Exploring the Dynamic Measurement of Position[*]

Cliff Randell Henk L. Muller
Department of Computer Science, University of Bristol, U.K.
cliff@cs.bris.ac.uk

Abstract

This paper describes the development of several methods for the dynamic measurement of error distribution for positioning systems. When a user moves along a path at an unknown speed, it is no longer feasible to calculate the distance between the measured and real positions. We outline three applications using different positioning systems and apply our methods to the evaluation of the performance of these systems. Two indoor ultrasonic positioning systems are described and compared with the use of the GPS system outdoors. Alternative options for the placement of sensors/antenna on the human body are evaluated.

Keywords *- position measurement, location sensing, GPS, ultrasonics, wearable computer,*

1 Introduction and Background

The measurement of position is one of the fundamental requirements of context aware mobile and wearable computing. The establishment of the U.S. Department of Defence's Global Positioning System (GPS) has been a significant factor in the development of external location aware computer applications [1, 2]. Much research has also been carried out into providing indoor positioning systems using ultrasonic and RF technologies [3, 4]. The performance of these systems is usually measured with the sensing device at a fixed reference position with sensed deviations from this position being recorded as the error distribution. The errors can be statistically analysed to produce specifications for the positioning system. Commonly used measures are the Circular Error Probability (CEP) which defines the radius of a circle that represents a 50 percent probability of a position lying in that circle; standard deviation (σ); and the root mean square (RMS) value. Additionally a 95% figure is often used in surveying and mapping. Assuming a normal error distribution this represents a 2 Sigma error. A review of positioning systems by Hightower and Borriello [5] 'strongly encouraged the location research and development community to investigate how to best obtain and represent such error distributions'.

Measurements are easily obtained for devices which are static, and can be reliably used for performance comparison. The measurement techniques usually require the sensing device to be placed in an optimum position, for example with a clear view of the sky for GPS, or with minimal obstructions for indoor systems. Readings from this sensing device are compared with a reference position at which the device is supposedly placed to generate the error data. However these conditions do not reflect how the devices are actually used in practice - signals are often blocked by the human body, and filtering using a model of the user's behaviour is used by many off the shelf systems e.g. by using a Kalman filter. Factors such as these can significantly affect the performance of positioning systems.

The object of the research described in this paper is to better understand and to quantify positioning errors as actually experienced by users. Using this knowledge we are able to improve and develop our current position sensing systems. The initial research was carried out in the Mackintosh Room of Scotland's Centre for Architecture, Design and the City in Glasgow, where an ultrasonic positioning system [6] has been installed. The constraints of this room provided a particularly challenging enviroment for position sensing. To support the research in the Mackintosh Room, two further studies were carried out. One in a near ideal location - in an atrium at the Hewlett-Packard Laboratories in Bristol U.K., and the other outdoors using GPS in the Millennium Square, also in Bristol.

At all three locations the positioning systems were used to support location aware applications using wearable computers. User tests were also carried out enabling subjective impressions to be gathered on the effectiveness of the location sensing at each site, but reporting on these falls outside the scope of this paper.

---

[*] Funding for this work is received from the U.K. Engineering and Physical Sciences Research Council, Grant No. 15986. Further support is provided by Hewlett Packard's Art and Science programme.

2 Test Procedures

We considered various methods to ensure that the sensor accurately completed a measured test route at a predetermined speed, or alternatively ways of making an accurate record of the path actually followed and its timing. The use of cameras to record the test was impractical due to the configuration of the Mackintosh Room; and robotic techniques would have required complex mechanisms to minimise inertial factors and variable drive characterisitics. Another option was to lay a track for a motorised carriage to run on. These mechanical approaches were rejected as being impractical, and additionally we preferred to use a human tester as this would more realistically reflect the use of the positioning system in practice. Pressure sensing mats were also rejected as being unable to give the resolution we required, and also as being a logistic challenge to safely cover a public area of over 200 m^2.

All the tests were thus based on a tester carrying or wearing the appropriate sensor around a predetermined trail with a total distance of around 60 metres in the area in which the positioning system is to be tested. The tester is required to complete the route as closely as possible to a period of two minutes. The trail is designed to include areas which are known to have problematic position sensing. Though the route might be typical of a user, the overall speed at which it is undertaken is faster than would be normal. This has the advantage of making testing more efficient, and it makes the test more challenging. The tester stops for five seconds at each of eight marked reference spots around the area and tries to walk at a pre-determined speed between them. The route and timing are designed to simulate a typical user moving between points of interest and stopping at each of these points.

Synchronisation of the tester and the desired route is assisted by the use of a 1 Hz metronome. The tester uses the metronome beats to check that progress corresponds with the specified route and timings. The trail is shown as the continuous line in Figures 1, 2 and 3. The diamond markers show the actual positions as measured by the positioning system; they are connected sequentially by the dotted lines.

This procedure was adopted for testing positioning accuracy for the three installations. These were designed to enable the comparison of direct and indirect ultrasonics and GPS; to establish the effect that the position of the receiving sensor/antenna on the body has on the accuracy and the effect of various correction algorithms.

Below we describe the three tested environments in detail.

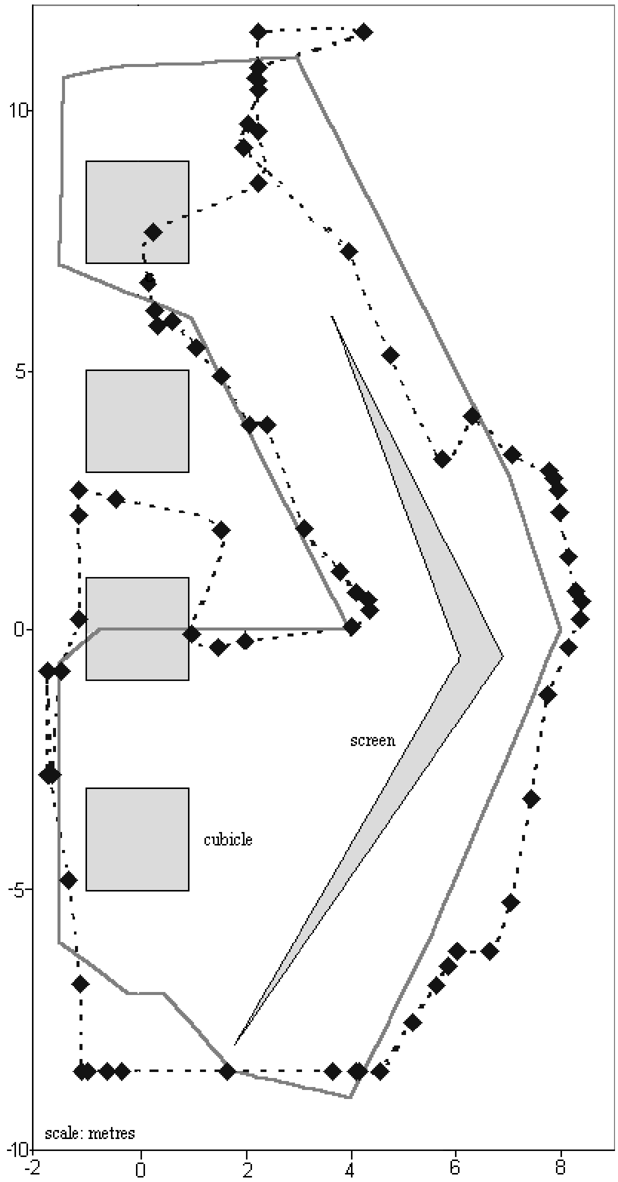

Figure 1. Mackintosh Room Test Path.

2.1 The Mackintosh Room

The Equator City project [7] aims to weave digital media into the physical spaces that people use, and to do this in meaningful ways which show and support their activity. The life and work of Charles Rennie Mackintosh provides a well documented theme which can be explored in and around the City of Glasgow. The Mackintosh Room in the Lighthouse [8] was chosen as a suitable starting point for this project.

The need to determine the physical position of the user in the Mackintosh Room was critical to determining both the user's interests and their path. The 12 by 22 metre Mackintosh Room presents considerable challenges - the 3.5 metre high room is subdivided by a 2.3 metre high screen, contains four 2.1 metre high open cubicles, and any installed positioning system must, for aesthetic reasons, be well hidden - precluding the ceiling mounting of transducers. While an RF system would seem to be appropriate, the achievable accuracy would not be adequate to identify which exhibits in the room were of interest to the user.

Positioning systems have previously been studied and built as part of the Wearable Computing Project at Bristol University [6], though the maximum area previously covered by their indoor system had been 4.2 by 6.5 metres. After exploratory laboratory testing it was found that it would be worthwhile to investigate a positioning system based on ultrasound signals reflecting off the ceiling of the Mackintosh Room. Eight transmitting transducers were subsequently positioned facing upwards on top of the screen and cubicles. This configuration met aesthetic considerations, though it relies on using the signals that bedevil other ultrasound positioning systems. The user is provided with the choice of a handheld receiving sensor or a jacket based wearable computer fitted with ultrasonic receiving sensors on each shoulder, and a Jornada PDA display.

It was recognised that there would be areas where the positioning would not work satisfactorily, for example inside the cubicles and in areas shadowed from the reflected ultrasound. This installation provided an ideal environment to test the limitations of an ultrasound system - and to develop a set of metrics to be able to assess the value of modifications to improve performance. An initial installation demonstrated that the use of reflected ultrasound signals worked in practice, however there were serious problems with range, deadspots, accuracy and jitter.

The test trail was designed to include areas which were known to have inadequate ultrasonic coverage - including the very limits of coverage and passing through one of the cubicles. The trail and results are shown in Figure 1.

In addition static measurements were taken, and finally the timings between physically entering a trigger zone, the trigger being generated, and corresponding web-page refresh signals were measured. Tests were carried out comparing different receivers, correction algorithms and the effectiveness of triggering.

2.2 The Wired Woods

As part of an exhibition of photographs of local woodlands, titled *A Walk in the Wired Woods* hosted by Hewlett-Packard's European Laboratories in Bristol, soundtracks were created which could be associated with each photograph. These soundtracks comprised of various pieces of music, woodland sounds and spoken narrative. The visitor to the exhibition was provided with a belt worn computer which housed a Jornada PDA connected to an ultrasonic receiver. The transducer for the receiver was mounted on the top of headphones worn by the visitor. Using the Bristol University ultrasonic positioning system with conventional overhead transmitting transducers, it was possible to play the soundtrack associated with each photograph when the visitor stood in front of it. The exhibition was situated in a 12 by 20 metre atrium, which apart from eight columnar decorative trees, provided a near ideal environment for a direct ultrasound installation. The floor area was subdivided into 36 zones or 'aura' ranging in size from half a metre in diameter to 4 metres. A different soundtrack was associated with each zone and played when the user entered the zone.

The ultrasonic installation was comparable with that of the Mackintosh Room, both used eight transducers to provide coverage over a similar area, though in this case it was possible to mount the transducers facing downwards from a height of 4.8 metres.

To test the system a similar trail to that used in the Mackintosh Room was devised with eight reference points where the tester stopped for five seconds at each one. The total test time was again 120secs. Ten sets of readings were recorded using a wearable computer again with shoulder mounted sensors. The trail and a sample result can be seen in Figure 2.

2.3 The Millennium Square

In this senario we employed GPS to provide position sensing for a wearable computer in Bristol's Millennium Square - a 65 by 55 metre public square in the centre of the city. A spatialised sound experience has been designed to simulate musicians playing at different points around the square. The GPS system is used to track users wearing 'CyberJackets' containing a StrongARM based processor interfaced wirelessly to a server. The position and orientation of the user's head is used to control a spatialised sound mix of audio streams being sent to the user's headset as the square is explored.

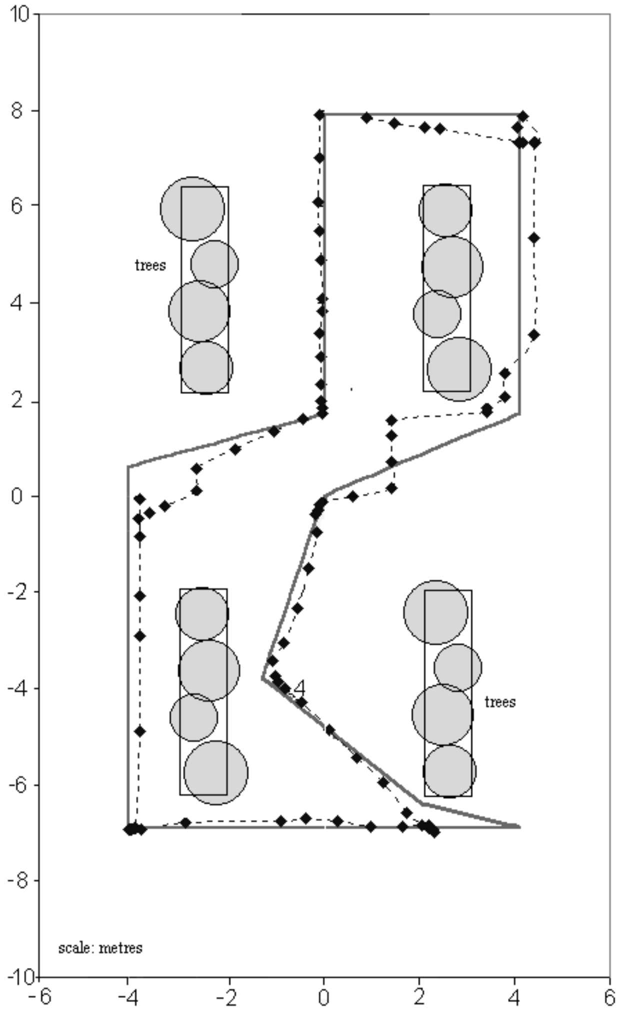

Figure 2. H-P Labs Trail.

A feature of the Millennium Square is the Zenith artwork - a 50 metre analemma. An analemma is a figure of eight construction which results from using the shadow lines produced by a gnomon recorded at noon throughout the calendar year. This provides a test route similar in scale to those used in the previous positioning experiments. Before carrying out the tests a precise survey of the Millennium Square was carried out by placing a GPS receiver at four positions around the square and recording the mean readings over 24hours at each position. The GPS accuracy was enhanced by using differential GPS (dGPS) error correction. By using four reference points we estimated that accuracy of the resulting positions was better than one metre. This was consistent with our static dGPS 24 hour tests which showed a 50% error of two metres. Experiments were carried out with differential GPS error correction messages sent wirelessly to the wearable computer using TCP/IP over 802.11b. Tests were carried out with the GPS antenna placed on different parts of the body - the outstretched hand, on the head, shoulder and waist.

Each test was carried out by walking once around the analemma carrying a GPS receiver. The total time taken for each test was again two minutes. The tester stopped for five seconds at the eight predetermined positions on the analemma. Ten sets of readings were recorded using the cyberjacket. The trail and a sample result can be seen in Figure 3.

3 Evaluation Methods

System errors were identified as those caused by dead spots, range and shadowing limitations, geometric inaccuracy, jitter and sensor placement. However other errors were also introduced by the testing procedure. These were primarily:

- Tester positioning errors caused by poor synchronisation of the tester and the specified test route - or right place, wrong time;
- Device positioning errors caused by poor correspondence between the test receiver and the test route - or simply wrong place;
- Static errors caused by the limits on positioning of the sensors, and of the reference positions - or surveying inaccuracies .

The determination, and reduction, of the system errors are our main objective. To achieve this, and to produce results that may be comparable with other installations, we need to understand and to take account of these other components itemised above.

If the tester is one second late in reaching a reference spot when travelling at 1m/s, this will result in a one metre error in a synchronised analysis. The time taken to walk the

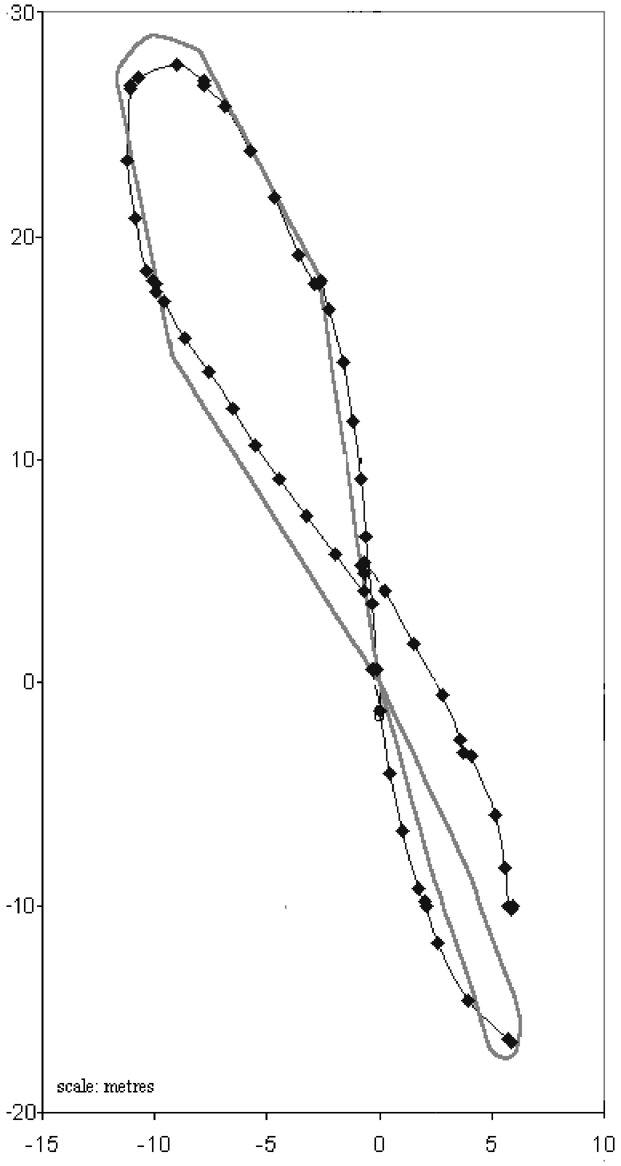

Figure 3. Analemma Test Path.

test route in the Mackintosh Room ranged between 116s and 122s. The average speed required over the route is 0.48m/s. It is likely thus that mistiming generated errors in the order of 2m.

The position of the receiving sensor in handheld mode is around 40cm from the centre of the body. For the wearable computer with its two receiving sensors one on each shoulder the potential mean error is around 20cm.

Whilst every effort was made to position the sensors and reference spots accurately, the limitations or our measuring techniques resulted in further errors of around 10cm (in the order of 1% of the dimensions of the test area).

Initially we attempted to synchronise the measured readings from the Mackintosh Room with the prescribed path. The difference between the current reading and the point at which the tester should have been at on the test route at that time was calculated. This included all the error components noted above, but in particular it was subject to the largest of these - the errors generated by the tester being in the right place at the wrong time.

These errors were analysed to produce 50% readings of around 0.9m and 95% reading of around 2.9m. The 50% figure represented the performance of the system working correctly, and the 95% reading illustrated the performance of the system when the tester entered an area of no coverage. Neither of these figures adequately reflected the overall performance of the system.

These initial measurements provided two critical axioms which needed to be taken into consideration for our proposed measurement techniques.

- The measurement should include all readings without giving overdue emphasis to particular circumstances such as areas of exceptionally good or poor accuracy;
- The measurement should be independent of time related factors introduced by the pace of the tester.

Using these axioms, three different approaches to analysing the error data were explored. First we simply analyse the distance from the readings to the path; secondly a standard deviation measure is derived which takes into account errors along the path as well as the distance from the path; and thirdly a genetic algorithm is employed to optimise the probable locations on the path where the tester should have been at the time the readings were taken.

3.1 Path Deviation

The first approach was to compute the deviation from the specified path. This eliminates timing errors, but does not give an error reading if there are errors along the path. This gives artificially favourable results and in effect is a one dimensional reading, the dimension being the normal to the path.

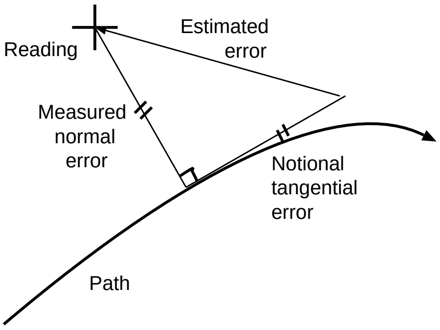

Figure 4. Estimated error calculation.

The results for the Mackintosh Room were 0.52 metres 50% and 1.83 metres 95%. The 50% result appears overly favourable. The 95% result is much harder to analyse because readings are included from a cubicle where there is no effective signal present, and from corners of the room where the signal is intermittent due to range/shadowing problems. The large error in these areas represent the performance when the positioning system is required to handle the complete loss of signals.

This technique is fallible as it makes allowance for errors which are along the tester's path, and can produce excellent results if the tester simply remains motionless at any single point on the path.

3.2 Projected Standard Deviation

The use of the standard deviation measure addresses the first of the stated axioms - that all readings should contribute to the resulting figure. The second axiom is addressed by considering the path of the tester as a measure of time. By using the normal distance from the path we have a time independent measure. This, however, becomes a one dimensional measure and does not include any error tangential to the path as in the path deviation method. Given a symmetrical deployment of transmitting transducers or satellites, it is probable that the tangential error, when expressed as a standard deviation σ_t, is equal to the standard deviation measurement in the direction of the normal σ_n .

We thus propose that for dynamic positioning tests 'projected standard deviation', σ_p, should be measured as the normal standard deviation from the path, σ_n, multiplied by the square root of two. This is illustrated for a single reading in Figure 4 and for the overall standard deviation in the equation:

$$\sigma_p = \sqrt{\sigma_n^2 + \sigma_t^2} = \sqrt{2}\sigma_n$$

As σ_n is calculable we have a measure of the performance of the system independent of the potentially greatest source of error - the testers timing precision - and which encompasses the performance of the system in all parts of the area under test.

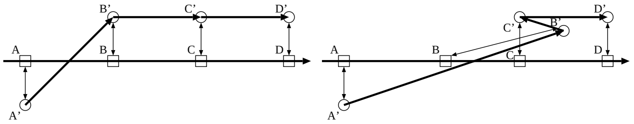

Figure 5. Use of the Genetic Algorithm.

This approach has the merit of being relatively easy to calculate with readily measurable components. Unfortunately it is still not difficult to achieve misrepresentative results by simply remaining static at any point on the path. For this test to be entirely robust there must be some factor which enforces the reference points to move along the path. The use of genetic algorithms to optimise a controlled progression along the path is explored next.

3.3 Genetic Algorithm

In order to establish the distance between the path and the points we would ideally match each measured point up with a point on the path. The basics of the Genetic Algorithm (GA) approach is to estimate points along the path where the wearer can reasonably be expected to be, and then calculate a standard deviation.

If we assume that the wearer is never walking backwards along the path, and is bound by an upper speed, then we can use a genetic algorithm to massage the points along the path in such a way that we minimise the standard deviation. That is we try to find the best possible distribution of points along the path which do not violate the constraints of moving forwards only below a maximum speed.

As an example consider Figure 5. On the left hand side we have shown an observed path (A',B',C',D') and the expected path. The points A,B,C,D are the closest points on the expected path to the observed path and are used for the calculation of the standard deviation. On the right hand side we have shown the situation where observation B' is way out of line. Instead of taking the shortest distance to the path, the closest point where the observer could have been realistically must be between A and C, as close as possible to B yet not too far from A in order not to violate the maximum speed limitation. This will give us a more realistic error measure.

Though not a critical factor it is worth noting that the Genetic Algorithm was computationally expensive and required around one hour of processing time to complete the analysis for each test.

	50%	95%	σ_p	GA
Mack Rm	0.52m	1.83m	1.29m	1.42m
HP Labs	0.25m	0.62m	0.45m	0.70m
Mill. Sq.	1.34m	2.45m	2.16m	1.54m

Table 1. Summary of measurements

4 Results

Consistent and reliable results were obtained for the ultrasonic test paths and samples of these are summarised in Table 1. Results for the dGPS tests are also shown however these should be treated with caution as satellite configuration varies continuously and can affect accuracy. To properly evaluate the measurement strategies we would need more data - more paths and more runs over each path. However, even with this limited data, we can already see the weaknesses of various approaches.

For the HP test, the σ_p value is 40% smaller than the GA value, whereas for the Millennium Square experiment, the σ_p value is 40% higher than the GA value. This is to be expected, since a large part of the deviation in the HP test is in the lower right hand corner of Figure 2: the user appears not to be progressing, but stays on the path. This will not count as an error in the first three columns, but the GA picks it up. In the case for the Millennium Square experiment, the user is progressing roughly along the path, and the closest distance between the paths are very close to what the GA calculates as the most likely distances. In the mackroom, the error is dominated by the large discrepancies in the top left hand corner, which are flagged by both the σ_p and GA measures.

In the Mackintosh Room it was shown that it is possible to track a user with around 0.52 metre accuracy 50% of the time. We also generated reliable information triggers within one second of entering a trigger zone, however, we failed to find an algorithm which will track users while they are inside the cubicles. We are also operating at the limit of the current system when the user is in the corners of the room. Under these conditions spurious readings were obtained and are reflected in the 95% figure of 1.83 metres. Neither the 50% or the 95% readings reflected the overall performance of the system which users found ‘useful’ though some ‘funnies’ were reported. The σ_p and GA measures both gave a similar reading of around 1.4 metres which better reflect the overall performance.

Using the direct ultrasound system at H-P Labs, the Path Deviation 50% and 95% and Projected Standard Deviation σ_p values were derived for similar conditions as those explored in the Mackintosh Room. Again, the 95% figure for the direct ultrasound signal gives a measure of the system’s performance when there is a poor signal, that is when blocked by the columnar trees, and the 50% figure represents the areas where there is good coverage. The σ_p derived using the $\sqrt{2}\sigma_n$ formula gives a more representative value. The result from the GA appears overly pessimistic and it should be noted that in formal user tests feedback from the visitors was overwhelmingly positive. A majority of interviewees reported feeling immersed and compared the experience to be most like ‘a walk in the woods’. The tests confirmed that it is possible to create a convincing and compelling experience with these levels of accuracy [9]. Here we are especially interested in comparing the results from a direct ultrasound system with the reflected signals in the Mackintosh Room. The 50%, 95%, σ_p and GA figures show comparable improvements from which we can conclude that using reflected signals worsens the accuracy of the system by approximately 0.75 metres.

Mounting	50%	95%	σ_p
Handheld	1.03m	3.24m	2.25m
Shoulder	1.49m	3.46m	2.63m
Head	1.34m	2.45m	2.16m
Waist	2.88m	9.84m	6.81m

Table 2. Results of different GPS antennae mounting positions

The dGPS tests at the Millennium Square analemma gave surprisingly good results. The 50% figure of 1.34 metres and the GA figure of 1.54 metres certainly did not seem representative, while the 95% and σ_p were certainly more realistic. The system accuracy was sufficient to determine which side of the ‘big’ end of the analemma that the user is on, but was insufficient to do this at the ‘small’ end. The 50% and GA figures would imply otherwise. The performance of the dGPS positioning exceeded our expectations and that necessary for the application. The characteristics of the GPS system - with continuously varying satellite configuration - are such that these results cannot be regarded as absolute, however they do have value for comparative purposes. For consistent results it would be necessary to perform the testing over a period of 24 hours, as is done with static testing. Unsurprisingly this was considered impractical.

A further example of our tests is shown in Table 2. We are interested in the difference of mounting antennae on different parts of the body. We have applied our measuring techniques to show that while head mounting generally gives the most accurate results, shoulder mounting is nearly as good, and is recommended as a satisfactory compromise for our wearable computer applications.

We also explored the use of speed limiting algorithms which exhibited a clear tendency to worsen the performance of the system when the speed was set too high or too low. As an example the accuracy at 2 metres/sec was 50% better than at 1 metres/sec. This is a further example of how these techniques may be used to optimise positioning systems.

5 Conclusion

The three scenarios under which the measuring techniques were tested, provided useful comparisons between differing technologies and user applications. They also provided a useful framework for comparison of measurement techniques.

The use of the genetic algorithm was found to yield similar results to the simply calculated projected standard deviation. There was a significant difference in the case of the direct ultrasound installation at H-P Laboratories where the GA result was overly affected by a single problematic part of the path. The genetic algorithm approach is certainly more robust, however the projected standard deviation technique will give similar results if it is carried out in a conscientious manner. The simple safeguard of ensuring that each reference point on the path follows the previous point improves the robustness of this technique. The 50% and 95% path deviation measurements clearly did not give fully supportable results, however as a quick measure of performance they were useful if their limitations were taken into account.

Using these measurement techniques we were able to refine our correction algorithms, and to compare the two ultrasonic installations with the conclusion that the use of the ultrasonic reflections introduces an error in the order of 0.75 metres. The placement of the GPS antenna on the shoulder of the user was confirmed as the most satisfactory compromise between accuracy and comfort. We were also able to achieve dGPS accuracy of a similar order to the reflected ultrasound accuracy, and confirm with user tests that both of these were satisfactory for our particular sample applications.

Further system testing continues using the projected standard deviation measure with the additional safeguard of ensuring that the points on the path used for error measurement are sequential. While being useful in establishing the position of our static reference points, the five second stops have been removed from the test procedure. These stops prevented simple checking to ensure that the readings occurred sequentially along the path. Though omitting the stops makes the testing less realistic, it should provide for a more robust analysis to be carried out with relative ease.

6 Acknowledgements

The research in the Mackintosh Room would not have been possible without the assistance and participation of Ian MacColl of the University of Glasgow. Thanks also go to Lynne Campbell and the other Lighthouse staff, and visitors, without whose assistance and cooperation the tests could not have taken place.

In addition the authors wish to recognise and thank the many people involved in the projects described in this paper. These include Matthew Chalmers, Barry Brown and Areti Galani from the University of Glasgow and Anthony Steed from University College, London, for their work in the Mackintosh Room; Phil Stenton, Richard Hull, Jo Reid and Erik Geelhoed of Hewlett-Packard for their contributions to the installation and user testing at the Hewlett-Packard Laboratories in Bristol; and Mike McCarthy and Andrew Moss of the University of Bristol, for their assistance and contributions in the Millennium Square.

We are also very grateful to the anonymous paper reviewers for their helpful comments and suggestions.

References

[1] S. Log, D. Aust, G. Abowd, and C. Atkenson. Cyberguide: Prototyping context-aware mobile applications. In *CHI 96*, pages 293–294, April 1996.

[2] S. Feiner, B. MacIntyre, T. Hollerer, and A. Webster. A touring machine: Prototyping 3d mobile augmented reality systems for exploring the urban environment. In *Proceedings of The First International Symposium on Wearable Computers*, pages 74–81, October 1997.

[3] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. In *IEEE Personal Communications Magazine*, volume 4 no.5, pages 42–47, October 1997.

[4] Polhemus home page. http://www.polhemus.com/.

[5] J. Hightower and G. Borriello. Location systems for ubiquitous computing. *Computer*, pages 57–66, August 2001.

[6] C. Randell and H. Muller. Low cost indoor positioning system. In *UbiComp 2001: International Conference on Ubiquitous Computing*, pages 42–48, September 2001.

[7] The equator interdisciplinary research challenge. http://www.equator.ac.uk/.

[8] The lighthouse, scotland's centre for architecture, design and the city. http://www.lighthouse.co.uk/.

[9] R. Hull, J. Reid, and A Kidd. Experience design in ubiquitous computing. Technical Report HPL-2002-115, HP Labs, 2002.