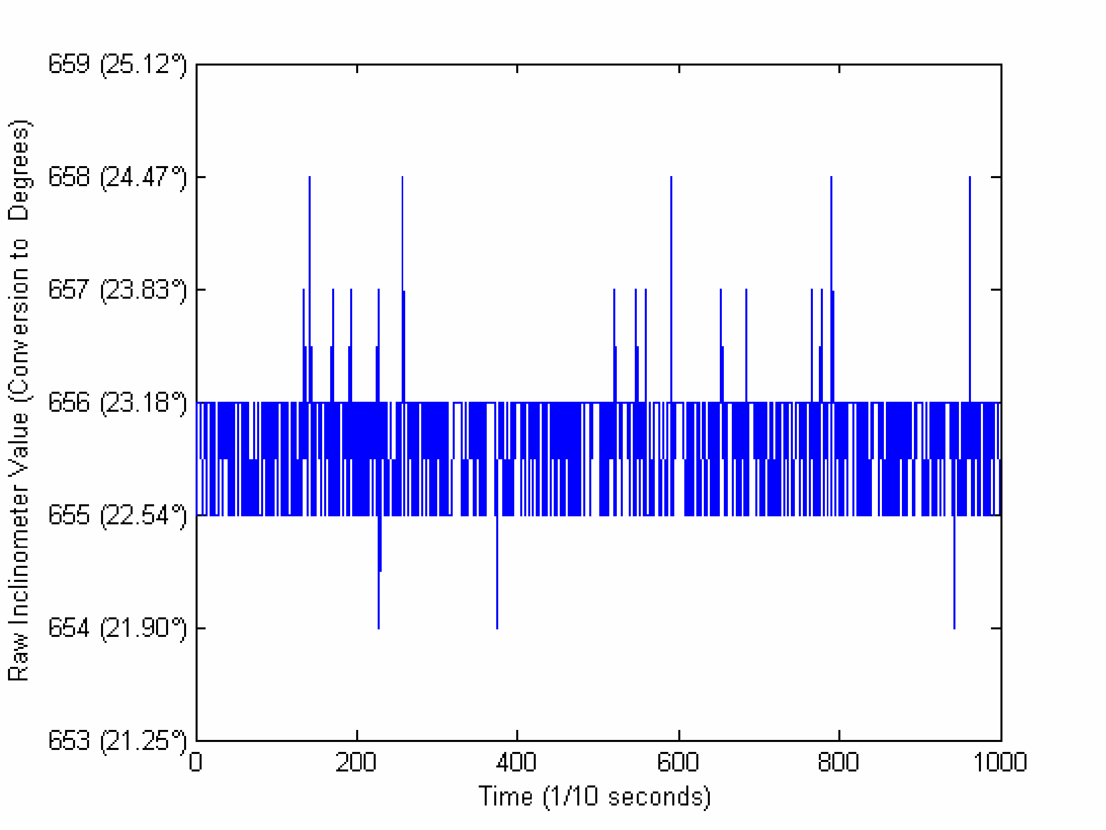

Figure 5.1.2: Stationary Inclined Inclinometer Test

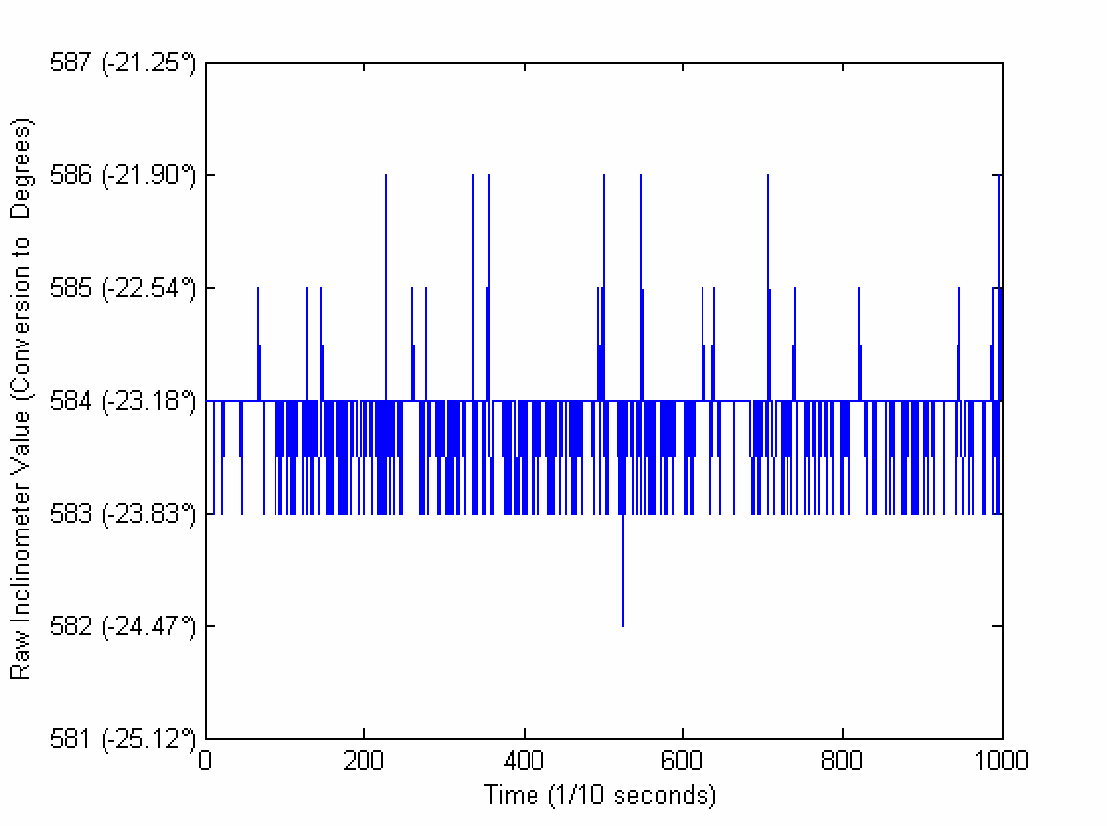

Figure 5.1.3: Stationary Declined Inclinometer Test

The average value for the incline was 656 units and 584 for the decline. For both the incline and the decline, the readings did not vary by more ±2 units which translates to an error of ±1.29°. Again, this is the maximum error that was recorded during the test so this would not occur frequently. Similar to the horizontal test, the percentage of time that the reading was equal to the average was 63% and 80% for the inclined slope and declined slope respectively. Although the inclined slope had a lower percentage of time that the value was equal to the average, 98% of the time the value was either 655 or 656. Table 5.1.1 below shows the standard deviation of the various tests. This demonstrates that the accuracy of the inclinometer in a stationary state is excellent.

Test	**Average**	**Standard Deviation**
Horizontal	620 (0°)	0.45
Inclined	656 (23.2°)	0.45
Declined	584 (-23.2°)	0.53

Table 5.1.1: Stationary Inclinometer Test Statistics

The values shown in Table 5.1.1 can now be used to calibrate the inclinometer with a modified version of the standard linear equation, shown in Equation 5.1.1.

$$\theta = m(r - c)$$

Equation 5.1.1

Where θ is the angle in degrees

m is the gradient

r is the raw inclinometer value

c is the point at which graph crosses the r axis, also known as the zero value

Calculating the gradient is simply done by dividing the difference between the inclined and horizontal raw values by the difference in the corresponding angles. Similarly, this is done with the declined value in place of the inclined value and in this case the gradient is equal to 0.6444. The point at which the graph crosses the r-axis is known from the horizontal test which is 620. Equation 5.1.2 shows the final equation with the values substituted.

$$\theta = 0.6444(r - 620)$$

Equation 5.1.2

This equation is used to calculate the angle using the raw value from the inclinometer. The value for c, in this case is 620, is the raw inclination value for which the robot is horizontal. The dynamic tests will adjust this value to increase the accuracy of the terrain mapping.

For all of the stationary tests, the inclinometer performed exceptionally well, as would be expected. However, these results do not have much bearing of the accuracy of the terrain mapping because all the readings taken during the mapping are taken when the robot is in motion. This means that testing of the inclinometer in this dynamic state is essential.

5.2 Dynamic Performance

The dynamic tests performed on the robot were used to determine how the inclinometer functions under certain conditions. The values read from the inclinometer in these tests will underline the effectiveness of the inclinometer as a means of determining the altitude of each segment. The result will influence many factors in which the robot will undertake its mapping and also adjust the zero value for calculating heights.

Initially, testing of the inclinometer in the dynamic state was performed simply by recording the raw value for a period of 10 seconds. For the first 3 seconds, the robot would be stationary and then would drive forward for the remaining 7 seconds over a horizontal surface. This test would show what the reading should be for the surface and then show how the inclinometer is affected when the robot accelerates and reaches its cruising speed. The results of this test are very interesting and are shown in Figures 5.2.1 through to 5.1.4 below. Note that the angle in degrees has been calculated using Equation 5.1.2 found in the static tests, without any adjustments from the dynamic tests.

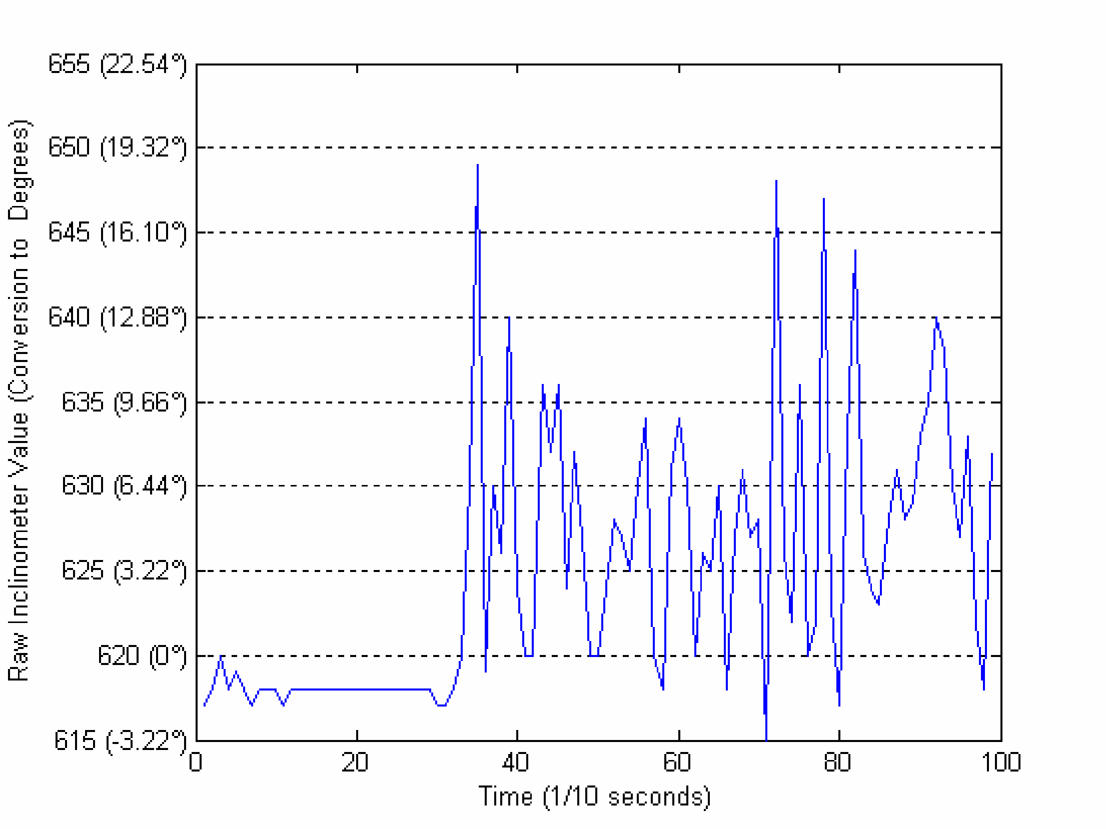

Figure 5.2.1: Dynamic Horizontal Test 1

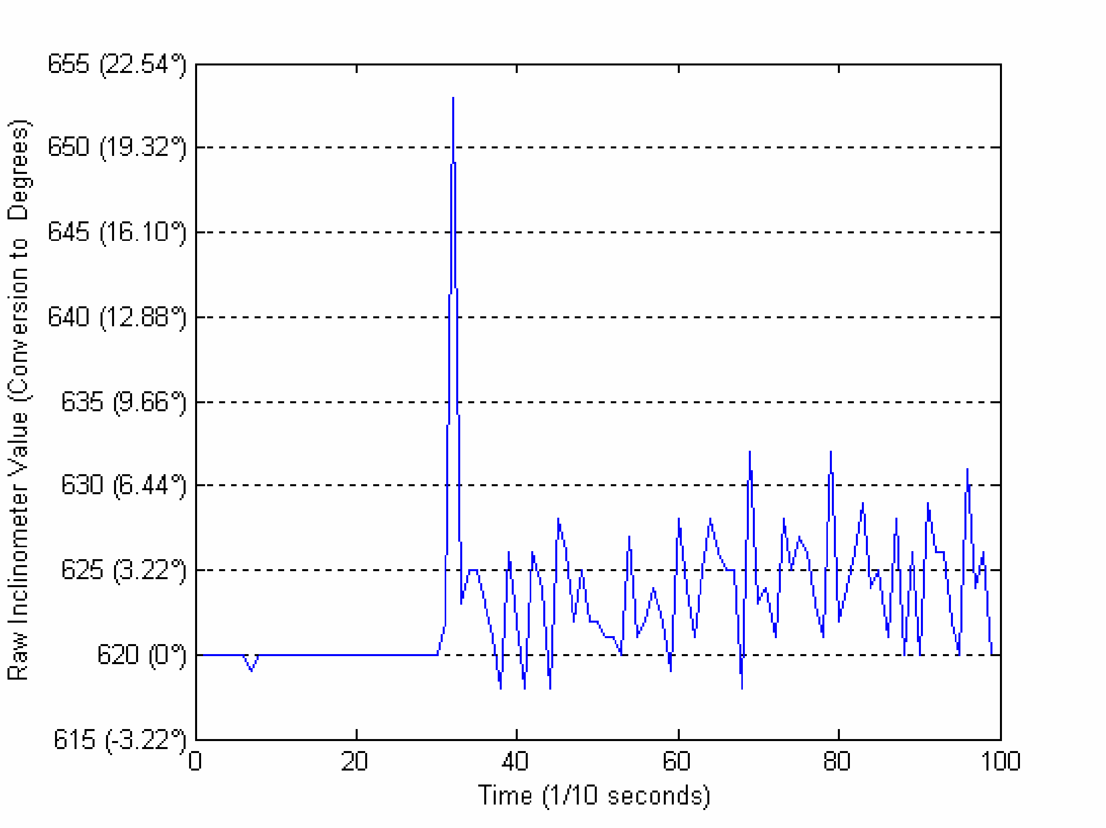

Figure 5.2.2: Dynamic Horizontal Test 2

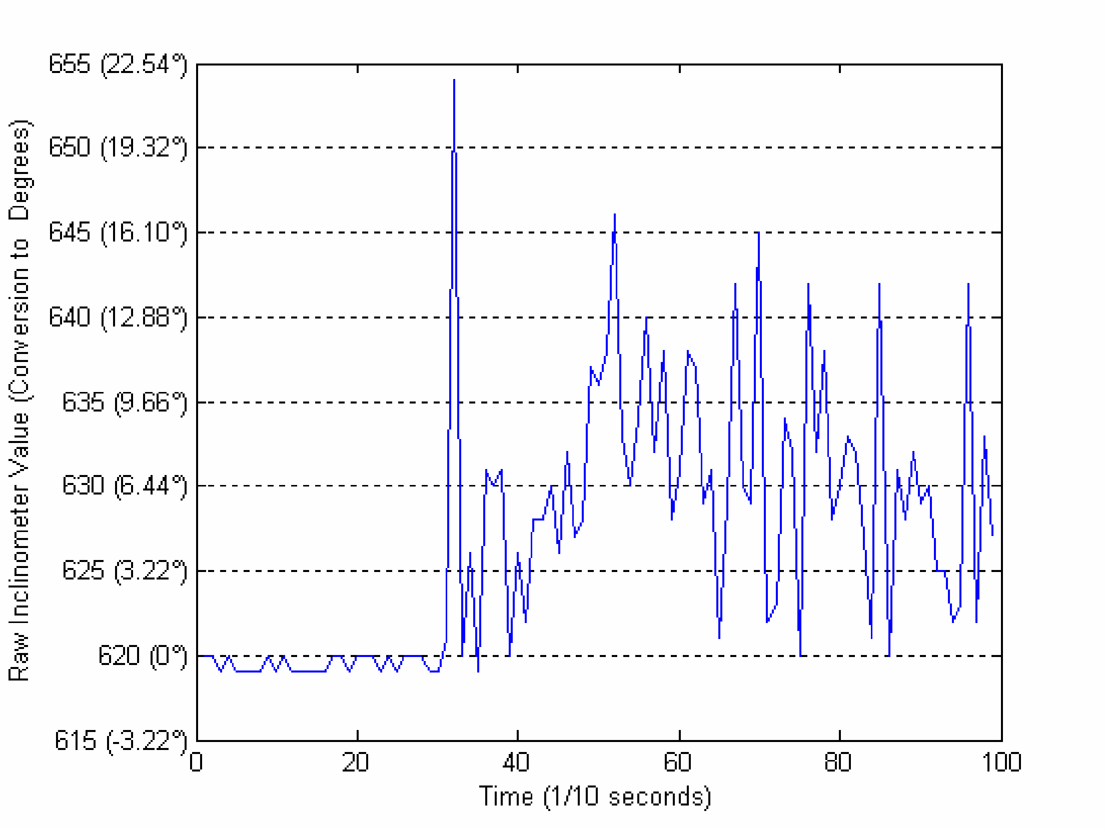

Figure 5.2.3: Dynamic Horizontal Test 3

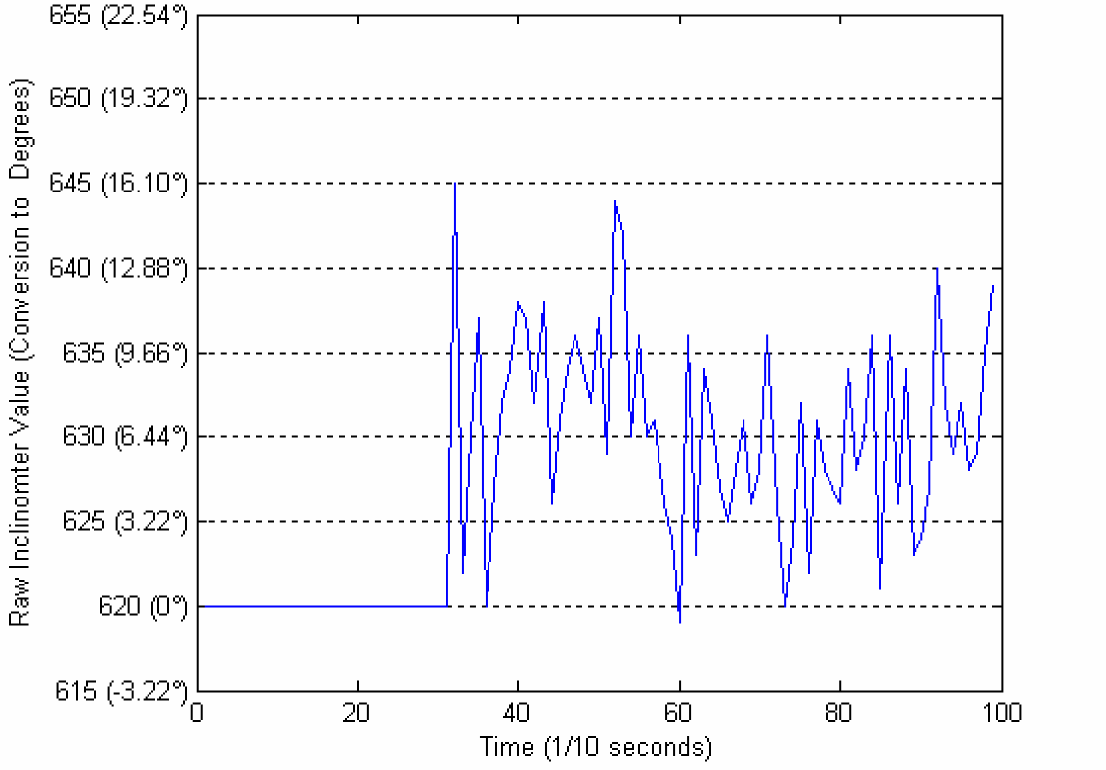

Figure 5.2.4: Dynamic Horizontal Test 4

For the first 3 seconds in all cases, the value of the inclinometer was approximately 620 without much variation; this corresponds to results of the static tests. At the time of 3 seconds when the robot begins to accelerate, the raw value read from the inclinometer suddenly spikes a significant amount. In some cases, this corresponds to an angle of about 21° which is a substantial error. This value quickly drops back to a more desired reading but continues to fluctuate significantly as the robot drives at cruising speed. From the figures shown above, it can easily be seen that the amount of oscillation can vary greatly even between tests performed one after another. The results of the four tests performed are shown in Table 5.2.1.

Test Number	Average after spike	Standard Deviation after Spike	Size of Spike	Conversion of spike to degrees
1	628	7.20	29	18°
2	624	3.32	33	21°
3	630	6.54	34	21°
4	630	5.60	25	16°

Table 5.2.1: Dynamic Test Results

These results show how much the motion of the robot actually affects the readings from the inclinometer. The amount of oscillation in the readings has a substantial bearing on the accuracy of the altitude calculations which is described by the standard deviations in the table. The standard deviations are calculated for the values recorded after the first spike. The deviations found are considerably greater than those found in the stationary tests. These oscillations are very difficult to eradicate, but perhaps using the average over the interval of the oscillations would increase the accuracy of the altitude computations.

There is a simple explanation for these errors when the robot is in motion. The initial spike in the raw reading is caused by the sudden application of power to the motor causing the robot to lift at the front end. This happens with many robots and other vehicles and is impossible to completely remove. Reduction of this error could be done mechanically by stiffening the robot with respect to its tracks but this is probably not the best option. The reason this would decrease the size of the spike is because it would reduce the amount of travel between the tracks and the robot and hence reduce the size of the spike. The size and frequency of the oscillations also depends on the speed and the surface which the robot is traversing. Another option is to reduce the rate at which the robot accelerates.

When the robot is instructed to drive forward, application of power to the motor is effectively instantaneous. By progressively increasing the amount of power to the motor, the effect of the sudden burst of power is greatly reduced, therefore reducing the size of the initial spike in the reading of the raw inclinometer values. This would make the robot’s terrain mapping more accurate but at this point in time, this spike is almost irrelevant.

The results from the dynamic tests were very inconsistent. This means that the zero value for the angle equation could not be adjusted based on these results. Firstly, the effect on the segments of an erroneous zero value must be investigated. Following this, a simple test of driving the robot over a mound, with known dimensions, must be performed. The results from this test performed with various zero values would then determine the offset required for the calculation of the angle.

5.3 Effect on Segment Calculations

Since the robot calculates the altitude every 4cm, this initial spike does not affect the calculation because it has no bearing on the reading when the calculation is made. However, to get more resolution into the map by reducing the size of the segments, as previously discussed, this spike may have an effect on the reading at the end of the segment depending on how small the segments are. Although this spike may affect the first altitude for each lap, reducing the amount of error when the robot is at cruising speed is far more important.

The oscillation of the readings from the inclinometer when the robot is at cruising speed must be accounted for. This is because for the vast majority of the time, when the robot is computing the altitudes, the robot is at cruising speed, so the oscillations in the inclinometer's measurements produce a significant amount of the error in the terrain map. By taking the average or the inclinometer readings over the 7 seconds that the robot is in motion, excluding the initial spike, and calculating the height over one 4cm segment using this average value, the height for one segment in each of the four dynamic test cases is shown in Table 5.3.1.

Test Number	Height of 1 Segment, using average
1	0.35 cm
2	0.17 cm
3	0.44 cm
4	0.44 cm

Table 5.3.1: Error in Height Calculation Over Segment

This table shows the amount of error that exists in each segment without any adjustments being made in the way that the heights are calculated. With the robot driving over a horizontal surface in each of these tests, the height at the end of each segment should be zero. Since each height calculation is based on the previous segment's height, the error obtained in each segment will accumulate throughout the terrain map. This means that action has to be taken to reduce this error as much as possible.

5.4 Mapping Performance over Known Mound

Testing of the robot's terrain mapping abilities over a mound with known dimensions was used to determine a zero value for the angle equation. Performing these tests with various zero values and determining the best option, would prove to produce more accuracy in the terrain map.

To reduce the effect of the oscillations in the readings from the inclinometer, when the robot is in motion, has proven to be quite difficult. There are two options to explore to reduce oscillations, the first of which is to change the raw value that corresponds to an angle of zero. Initially, the raw value corresponding to 0° of inclination for calibration, was the average value for the inclinometer when the robot was stationary. This is intuitively how the inclinometer would be calibrated but produces misleading measurements. The calibration should be made when the vehicle is driving forward at cruising speed because this is the state that the robot is in when the altitude calculations are made. Using the average of the inclinometer reading, over the 7 seconds that the robot is in motion, as the 0° calibration point would make more sense. Deciding on which average from the tests should be used can only be done by performing tests on each of them. To produce the best results, the robot's terrain mapping accuracy was tested by driving the robot over a mound, so that the initial height and the end height were the same.

Testing and adjusting the robot's terrain mapping parameters until the best combination for accuracy is found would greatly increase the accuracy of the terrain map. This test was accomplished by driving the robot over a mound with known dimensions and analysing the results. By changing the value for the zero level of the inclinometer, an optimum point could be found so that the height at the beginning and the end were equivalent. Also, the height along the horizontally flat surfaces would become relatively constant with minimal variation. Looking at the height that the robot calculates on top of the mound and comparing it to the physical value would also be an important result. Finally, determining how well aligned were the transitions from sloped to horizontal surfaces was another factor in determining the zero value. Analysing all of these with a range of tests, were performed and an optimum level was found.

Optimising the zero value was the first concern in these tests. Figures 5.4.1 through to 5.4.3 below show the results of these tests. Each figure describes the height of the robot as it drives from left to right over the mound. A video of this test is included on the enclosed CD. In Figure 5.4.1, the zero value is 622 and the figure shows the height of the robot increasing even when it is driving on a horizontal surface. This is not desirable as it affects the height of the robot after it has negotiated the mound. This figure also shows the altitude on top of the platform to be rather erroneous. The error in the calculated height along the platform is 5.15±0.22cm. This is a large amount of error considering that the height of the platform is 19.70cm. The difference in the altitude at the beginning and end was found to be

2.25cm. This would lead to large amounts of accumulated error if this value were used for the terrain mapping. The alignment of the transitions is inadequate as well, leading to this value being eliminated.

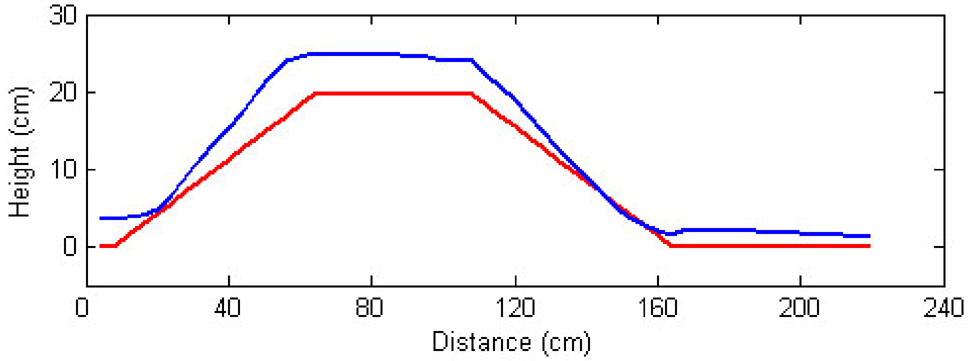

Figure 5.4.1: Mound Test with 622 as 0°

Figure 5.4.2 is the result of the test using 624 as the zero value. This figure shows the calculations maintaining consistency quite well for the horizontal surface but the beginning height and end height do not align well enough. In fact they differ by 2.8cm. The alignment of the transitions has improved over the case using 622. The error in the height on top of the platform in this case was found to be 1.6±0.27cm which is far better than the previous case. However, the most important factor in the decision of the zero level is the beginning and end height.

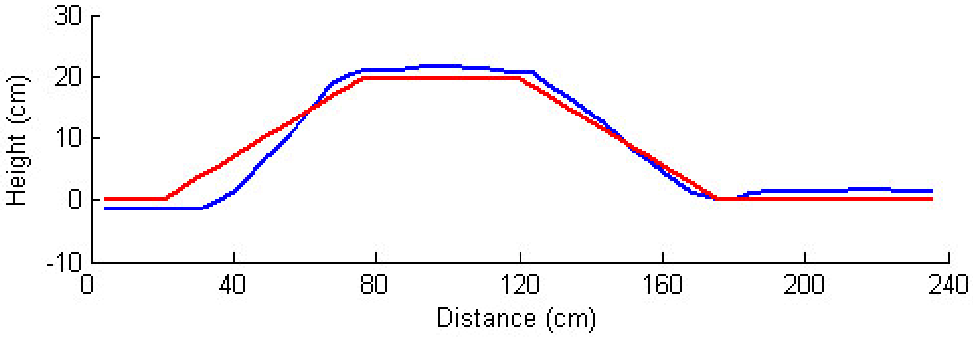

Figure 5.4.2: Mound Test with 624 as 0°

Figure 5.4.3 shows the beginning and end heights being much closer and the flatness of the horizontal sections being acceptable. In fact the difference in the height of the beginning and end altitudes was found to be only 0.54cm which is far better than the previous two tests. The alignment of the transitions is comparable to those found in the 624 case. The error in the height calculated along the top of the platform was found to be 2.01±0.26cm which is close to that found in the 624 case. The flatness in the horizontal sections of the mound test are the key factors. This flatness and the consistency between the beginning and final altitudes led to 623 being chosen as the zero value. This leads to the least accumulation of error for the entire terrain map. Although there is a 10% error in the height at the top of the slope, this will not increase the accumulation of error, as the height before and after the mound is almost equal.

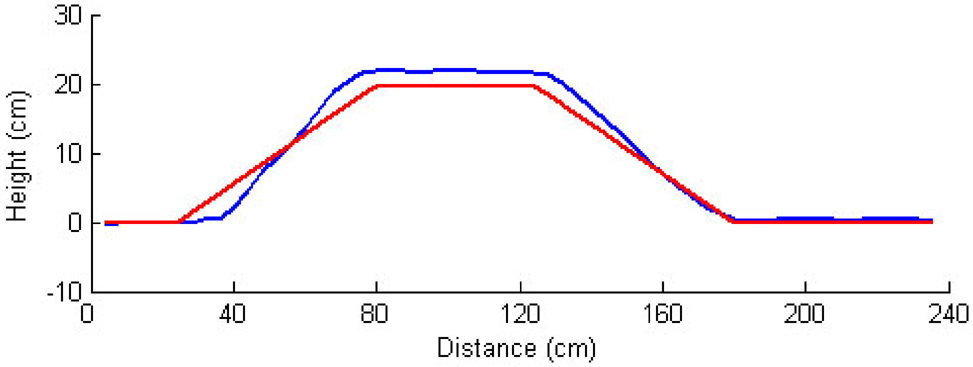

Figure 5.4.3: Mound Test with 623 as 0°

Since the height calculated at the beginning and end of the test, with 623 as the zero value, is consistent, the altitude on top of the mound must be due to the gradient in the equation for calculating the inclination angle. The reason for this error would mainly be due to error in the measurement of the slope of the mound. However, the error does not influence the overall result of the terrain map and is within 10% so at this stage this error has been noted but not resolved.

To further improve the accuracy of the terrain mapping, the resolution of the map was increased. This means that more computations are made over the same distance, therefore increasing the level of detail in the representation of the terrain, especially in transitions from horizontal surfaces to sloped surfaces and vice versa. This also increases the precision of the mound tests and ultimately the final terrain map. Firstly, the segment size was reduced from 4cm to 2cm. This produced a representation shown in Figure 5.4.4.

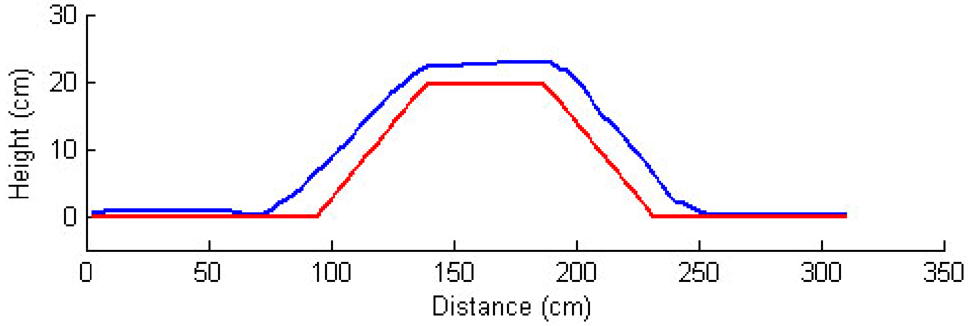

Figure 5.4.4: Mound Test with 2cm Segments

Noticing the difference between the two mound tests with 2cm and 4cm segment sizes is a little challenging. The main areas where the difference is most apparent are between transitions of slopes and the effect of large oscillations of the inclinometer. It can be seen that the transitions are smoother due to the number of calculations being made as the robot negotiates the transition. The smoothness of the transitions, as well as other points along the curve, is also due to the fact that having more frequent computations, the effect of a large oscillation in the inclinometer's reading is now reduced. This is because the oscillation will affect a segment that is now half the size, as the oscillation has no bearing on the next segment. The performance increase gained from the reduction of the segment size was significant, so another test was performed with 1cm segments. The results are shown in Figure 5.4.5 below.

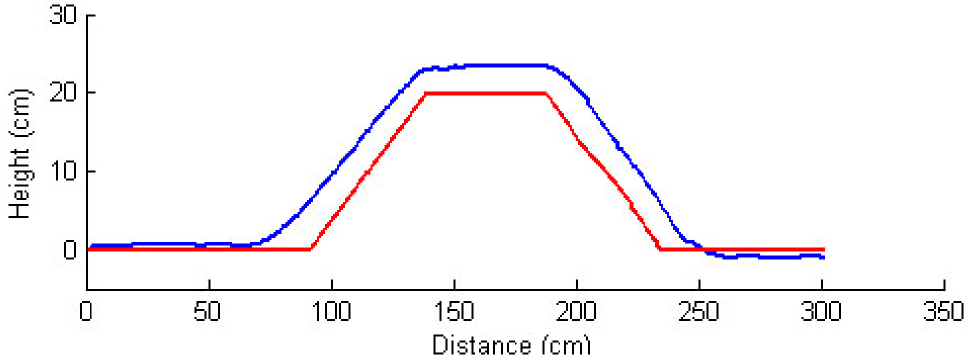

Figure 5.4.5: Mound Test with 1cm Segments

As with the previous reduction of the segment size, the performance gained from halving the segment size was worthwhile. Again, the area in which this gain is more noticeable is in the transitions and inclinometer oscillations. Due to the increased resolution of the mapping function, both the 2cm and 1cm cases have also recorded, and incorporated into the graph, the distance that the robot has traversed for each calculation that has been made. This increases the quality and precision of the curves shown in these figures. With the initial cases that used 4cm segments, the effect of the varying segment was negligible for the purposes of the graph but was found to be important for smaller segments. These segment sizes will also be incorporated into the final terrain map.

These tests produced a value for the zero level for use in the Equation 5.1.2 for calculating the angle of inclination. The value that was used was 623 and the performance of the terrain mapping for the test area was performed using this value.

5.5 Terrain Mapping Performance

A section of the representation of the physical area produced by the robot is shown in Figure 5.5.1. This figure also shows the path that the robot took to perform the mapping, represented by a black line. Underneath this is Figure 5.5.2 showing the ideal terrain map. This terrain map is the initial terrain map produced by the tracked robot which does not include the adjustments made to the calibration of the inclinometer and also does not include the increased resolution along with the actual segment size being recorded, as discussed in the previous paragraph. A video of the robot performing terrain mapping is included on the CD attached to this thesis.

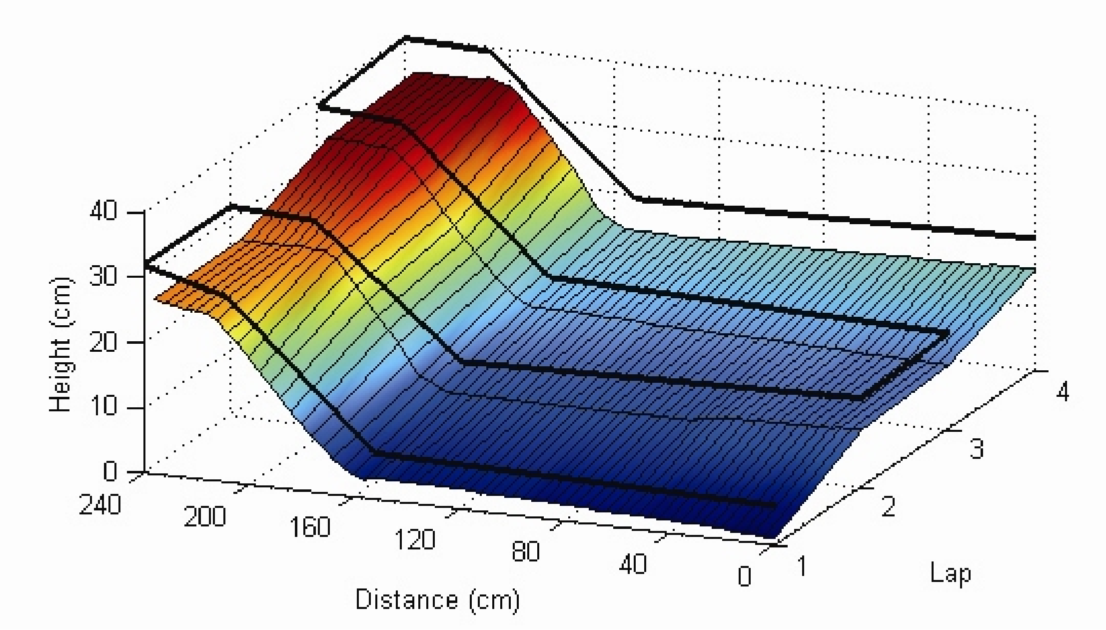

Figure 5.5.1: Initial Terrain Map

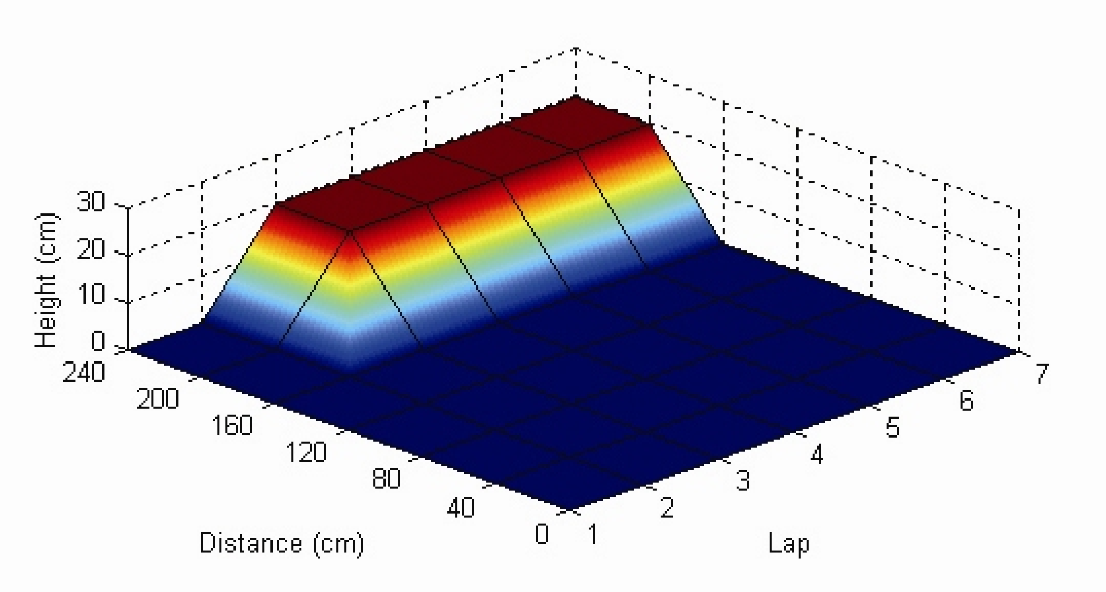

Figure 5.5.2: Ideal Terrain Map

In this terrain map, there are a few things to note with respect to determining where errors have occurred. Firstly, it can be seen that the edge of the platform is not a clean edge between each lap. This is highlighted by the brown areas on top of the platform not lining up with each other. This is mainly due to the resolution of the segments, so decreasing the size of the segments would make this edge cleaner. This also occurs at the bottom of the slope and can be resolved in the same way.

Another problem is along the lap itself. Along the horizontally flat part of each lap, the altitude is slightly increasing in the direction that the robot traversed the lap. The effect of this error is very apparent as the lap is followed in the representation. This not only affects the altitude of the segments along the horizontal surface but also affects the height at the top of the slope. Subsequently, each lap succeeding the current lap is influenced by the errors obtained in the current lap. Therefore, all errors in altitude calculations accumulate through the remainder of the map. By implementing the methods or reducing the amount of error previously discussed, the final terrain map should be much more concise and accurate.

Implementation of the methods for reducing the error in the terrain map representation does in fact improve the quality of the map. To improve the accuracy of the map, the size of the segments was reduced to 10mm, the calibrated zero value was changed from 620 to 623 and the segment sizes were recorded at the time of the computation of the altitude. The results of these modifications produced a more accurate terrain map illustrated in Figure 5.5.3 below. This figure also includes the path used to map the terrain, shown with the black line.

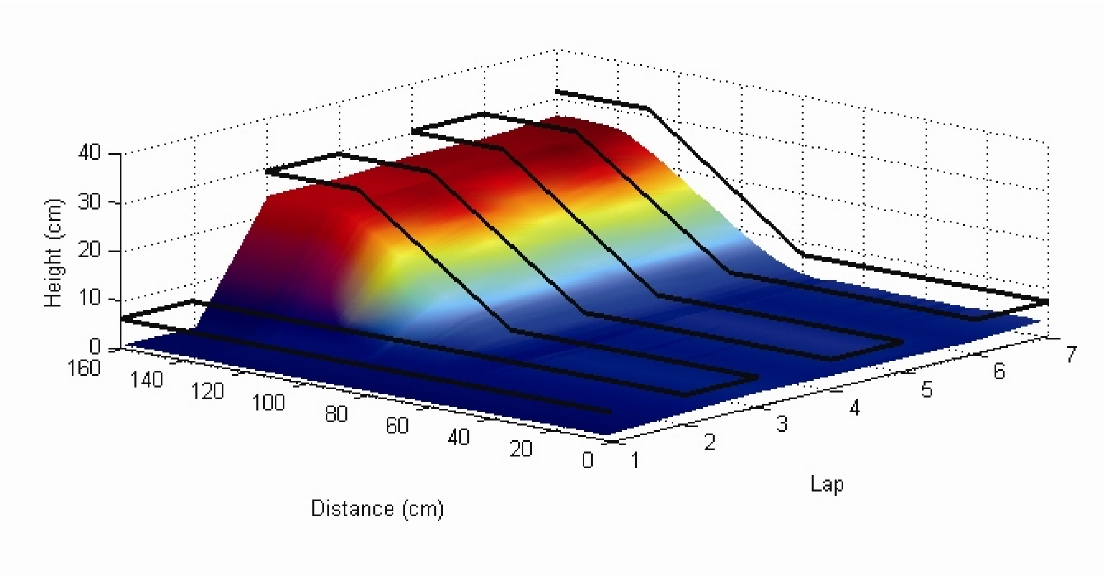

Figure 5.5.3: Final Terrain Map

This terrain representation has vastly increased accuracy and precision compared to the initial terrain map produced by the robot's mapping of the test area. Figure 5.5.3 above, demonstrates the reduction of error accumulated along the horizontal sections of each lap. These sections are considerably smoother, meaning that error accumulation throughout the remainder of the map is reduced. For the first lap, there are no slopes and Figure 5.5.4 depicts the smoothness of this lap.

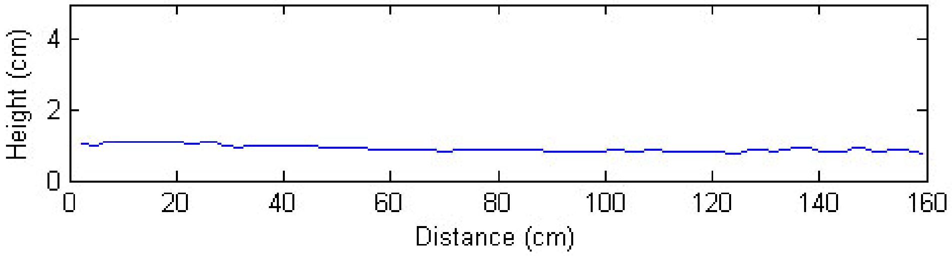

Figure 5.5.4: First Lap

The overall effect of this increase in flatness is more noticeable when observing the difference in altitudes calculated between adjacent laps. In the image of the terrain map, determining the difference in heights between adjacent laps is very difficult. In fact, the difference in height between the first lap and the final lap is only 1.8cm. This is extremely good when realising that the distance between the first lap and the final lap is approximately 2m. The robot has actually travelled over 15m from the start point to the last lap, meaning that the error in the altitude is about 2cm over 15m. This is an excellent improvement on the first terrain map produced in Figure 5.5.1. The original terrain map had a difference of 14cm. This improvement also produces a flatter platform at the top of the slope. The difference between the edges of the platform, parallel with the laps, is 2.1cm. This demonstrates how improved calibration techniques for the inclinometer and reduced segment sizes increase the accuracy of the terrain map; however these are not the only improvements.

There is a noticeable enhancement of the sharpness of the transitions between horizontal surfaces and sloped surfaces. Focusing on the base of the slope, it can be observed that the alignment between laps is clear and defined, improving the representation of the slope. This is also apparent at the top of the slope, however, there are more considerations to acknowledge.

With the increase in the resolution of the terrain map representation, previously unidentified sources of error become apparent. The most obvious of these errors can be observed when the robot begins to descend the slope. A bump exists on the point of transition from the horizontal platform to the slope. This bump is quite significant but is known and can be accounted for by means of post-processing. The contour of this bump is more apparent when viewing the terrain map from a different angle, as shown in Figure 5.5.5.

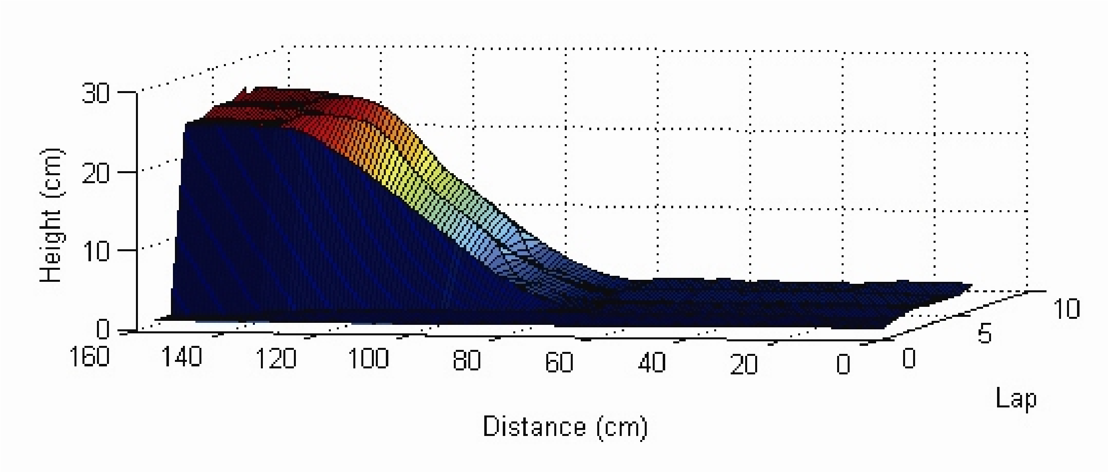

Figure 5.5.5: Emphasised Bump

This bump exists because, as the robot approaches the slope, the front of the robot will overhang the slope. When the centre of gravity of the robot travels beyond the transition point of the platform to the slope, it will suddenly fall. Since the inclinometer uses liquid to perform its measurement, there is a delay in the liquid becoming horizontal causing an error in its reading. The reason that the error is an increase in height is because the actual pitch angle of the inclinometer is greater in declination than the level of the measurement device within the inclinometer, due to the delay. This produces a positive result from the inclinometer and therefore produces a positive height. Reducing the effects of this source of error is difficult to perform whilst the robot is in motion. The robot could possibly detect a slope, using the downward angled front PSD or even a gyroscope, and then reduce the speed as the robot manoeuvres over the slope. This would reduce the amount of momentum produced in the levelling device within the inclinometer but at this point the amount of error is reasonable and post-processing would be more appropriate. Comparing the performance of a wheeled robot with the tracked robot at this transition would be interesting as the wheeled robot, with four wheels, does not drop off the edge of the platform onto the slope. This drop also causes errors as the robot descends the slope.

The bump is not the only source of error on the descent down the slope. As the robot traverses the edge, the robot will bounce slightly when it makes contact with the slope. This will affect the flatness of the slope which can be seen in the final terrain map displayed previously (Figure 5.5.5). Part of this error is due to the robot reducing its speed once it is on the slope. This will cause another small bump in the slope but not as extensive as the one obtained at the top of the slope. This phenomenon also exists at the bottom of the slope and when the robot climbs up the slope but, due to the decreased speeds at these points and smoother transitions, the effects are not as noticeable but are worth investigating further.

5.6 Error Accumulation and Post-Processing

As discussed in Section 4.3 for recording data, the robot requires the previous altitude calculation in order to compute the height for the current segment. This means that all the calculations are based on the initial height and accumulate more and more error as the robot traverses the terrain. Figure 5.6.1 shows the amount of error at the beginning of each lap, relative to the starting point. This demonstrates that as the robot travels farther, the amount of error increases. In the test arena, the beginning of each lap is known to be zero as it is horizontally aligned with the starting point. This means that the altitude calculated at these points by the robot should be zero, so the amount of error can be determined.

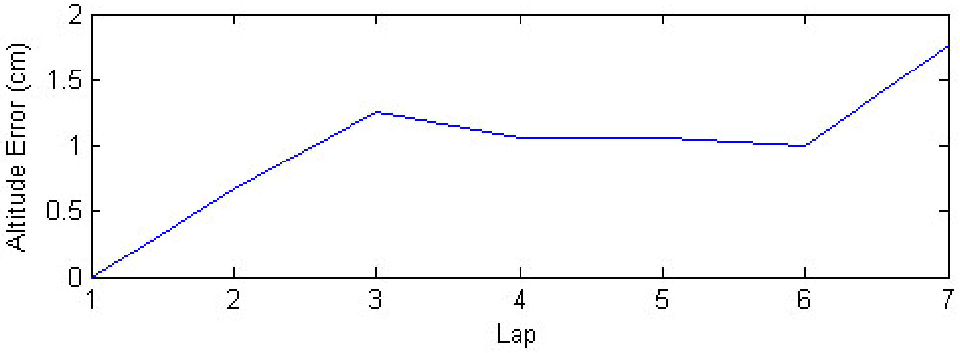

Figure 5.6.1: Error at the Beginning of Each Lap

Removing this accretion of error is rather difficult given the limited sensors available on the robot. However, elimination of this error may not necessarily be imperative in the context of the robot's function. Although it may not be a problem of high importance, it should still be acknowledged and reduced where possible.

There are a couple of ways to combat this problem of error accumulation, without adding more sensors. One of them is to run the robot's laps in one direction and then to have it remap the terrain with a 90° shift in the laps. For example, have the laps run from east to west and then rerun the mapping with laps that run from north to south. This would obviously take at least twice the amount of time and may not be desirable in certain situations. This would also require extra processing either during, but more likely after, the creation of the maps so that one map is imprinted on top of the other, hopefully giving a more accurate representation.

Post-processing could also be used on the single representation generated by the robot. The correlations between laps, especially in the test area, are extremely high meaning that on one lap, if the robot climbs a slope then on the next lap the robot would have to descend an equivalent slope. This means that the beginning and ending of the slope can be averaged over the laps, giving sharper edges. The altitudes on each lap could be adjusted according to the previous lap or both adjacent laps giving greater consistency over the features of area. This method would reduce the detail of the mapping but would provide consistency and sharp detail of the features of the area. Again, the type of post processing or error filtering used greatly depends on the requirements of the robot depending mainly on where the robot is to be used.

As shown, reduction of the error accumulation is not simple and requires more than just a method of how to reduce the error from accruing. It also requires a decision to be made on which of the available methods is implemented, depending on the environment and the purpose of the robot. In the test arena, a very simple filtering method of the inclinometer reading can be implemented, greatly increasing the accuracy of the map. This would simply be to assume that the inclination is zero when the average reading over the segment is between 3 and -3 degrees. This would mean that each time the robot drives on a horizontal surface, the error accumulation would be reset since the robot would record altitudes of zero in these segments.

Filtering the inclinometer's average readings, and assuming it to be zero when it is within a range, does work in the test area but defeats the overall objective of the robot. This can only be done in a known area that meets a certain criteria. The area must be fairly simple, with the angles of the slopes exceeding the range of the horizontal surface filter. The objective of this robot is to perform terrain mapping, not "test area" mapping so this filtering method is impractical. The robot must be able to determine if it is climbing very long slopes that would exist in reality and also realise that the correlation between laps is much less than that in the test arena.

With all these ideas in hand, it was found that post processing and running the mapping twice with one perpendicular to the other was not required. The reason for this is that the representation created on a computer, after the robot has performed the mapping, shows all significant features on the terrain with an acceptable level of detail. This means that if someone used this robot in an unknown area, they would be able to obtain a good understanding of the area without having seen it. Another reason why this decision is made is because the robot has to be used in the way that it was designed and used within the limitations of the sensors. This means that if another robot, or even this robot, were to use this mapping information to navigate through the terrain, there is ample information to be able to make a judgement on the path to take. If the robot were following laps, it would be able to see that there is a slope approaching, with good accuracy, and navigate accordingly. With regards to driving across laps, the amount of error between laps is insignificant because each lap is spaced by the width of the robot, approximately 25cm. Even if the error between the altitudes from the two adjacent laps were extremely erroneous, say 5cm, this translates into an incline of only 11° which does not affect the robot's ability to traverse the terrain as all mobile robots should be able to climb at incline of this degree.

5.7 Performance Comparison

Comparing the tracked robot's method of terrain mapping with other mobile robots introduces some interesting points. As seen in the Terrain Mapping Performance section, the level of detail of the terrain map obtained by this robot is extremely good, with an acceptable level of accuracy. The level of detail that other robots obtain can now be compared with that of the tracked robot. The robots discussed in a paper titled "Terrain Aided INS Robot Navigation", use two methods of obtaining the altitude [13]. The comparison produced some interesting outcomes.

The degree of accuracy of the altimeter would greatly affect the terrain mapping precision level. The altimeters used on these robots have an absolute accuracy range of ±4m over its operating range [13]. For large, wide open areas this level of accuracy would be acceptable but if it were to be used on a similar scale to that of the tracked robot, the terrain mapping abilities would be deficient. These types of altimeters are commonly used on aeroplanes where a difference of 50 metres has no bearing on its ability to fly. These sensors require calibration and use barometric pressure to measure the altitude. They are also subject to many other factors that influence the accuracy. This shows that on a small scale, the tracked robot's method of terrain mapping is excellent but is only applicable to small areas.

The other method these robots used to obtain the altitude was achieved using Global Positioning System (GPS). For these all terrain mobile robots, they used differential corrections in conjunction with the GPS. This greatly increases the accuracy of the solution but is still inadequate for making precise altitude measurements. The accuracy obtained by the GPS was ±1m which is of the same order of magnitude as the altimeters [13]. Similarly, this does not provide the accuracy required in order to perform terrain mapping in small areas.

These points demonstrate how the tracked robot obtains an accurate terrain map in small areas compared to other mobile robots. The tracked robot at this point is limited to small areas but could be tested over large areas compared to that which these robust, outdoor robots achieve in terms of terrain mapping. The major point to understand from this comparison is that the robots should only be examined in the field for which they are designed. That is, these large outdoor robots should not be used in a nine square metre area and similarly, the tracked robot should not be instructed to map a square kilometre. However, it is interesting to see how the differing approaches compare with one another.

5.8 Summary

Overall, the terrain mapping performance of the tracked robot was highly accurate and precise. The method of obtaining a terrain map proved to be exceptional with a minimal amount of error. Performing the various tests to reduce the amount of error accumulation was essential in ensuring the precision of the representation of the test area.

A key advantage of the tracked robot is its cost effective method of obtaining a detailed terrain map. The robot requires a minimal amount of sensors which are relatively cheap. Altimeters and GPS devices can get rather expensive and do not provide any benefit in performing terrain mapping in this case; in fact they would not be able to produce a terrain map. Differential corrections for GPS solutions also require subscriptions from providers, in most countries, which can cost in the order of thousands of dollars a year. Due to the simplicity of the tracked robot’s method of obtaining altitudes, there are only two sensors that are essential in adding similar functionality to other robots. To implement terrain mapping on an existing robot would only require an inclinometer and an optical encoder, or any other means of obtaining inclination and distance, assuming the robot already has the ability to traverse an area whilst avoiding objects. Then simply programming the robot to perform height calculations and record position information would add the functionality to existing robots.

The accuracy and cost effectiveness of the tracked robot makes for a practical solution for the production of terrain maps. The method is easily transferable to existing robots and its simplicity holds great value as it reduces the chances of unexpected behaviour of the robot.

6. Conclusion

6.1 Final Remarks

The method of driving a tracked robot in a lawnmower pattern to map the terrain of a test area proved successful. The calibration of the inclinometer, especially with the mound tests, was instrumental in reducing the amount of error accumulation throughout the mapping process. This method of terrain mapping also proved to be cost effective and easily transferable to existing robots.

The lawnmower pattern proved sufficient in traversing the entire test area. This allowed calculations of the altitude to be made for each segment within the area, but unfortunately introduced error accumulation. This occurred because each calculation was based on the previous altitude which meant that errors in all of the prior computations were evident in succeeding calculations. This problem could not be avoided because each height calculation had to be based on the previous height in order to produce a terrain map. The process for completing the turns at the end of each lap also had a strong bearing on the accuracy of the map. The turns had to leave the robot perpendicular to the wall to ensure that extra distance was not covered by driving at angles. The turns were completed successfully allowing for accurate distance calculations to be made.

Testing of the inclinometer's characteristics in various situations was pivotal in reducing the amount of error in the altitude calculations. The static tests provided a strong method for obtaining the linear characteristics of the inclinometer, enabling an accurate conversion from the raw inclinometer value to an angle in degrees. The dynamic tests illustrated the amount of variance in the inclinometer's reading. It demonstrated how the inclinometer's measurement spiked as the robot began acceleration, and showed how the reading varied as the robot traversed a flat terrain. This led to the robot being tested by driving over a mound and methodically adjusting the zero-value parameter. This was done in order to obtain the most accurate altitudes, shape and alignment of the resulting graph. Using the parameters obtained in these tests greatly reduced the amount of error accumulated over the entire terrain, to a point where the error became insignificant in respect to the size of the test area.

The resulting terrain map produced with the data obtained from the robot's mapping was exceptionally precise and informative. It reproduced all of the key features of the test area with a high degree of accuracy. This information was represented in a way that could easily be ported to other mobile robots in order for them to navigate terrain with ease.

Due to the simplicity of the terrain mapping technique, this functionality can be added with ease to existing robots. The process is very cost effective when compared to other means of obtaining altitude measurements, such as altimeters and GPS receivers, making it desirable and practical.

Overall, the tracked robot project was successful in proving a technique for terrain mapping which can now be used on other robots with minimal difficulty. The project discovered interesting characteristics of robot locomotion and inclinometer performance.

6.2 Future Work

There are still improvements that can be made on the tracked robot to allow increasingly precise altitude measurements and manoeuvrability.

The altitude calculated at the top of the platform on the mound tests could be more accurate. At the moment, an error of 10% exists which should be reduced. This would require more accurate measurements of the dimensions of the slopes, allowing the angle of the slope used in the calibration of the inclinometer to be more accurate.

Obstacle avoidance should be implemented which would allow the robot to map the entire terrain by navigating around objects. This would require the lawnmower pattern methodology to be reassessed. The robot would then be able to calculate the altitude at each reachable segment of the area. This would make the tracked robot more practical in the real world as objects will often block a robot's path.

To increase the level of detail of the terrain map, the path taken by the robot to perform a turn could be modified. The modified path would allow the robot to manoeuvre into a position to begin the next lap, remaining a desired distance from the last lap. That is, reduce the distance between each lap allowing more detail of the terrain to be mapped.

Using another inclinometer to measure the roll of the robot would also increase the amount of information obtained by the mapping. This would mean that the robot could drive on left or right sloping surfaces and obtain the slope for that surface, consequently producing more detailed information about the terrain..

Implementation of a method to slowly accelerate would also be desirable. This would reduce the size of the spike in the inclinometer's reading and, in-turn, reduce the amount of error. An optimal cruising speed could also be found that would diminish the intensity of the oscillations. Using either the gyroscopes, that are currently unused, or the anterior down-angled PSD, used to detect a slope, would allow the robot to slow down before declining slopes. This would reduce the significance of the bump that was evident in the tests.

Finally, by investigating how the robot manoeuvres over certain features of the terrain, post-processing can be performed to remove abnormalities in the terrain map. Some of these abnormalities are already evident in the results from the various tests, but implementation of a post-processing program has not been undertaken. This would increase the sharpness of the terrain map and greatly increase its credibility.

7. References

[1] A. Ellery, Environment-robot interaction--the basis for mobility in planetary micro-rovers, Robotics and Autonomous Systems, Volume 51, Issue 1, 30 April 2005, Pages 29-39. URL: http://www.sciencedirect.com/science/article/ B6V16-4F8TKDD-1/2/77aaa9238f98fae1b2f97eada3e0cfd9

[2] Z. Fan; Y. Koren; D. Wehe, Tracked mobile robot control: Hybrid approach, Control Engineering Practice, Volume 3, Issue 3, March 1995, Pages 329-336. URL: http://www.sciencedirect.com/science/article/ B6V2H-3YK00PN-54/2/fe0f67f7bd62eccde96fec42b4457eeb

[3] L. Matthies; Y. Xiong; R. Hogg; D. Zhu, A. Rankin, B. Kennedy, M. Hebert, R. Maclachlan, C. Won, T. Frost et al., A portable, autonomous, urban reconnaissance robot, Robotics and Autonomous Systems, Volume 40, Issues 2-3, 31 August 2002, Pages 163-172. URL: http://www.sciencedirect.com/science/article/ B6V16-46H7WDX-6/2/0ad3016e6ad8a6888e9e8199f52704dd

[4] M. Weigl; B. Siemiaatkowska; K. Sikorski; A. Borkowski, Grid-based mapping for autonomous mobile robot, Robotics and Autonomous Systems, Volume 11, Issue 1, May 1993, Pages 13-21. URL: http://www.sciencedirect.com/science/article/ B6V16-4808W88-1G/2/ad43604fd67a7a54e71b3d0266767c64

[5] T. Okui; Y. Shinoda, An outdoor robots system for autonomous mobile all-purpose platform, Robotics and Autonomous Systems, Volume 17, Issues 1-2, April 1996, Pages 99-106. URL: http://www.sciencedirect.com/science/article/ B6V16-3VS5JCB-F/2/b3a519553823b8e11fc8174bcb9d5bc4

[6] H. Schempf; E. Mutschler; B. Chemel; S. Nayar; C. Piepgras; W. Crowley; S. Boehmke, Evolution of urban robotic system developments, Advanced Intelligent Mechatronics, 1999. Proceedings. 1999 IEEE/ASME International Conference on,

Vol., Iss., 1999 Pages:689-694 URL: http://ieeexplore.ieee.org/iel5/6508/17391/00803251.pdf?isnumber=17391∏=STD&arnumber=803251&arnumber=803251&arSt=689&ared=694&arAuthor=Schempf%2C+H.%3B+Mutschler%2C+E.%3B+Chemel%2C+B.%3B+Nayar%2C+S.%3B+Piepgras%2C+C.%3B+Crowley%2C+W.%3B+Boehmke%2C+S

[7] M. Iagolnitzer; F. Richard; J. Samson; P. Tournassoud, Locomotion of an all-terrain mobile robot Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on, Vol., Iss., 12-14 May 1992Pages:104-109 vol.1 URL: http://ieeexplore.ieee.org/iel2/399/5764/00220327.pdf?isnumber=5764∏=STD&arnumber=220327&arnumber=220327&arSt=104&ared=109+vol.1&arAuthor=Iagolnitzer%2C+M.%3B+Richard%2C+F.%3B+Samson%2C+J.F.%3B+Tournassoud%2C+P

[8] Stability of a multi tracked robot travelling over steep slopes, Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on, Vol.5, Iss., 26 April-1 May 2004 Pages: 4701- 4706 Vol.5 URL: http://ieeexplore.ieee.org/iel5/9126/28923/01302459.pdf?isnumber=28923∏=STD&arnumber=1302459&arnumber=1302459&arSt=+4701&ared=+4706+Vol.5&arAuthor=Shoval%2C+S

[9] C. Chen; Reactive locomotion control of articulated-tracked mobile robots for obstacle negotiation, M.M. Intelligent Robots and Systems '93, IROS '93. Proceedings of the 1993 IEEE/RSJ International Conference on, Vol.2, Iss., 26-30 Jul 1993 Pages:1349-1356 vol.2 URL: http://ieeexplore.ieee.org/iel2/4461/12646/00583774.pdf?isnumber=12646∏=STD&arnumber=583774&arnumber=583774&arSt=1349&ared=1356+vol.2&arAuthor=Chen%2C+C.X.%3B+Trivedi%2C+M.M

[10] T. Matsuda; K. Takagi; T. Sakamoto, Rough terrain robot with 4 turn-over track frames-a feasibility study on the mechanism, Systems, Man, and Cybernetics, 1999. IEEE SMC '99 Conference Proceedings. 1999 IEEE International Conference on, Vol.4, Iss., 1999 Pages:842-847 vol.4
URL: http://ieeexplore.ieee.org/iel5/6569/17543/00812517.pdf?isnumber=17543∏=S

TD&arnumber=812517&arnumber=812517&arSt=842&ared=847+vol.4&arAuthor=Matsuda%2C+T.%3B+Takagi%2C+K.%3B+Sakamoto%2C+T

[11] Mars Exploration Rover, Jet Propulsion Laboratory, National Aeronautics and Space Administration, October 2004, URL: http://marsrovers.nasa.gov/newsroom/factsheets/pdfs/Mars03Rover041020.pdf

[12] W. Warner, Great Moments in Microprocessor History, IBM, 22 December 2004, URL: http://www-128.ibm.com/developerworks/library/pa-microhist.html?ca=dgr-lnxw01MicroHistory

[13] S. Bruder; K. Wedeward, Terrain Aided INS Robot Navigation: A Deferred Decision Making Approach, Department of Electrical Engineering, New Mexico Tech, URL: http://isrg.nmt.edu/~ms/publications/f/MWS99.pdf

8. Appendices

8.1 Appendix I: CD Contents

Documents Folder:

Contains this dissertation paper, seminar slides, interim report and introductory research.

Photos Folder:

Contains various photos of the robot and related components.

Source Code Folder:

Contains the program code and associated programs for the project.

Tech Specs Folder:

Contains technical specifications of various components.

Test Results Folder:

Contains data resulting from the inclinometer tests and terrain mapping tests.

Videos Folder:

Contains videos of the robot performing terrain mapping as well as the mound test.

8.2 Appendix II: Tracked Robot Photos

Figure 8.2.1: Robot - Top View

Figure 8.2.2: Robot - Front View

Figure 8.2.3: Robot - Side View