

# User-Friendly Surveying Techniques for Location-Aware Systems

James Scott<sup>1</sup> and Mike Hazas<sup>2</sup> \*

<sup>1</sup> Intel Research Cambridge

<sup>2</sup> Lancaster University

james.w.scott@intel.com

**Abstract.** Many location-aware applications rely on data from fine-grained location systems. During deployment such systems require a *survey*, specifying the locations of their environment-based components. Most current surveying methods are time-consuming, and require costly and bulky equipment.

This paper presents the concept of *self-surveying*, i.e. methods by which a location system can survey itself. Such methods are user-friendly, fast, and require little or no extra equipment. Experimental results show self-survey accuracies comparable to the accuracy of the underlying location system.

## 1 Introduction

Ubiquitous computing applications often make use of location information. There has been a recent focus on *fine-grained* location systems, which are capable of accuracies down to several centimetres. Unfortunately, existing location systems, particularly those that are fine-grained, typically require much effort and expertise to deploy. This presents a significant obstacle to getting location systems (and consequently ubiquitous applications) out of research labs and into homes and offices. While it is foreseeable that miniaturisation, power management, and wireless communication will facilitate user-friendly *installation* of location system infrastructure, the requirement for easy *surveying* of location systems has not been widely discussed.

With regard to location systems, a survey is a collection of data concerning the physical configuration of the system infrastructure. Such infrastructure often comprises ceiling- or wall-mounted components, and a survey might include accurate positions and/or orientations of these components. In the fine-grained location systems developed and installed in research labs, the surveying procedures used often rely on specialised methods and equipment, which may not be suitable for wide deployment scenarios. Development of simple, inexpensive, and user-friendly surveying methods is therefore important.

This paper discusses the concept of *self-surveying*, which is the ability of a location system to perform a survey of itself, using its own infrastructure and components.

---

\* The experiments discussed in this paper were conducted while both authors were members of the Laboratory for Communication Engineering (LCE), University of Cambridge.

## 2 Related Work

This section presents some conventional surveying methods, and describes prior work in self-surveying. First, however, some fine-grained location systems are briefly reviewed.

In the *Constellation* [1] system, a mobile unit containing ultrasonic sensors is worn on the head and belt of the user, and transmitters are placed in the environment. By using the times-of-flight of ultrasonic pulses, the user is tracked with accuracies of about 5 mm. The *Bat* [2] system employs ultrasonic ceiling receivers to track small transmitting tags worn by people and attached to objects, with accuracies of about 3 cm. Adopting a more privacy-oriented approach, *Cricket* [3] uses independent ultrasonic transmitters in the environment and mobile receivers, with reported accuracies between 5 and 25 cm. Additionally, recent research in ultra-wideband radio systems indicates their potential for accurate tracking [4].

The *EasyLiving Person Tracker* [5] uses two stereo vision cameras to track up to three people in a room. However, this system is not able to directly identify the subjects. The *TRIP* system [6] uses a single camera to detect and identify circular barcode tags worn by people and attached to objects. Provided a tag is within 3 m of a camera, its position can be estimated with an accuracy of 10 cm. *HiBall* [7] is a location system designed for augmented and virtual reality. Users each wear a small device consisting of a cluster of lateral effect photo diodes. Arrays of ceiling infrared LEDs are flashed in sequence, allowing the position of each wearable to be estimated with typical accuracies of about 5 mm.

### 2.1 Conventional Surveying

For many fine-grained location systems described in the literature, there has been no significant discussion concerning the surveying procedures used. In practice, either the environmental units are mounted in a deterministic fashion so that their relative positions are known [3], or units are arbitrarily placed and then surveyed using manual measurements [8].

Another class of surveying techniques uses physical devices to aid the surveying process. Two such devices are the *crate* and *theodolite*. The *crate*, developed at AT&T Labs Cambridge, is a purpose-built device comprising a rigid frame with three spring-loaded reels of cable mounted on it; each reel can electronically report the length of its extended cable. By touching the end of each cable to each survey point, trilateration software can determine the points' locations. A theodolite is a device which accurately measures the range and horizontal and vertical angles between itself and a reflector. This tool, normally used by land surveyors, was used at the LCE to survey a *Bat* system deployment.

To assess the accuracy of the *crate* and theodolite, *Bat* system infrastructure in three rooms was surveyed using both devices. Distances between forty ceiling unit pairs were measured by hand using a tape measure. The difference between the hand-measured distances and the distances according to the surveys averaged 4 mm for both devices. While both surveying methods are accurate, they are time-consuming (a typical room took one person-hour to survey) and the devices are expensive and cumbersome, and not suitable for user-executed deployments in areas such as homes.

## 2.2 Previous Research on Self-Surveying

There has been some research pertaining to self-surveying with particular fine-grained location systems. Foxlin et al. [1] proposed an *auto-mapping* algorithm whereby three seed beacons are placed at measured locations and other beacons are placed arbitrarily. As a user moves through the environment the tracking system can slowly discover and position the arbitrarily-placed beacons. With regard to ultrasonic location systems, Ward commented that it would be possible to use a non-linear multilateration algorithm to estimate unknown ceiling receiver positions by gathering a number of transmitter-to-receiver distances from transmitting tags placed at known locations [8, page 59].

Gottschalk et al. proposed and implemented an *auto-calibration* method for the Hi-Ball system [9]. Using rough LED location estimates, and thousands of observations from unknown test locations, the system (1) estimates the wearable's position at each test location, and (2) calculates back-estimates of LED positions. These two steps are repeated until LED position estimates converge. In more recent work with HiBall, Kalman filters have been used to accomplish on-line auto-calibration [7]. The methods allow the system to continually update LED location estimates during normal operation.

The work described above is concerned with specific location systems. In contrast, this paper discusses general principles of self-surveying, and describes self-surveying concepts and techniques which are applicable to many types of location system.

## 3 Self-Surveying

Most fine-grained location systems operate by using the surveyed locations of *fixed units* in the environment in order to calculate the location of *mobile units*, such as people or equipment. The survey may include information on either or both of the fixed units' positions and orientations. For example, the vision-based TRIP system relies on knowing the orientation of the fixed units as well as their positions, whereas the ultrasonic Cricket system relies only on knowing the positions.

In normal operation, location systems use sighting data and known fixed unit locations to calculate a mobile unit's location; i.e. they use a number of known quantities to determine a few unknown quantities. However, sightings may contain more data than is strictly necessary for location purposes. For example, in the Bat system, a minimum of three distances would be necessary to locate a mobile unit, but a given ultrasonic pulse is often "heard" by well over three receivers. This *surplus data* is what makes self-surveying viable. Additional surplus data can be generated by constraining the locations of mobile units during self-surveying, thus reducing the number of unknowns, and allowing more data to be put toward finding the fixed unit locations.

To conduct a self-survey, a system must first *gather* sightings which include surplus data. This data must then be *processed* in some way to determine estimates of the fixed unit locations. In addition, there is often a requirement for *combining* data so that self-surveys of nearby physical spaces can be unified into a single coordinate space.

### 3.1 Data-Gathering Methods

Three data-gathering methods, known as the *people*, *floor* and *frame* methods, were investigated in this research, and are described below.

**People.** A convenient data gathering method is to log sightings from a location system while it is in ordinary use. Sightings can be gathered as people (equipped with mobile units if necessary) move around the space. This method potentially provides a completely transparent surveying process, whereby a new system would automatically be surveyed as it is used, with survey accuracy improving as the system ages. However, the people method provides less surplus data compared to the methods below, since the locations of the mobile units are completely unconstrained.

**Floor.** A second self-surveying method would be to place many mobile units on the floor of a space for a period of time. Unlike in the people method, where each sighting may be taken with the mobile unit at a different location, in the floor method, many sightings are gathered from each stationary mobile unit. This makes the self-survey solution set much more constrained (since the mobile units are approximately coplanar), and also provides the advantage that erroneous sightings can be more easily identified. This method requires a dedicated self-survey measurement phase before the system can be put into use, and the user is required to place and recover the mobile units.

**Frame.** In the frame method, mobile units are placed onto known points on a rigid frame. The frame method's advantage is that it provides the most surplus data of the three methods presented. Since the relative locations of the mobile units are known, the self-surveying algorithm need only solve for the locations of the fixed environmental units. Although this method incurs the disadvantage of requiring hardware for the frame, the frame might be manufactured using inexpensive, lightweight materials for easy assembly.

### 3.2 Processing Self-Survey Data

Once data gathering is complete, the data must be processed to find estimates of the fixed unit locations. Two procedures for accomplishing this are discussed below.

**Simulated Annealing.** One possible method for finding the fixed unit locations is to use an iterative algorithm to find better and better solutions for the mobile and fixed unit locations until a "best guess" is found. For each possible solution, there must be some method of *scoring* the solution against the gathered data. For example, in the case of an ultrasonic location system, a suitable scoring function might be found by taking the sum of the differences between the distances reported by the sightings and the corresponding distances according to the current location estimates.

Simulated annealing [10] is a mathematical optimisation technique which mimics the process of crystallisation in bulk matter through a randomised search procedure, using the notion of *temperature* to avoid being trapped by local minima in the solution space. Simulated annealing was chosen for the purposes of this research because it involves no assumptions about the solution space; it can thus be applied with little modification to all three data-gathering methods.

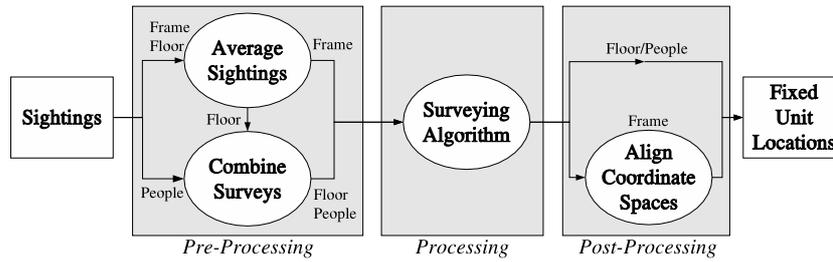


Fig. 1. Data flow for self-surveying

**Inverting the Location Algorithm.** Since the relative mobile unit locations are known when the frame method is used, it is as if the location system were “inverted,” in that a fixed unit must be located using known locations for mobile units. For some location systems, it may be possible to determine the locations of the fixed units by employing the same algorithm that is used for normal operation of the system.

### 3.3 Unification of Coordinate Spaces

One issue that must be dealt with in many surveying procedures is combining the results of multiple surveys with arbitrary coordinate spaces (such as separate rooms) into a single, unified space (such as a building). This is especially significant for a self-surveying system, as the aims of minimising the work required by the users for surveying would be compromised if they were required to manually measure the relative positioning of the various spaces surveyed, so that coordinate spaces could be combined.

For the people and floor methods, automatic unification of coordinate spaces can be achieved by combining the self-surveying data in a pre-processing stage; i.e. by merging multiple data sets so that they may be presented as a single data set to the processing algorithm. For the frame method, combining data from different surveys during pre-processing is not viable in all cases, since the processing method might rely on knowing the relative locations of the mobile units. Therefore, a post-processing combination of surveys must be performed. This may be achieved by calculating coordinate space transforms which map the points in multiple surveys such that the sum of distances between common points is minimised in the unified survey.

The flow of data during the self-surveying process, and the differences in this flow depending on the method used, are illustrated in Fig. 1.

## 4 Implementation

In order to determine the accuracy of self-surveying, the Bat system installation at the LCE was used to conduct self-surveys in a number of rooms of various sizes. In particular, two large offices (5 m × 6 m) and three small offices (5 m × 3 m) were used for data gathering; however, only two of these five data sets were used as test data during development of the self-survey algorithms.

## 4.1 Data Gathered

The data was gathered by capturing raw sighting information from the Bat system. The logging included all distances recorded between the surveying tags and the fixed ceiling units during the course of a survey.

For the people method, data was gathered by having two people, each with four tags, walk around a room for a specified period of time. The tags were worn on the chest and on the back, and one was carried in each hand, the hands being kept a forearm's length away from the body. Readings were taken for a period of ten minutes for a small room, and twenty minutes for a large room.

For the floor method, in small rooms twenty tags were placed randomly but uniformly over the floor area, and left for ten minutes. Large rooms were surveyed as if they were two adjacent small rooms, i.e. in two ten-minute runs each covering opposite halves of the room. The data for large rooms was combined at the pre-processing stage so that it appeared as a single forty-tag experiment.

For the frame method, a frame measuring  $1\text{ m} \times 1\text{ m}$  was equipped with twenty-one tags. In small rooms, the frame was placed in the centre of the room for ten minutes. Large rooms were surveyed by treating them as two adjacent small rooms; the survey results were combined during a post-processing stage, as described in Sect. 3.3.

## 4.2 Pre-processing

The floor and frame methods use a pre-processing stage to obtain an averaged value for each tag/fixed unit distance. This is accomplished using the following procedure: (1) ignore tag and fixed unit pairs which returned fewer than  $t$  distances; (2) for each remaining pair, find groups of distances which make up at least proportion  $p$  of the number of sightings between that pair, and which are within  $d$  of one another; (3) take the mean of the group representing the shortest distance (since longer peaks are likely to be due to reflections). By trial and error, suitable values for the parameters  $t$ ,  $d$  and  $p$  were found to be 10, 7.5 cm, and 40% respectively.

## 4.3 Processing

Both the simulated annealing and inverted processing techniques were implemented. For the simulated annealing algorithm, the starting location of the fixed units is a single location 2 m above the level of the tags. The tags also start at a single location, except when using the frame method, where the tag locations are specified according to the frame geometry. Ten thousand iterations are used for the floor and frame methods, and twenty thousand are used for the people method.

During each iteration, movements are potentially made to the position of each fixed unit and each tag, except that for the frame method, the tag positions are fixed, and for the floor method, the height of the tags is fixed. Each potential movement is chosen from a group of scalar multiples of a random vector; the scalars used are the integers from -4 to 4. For each of the eight possible new positions, errors in the new tag-fixed unit distances are summed; this provides the scoring function. If at least one of the potential movements gives a better score than the current position, the best-scoring

**Table 1.** Per-room self-survey accuracy for various data-gathering and processing methods. Accuracies represent the mean location error of all fixed units in the room. Rooms marked \* were used as test data when developing the self-surveying procedure. All values are in centimetres.

Method	Small Room 1	Small Room 2	Small Room 3*	Large Room 1	Large Room 2*	Mean accuracy	Mean accuracy (central units)
People	20	20	24	17	13	19	15
Floor	12	21	26	9	6	15	5
Frame (SA)	9	9	2	10	5	7	4
Frame (NR)	3	3	2	4	3	3	2

movement is taken. If not, then with probability equal to the temperature, a randomly chosen movement will be taken. The starting temperature is 0.02 for the people method and 0.01 for the frame and floor methods. The temperature declines linearly, reaching zero at the final iteration. At each iteration, if the average distance error is the lowest so far recorded, then the fixed unit locations are recorded.

The above steps are repeated twice, with a *culling* stage in the middle. This stage was introduced to handle erroneous sightings which represent reflections or other errors. While the algorithm largely ignores such errors, the results are still influenced slightly. The culling stage examines each tag’s distances, and discards distances which have a much higher error than the others. A second set of iterations is then run.

For the inverted processing technique, the non-linear regression algorithm used by the Bat system was applied. As described in Sect. 3.2, only the frame data may be used with this method. To distinguish between the frame-based results sets, the names *frame (SA)* for simulated annealing and *frame (NR)* for non-linear regression will be used.

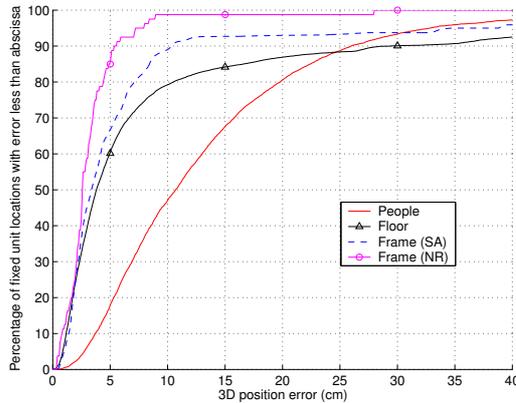
## 5 Experimental Results

For each self-survey, the resulting locations were compared with a theodolite survey. The self-survey accuracies for each room are shown in Table 1. Since simulated annealing produces slightly different results each time it is run, the people, floor, and frame (SA) accuracies given represent the ninetieth percentile.

The results show that the frame (NR) method performs best out of the self-surveying methods, with a survey mean error of 3 cm. Significantly, the large-room experiments exhibited similar accuracies to the small-room experiments, indicating that unification of coordinate spaces during the pre- and post-processing stages did not adversely affect the accuracies of the self-surveying methods studied.

### 5.1 Analysis and Optimisation

In order to analyse the results further, it is useful to examine an error distribution for individual fixed unit positions, as shown in Fig. 2. The diagram illustrates an “elbow” in the frame (SA) and floor results; while 75% of the fixed units are located accurately,



**Fig. 2.** Location error distributions for all fixed units in all rooms

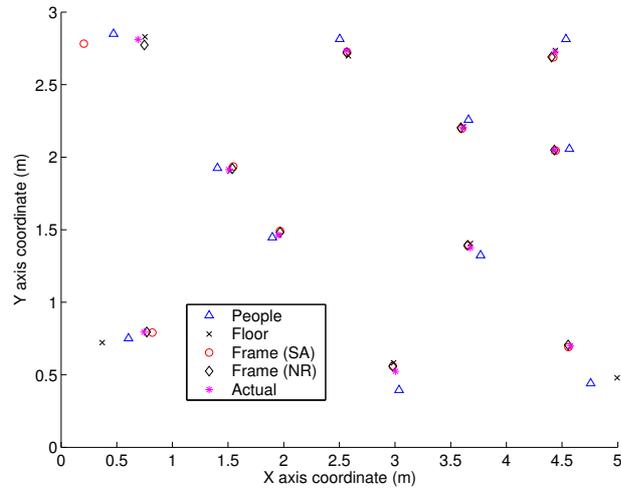
the remaining 25% are for some reason subject to high errors. This is explained by looking at a plan view of one of the rooms, shown in Fig. 3, which shows that large errors occur most often for fixed units located at room extremities. If such “edge units” are removed from consideration, the average floor and frame (SA) room accuracies are greatly improved, as shown in the rightmost column of Table 1.

The reason behind these errors lies in the fact that, especially for the floor and frame methods, the location solutions for edge units will tend to be less rigidly constrained by the data. Such units are likely to have a less favourable geometry with respect to the surveying tags; i.e. they have a poorer position dilution of precision. Also, edge units tend to be at further distances from the surveying tags, and it is known for the Bat system that ranging at distances greater than several metres succeeds less often, and is less accurate. A remedy to both of these problems may be to gather more surveying data over a wider area.

## 5.2 Accuracy Limits

The results, both for the frame (NR) case and for the central units in the frame (SA) and floor cases, show accuracies comparable to the Bat system’s accuracy of 3 cm. This is a very positive result for self-surveying, as it indicates that the accuracy of self-surveying techniques can approach the accuracy of the surveyed system. Note that it is not claimed that a self-surveyed location system will locate objects as accurately as a conventionally-surveyed location system, as the error in self-surveying and error in the location process may combine, either additively or otherwise. Characterisation of this effect is an area for future exploration.

The implication of these results for the deployment of location systems is that there is no single correct surveying method to use. For applications with low accuracy requirements, a minimal-cost, people-based survey may be appropriate. As accuracy requirements increase (e.g. with the deployment of a new application), a more accurate and more labour-intensive survey could be carried out.



**Fig. 3.** Plan view of typical self-survey results for a small room

## 6 Further Work

Some further improvements to specific aspects of the self-surveying procedure have already been proposed. This section looks at broader research directions.

One area for exploration is in different types of processing methods. In particular, the non-linear regression algorithm may be extended to be applicable to the floor and people methods. Alternatively, other processing algorithms, for example based on Kalman filters, may be a fruitful area for investigation.

Another area for future work is in validating the viability of self-surveying in other location systems. The Bat system used in the experiments above employs distances from transmitting tags to receiving ceiling units to determine location. It is easy to see how the same methods would be applicable to other ultrasound-based systems, including systems with roaming receivers and fixed transmitters such as Cricket, and to systems in general which utilise time-of-flight methods, such as ultra-wideband systems.

Vision-based systems can also make use of distance. However, they may instead rely on the angle from the camera's axis to the target, or its relative orientation. The viability of self-surveying in such cases may be seen by noting that vision-based systems can operate with roaming tags and stationary cameras, or vice versa. As with time-of-flight systems, it is therefore possible to conceive that the locations of the fixed units can be regarded as unknown values, and that a self-survey can take place, given enough "surplus data."

Finally, moving away from self-surveying, there are other issues in user-friendly deployment of location systems to be explored. One such issue is that of finding the physical configuration of a space outfitted with a location system. This includes finding the locations of walls and doors, of desks and other furniture, and of interactive objects such as computers and telephones. Without this information, location-aware ap-

plications cannot assist in users' interactions with their environment. Further work is therefore required in this area; one interesting thread of research involves the use of ray-tracing to model horizontal surfaces in the environment [11].

## 7 Conclusions

This paper has discussed the concept of *self-surveying* for indoor fine-grained 3D location systems, and has described several self-surveying methods. Experimental results using the Bat system have shown that self-surveying is capable of reaching survey accuracies similar to that of the underlying location system, in this case 3 cm.

Self-surveying facilitates the deployment of location systems by untrained personnel, with minimal overhead in time and hardware costs. By contrast, conventional surveying techniques have required trained use of high-cost and bulky equipment. Self-surveying is therefore a key step toward user-friendly deployment of location systems, and enabling the use of ubiquitous computing applications in everyday environments.

**Acknowledgements.** The authors would like to thank Rob Harle for his help with the Bat system and insightful comments on this work, Alastair Beresford and Alastair Tse for their work on theodolite surveying, the reviewers for their helpful comments, and Andy Hopper for his support.

## References

1. Foxlin, E., Harrington, M., Pfeifer, G.: Constellation: A wide-range wireless motion-tracking system for augmented reality and virtual set applications. In: Proceedings of SIGGRAPH '98, Orlando, Florida, USA (1998) 371–378
2. Addlesee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward, A., Hopper, A.: Implementing a sentient computing system. *IEEE Computer* **34** (2001) 50–56
3. Priyantha, N.B., Miu, A.K.L., Balakrishnan, H., Teller, S.: The Cricket Compass for context-aware mobile applications. In: Proceedings of MobiCom 2001, Rome, ACM/IEEE (2001)
4. Lee, J.Y., Scholtz, R.A.: Ranging in a dense multipath environment using an UWB radio link. *IEEE Journal on Selected Areas in Communications* **20** (2002) 1677–1683
5. Krumm, J., Harris, S., Meyers, B., Brumitt, B., Hale, M., Shafer, S.: Multi-camera multi-person tracking for EasyLiving. In: Proceedings of the Third IEEE International Workshop on Visual Surveillance, Dublin, Ireland (2000)
6. López de Ipiña, D., Mendonça, P., Hopper, A.: TRIP: A low-cost vision-based location system for ubiquitous computing. *Personal and Ubiquitous Computing* **6** (2002) 206–219
7. Welch, G., Bishop, G., Vicci, L., Brumback, S., Keller, K., Colucci, D.: The HiBall tracker: High-performance wide-area tracking for virtual and augmented environments. In: Proceedings of the ACM Symposium on Virtual Reality Software and Technology, London (1999)
8. Ward, A.: Sensor-Driven Computing. PhD dissertation, University of Cambridge (1998)
9. Gottschalk, S., Hughes, J.F.: Autocalibration for virtual environments tracking hardware. In: Proceedings of SIGGRAPH '93, Anaheim, California, USA, ACM (1993) 65–72
10. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. *Science* **220** (1983) 671–680
11. Harle, R., Hopper, A.: Building world models by ray-tracing within ceiling-mounted positioning systems. In: Proceedings of UbiComp 2003, Seattle, USA, Springer-Verlag (2003)