International Journal of Computer Integrated Manufacturing
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713804665

Sources of variability in the set-up of an indoor GPS

Carlo Ferri[a]; Luca Mastrogiacomo[b]; Julian Faraway[c]
[a] Via XI Febbraio 40, Castelli Calepio, BG, Italy [b] DISPEA, Politecnico di Torino, Torino, Italy [c] Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY, UK

Online publication date: 21 May 2010

To cite this Article Ferri, Carlo , Mastrogiacomo, Luca and Faraway, Julian(2010) 'Sources of variability in the set-up of an indoor GPS', International Journal of Computer Integrated Manufacturing, 23: 6, 487 — 499
To link to this Article: DOI: 10.1080/09511921003642147
URL: http://dx.doi.org/10.1080/09511921003642147

International Journal of Computer Integrated Manufacturing
Vol. 23, No. 6, June 2010, 487–499

Taylor & Francis
Taylor & Francis Group

Sources of variability in the set-up of an indoor GPS

Carlo Ferri[a]*, Luca Mastrogiacomo[b] and Julian Faraway[c]

[a]*Via XI Febbraio 40, 24060 Castelli Calepio, BG, Italy;* [b]*DISPEA, Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino 10129, Italy;* [c]*Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, UK*

(*Received 17 June 2009; final version received 5 January 2010*)

An increasing demand for an extended flexibility to model types and production volumes in the manufacture of large-size assemblies has generated a growing interest in the reduction of jigs and fixtures deployment during assembly operations. A key factor enabling and sustaining this reduction is the constantly expanding availability of instruments for dimensional measurement of large-size products. However, the increasing complexity of these measurement systems and their set-up procedures may hinder the final users in their effort to assess whether the performance of these instruments is adequate for pre-specified inspection tasks. In this paper, mixed-effects and fixed-effects linear statistical models are proposed as a tool to assess quantitatively the effect of set-up procedures on the uncertainty of measurement results. This approach is demonstrated on a Metris Indoor GPS system (iGPS). The main conclusion is that more than 99% of the variability in the considered measurements is accounted for by the number of points used in the bundle adjustment procedure during the set-up phase. Also, different regions of the workspace have significantly different error standard deviations and a significant effect on the transient duration of measurement. This is expected to affect adversely the precision and unbiasedness of measurements taken with Indoor GPS when tracking moving objects.

Keywords: large scale metrology; large volume metrology; distributed coordinate measuring systems; Indoor GPS; iGPS; uncertainty

1. Introduction

During the last decades research efforts in coordinate-measuring systems for large-size objects have led to a broadening of the range of instruments commercially available (cf. Estler *et al*. 2002).

These coordinate measurement instruments can be grouped into two categories: centralised and distributed systems (Maisano *et al*. 2008).

A centralised instrument is a measuring system constituted by a single hardware element that in performing a measurement may require one or more ancillary devices such as, typically, a computer. An example of a centralised instrument is a laser tracker that makes use of a spherically-mounted reflector (SMR) to take a measurement of point spatial coordinates and that needs to be connected to a monitor of environmental conditions and to a computer.

A distributed instrument is a collection of separate independent elements whose separately gathered measurement information needs to be jointly processed in order for the system to determine the coordinates of a point. A single element of the system typically cannot provide measurements of the coordinates of a point when standing alone. Precursors of these apparatuses can be identified in wireless indoor networks of sensors for automatic detection of object location (cf. Liu *et al*. 2007). These networks can be deployed for inspection tasks in manufacturing operations once their trueness has been increased. The term trueness is defined in BS ISO 5725-1:1994 as 'the closeness of agreement between the average value obtained from a large series of test results and an accepted reference value' (Section 3.7).

When inspecting parts and assemblies having large dimensions, it is often more practical or convenient to bring the measuring system to the part rather than vice versa, as is typically the case on a smaller scale. Therefore, instruments for the inspection of large size objects are usually portable. In performing a measurement task, a single centralised instrument, say a laser tracker, can then be deployed in a number of different postions which can also be referred to as stations. By measuring some fixed points when changing station, the work envelope of the instrument can be significantly enlarged enabling a single centralised instrument to be used for inspection of parts significantly larger than its original work envelope. To illustrate this concept, in Figure 1(a) the top view of three geometrical solids, a cylinder, a cube and an octahedron (specifically a hexagonal prism) is displayed. These

*Corresponding author. Email: info@carlo.comyr.com

ISSN 0951-192X print/ISSN 1362-3052 online

DOI: 10.1080/09511921003642147
http://www.informaworld.com

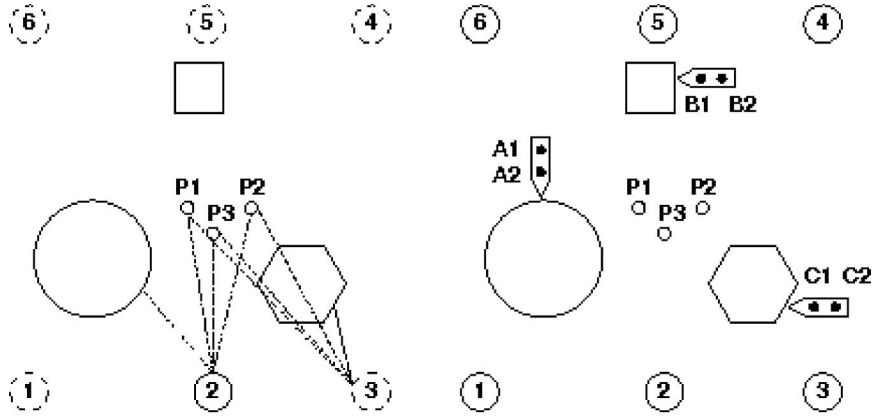

(a) Centralised. (b) Distributed.

Figure 1. Centralised and distributed measurement systems.

solids are inspected by a single centralised instrument such as a laser tracker, which is moved across different positions (1, 2, ..., 6 in the figure) from each of which the coordinates of the points P1, P2 and P3 are also measured. In this respect, a single centralised system appears therefore comparable with a distributed system, whose inherent multi-element nature enables work envelopes of any size to be covered, provided that a sufficient number of elements are chosen. This characteristic of a measuring system of adapting itself to suit the scale of a measuring task is often referred to as scalability (cf. Liu *et al*. 2007). The concept above can therefore be synthesised by saying that a centralised system is essentially scalable in virtue of its portability, whereas a distributed system is such due to its intrinsic modularity.

With a single centralised instrument, measurement tasks within a working envelope, however extended, cannot be performed concurrently but only serially. Each measurement task to be performed at a certain instant in time needs a dedicated centralised instrument. This is shown in Figure 1(a) where the cylinder is measured at the current instant with the instrument in position 2, whereas the hexagonal prism is going to be measured in a future instant when the instrument will be placed in position 3. With a distributed system this limitation does not hold. With a distributed system, concurrent measurement tasks can be performed provided that each of the concurrent tasks has a sensor or subgroup of sensors dedicated to it at a specific instant within the distributed instrument. In Figure 1(b), the same three objects considered in the case of a centralised instrument are concurrently inspected using a distributed system constituted by six signal transmitter elements (1, 2, ..., 6) and three probes, each carrying two signal receiving elements whereby the coordinates of the probe tips are calculated.

This characteristic of distributed systems is especially advantageous when concurrently tracking the position of multiple large-size components during assembly operations. The sole way of performing the same concurrent operation with a centralised system would require the availability and use of more than a single centralised instrument (laser tracker, for instance), with potentially-detrimental economic consequences on the manufacturing organisation in terms of increased fixed assets, maintenance costs and increased complexity of the logistics.

A number of different distributed systems have been developed recently, some as prototypes for research activities (cf., for instance, Priyantha *et al*. 2000; Piontek *et al*. 2007), some others with a level of maturity sufficient for them to be made commercially available (cf., for instance, Welch *et al*. 2001; Maisano *et al*. 2008). In this second case, the protection of intellectual property (IP) rights prevents users' transparent access to the details of the internal mechanisms and of the software implemented in the systems. This may constitute a barrier to a full characterisation of the performance of the equipment. This investigation endeavours to provide better insight into the performance of such systems by using widespread statistical techniques. The main objective is therefore not to criticise or evaluate the specific instrument considered thereafter, but to demonstrate the use of techniques that may be beneficially deployed also on other distributed systems. In particular, the effect of discretionary set-up parameters on the variability and stability of the measurement results has been analysed.

In the next section the main characteristics of the Metris iGPS, which is the instrument considered, are described. A cone-based mathematical model of the system is then presented in Section 3. The experimental set-up is described in Section 4 and the results

of the tests are analysed in Section 5. Conclusions are drawn thereafter.

2. Physical description of the instrument

The instrument used in this study is the iGPS (alias indoor GPS) manufactured by Metris. The description of such a system provided in this section is derived from publicly available information.

The elements constituting the system are a set of two or more transmitters, a number of wireless sensors (receivers) and an unit controlling the overall system and processing the data (Hedges *et al*. 2003; Maisano *et al*. 2008).

Transmitters are placed in fixed locations within the volume where measurement tasks are performed. Such a volume is also referred to as a workspace.

Each transmitter has a head rotating at a constant angular velocity, which is different for each transmitter, and radiates three light signals: two infrared fan-shaped laser beams generated by the rotating head, and one infrared strobe signal generated by light emitting diodes (LEDs). The LEDs flash at constant time intervals ideally in all directions, but practically in a multitude of directions. Each of these time intervals is equal to the period of revolution of the rotating head on which the LEDs are mounted. For any complete revolution of the rotating head a single flash is emitted virtually in all directions. In this way, the LED signals received by a generic sensor from a transmitter constitute a periodic train of pulses in the time domain where each pulse is symmetric (cf. Hedges *et al*. 2003, column 6).

The rotating fan-shaped laser beams are tilted by two pre-specified opposite angles, ϕ_1 and ϕ_2 (e.g. -30 and 30°, respectively) from the axis of rotation of the head. These angles are also referred to as slant angles. The fact that the angular velocity of the head is different for different transmitters enables each transmitter to be distinguished (Sae-Hau 2003). A schematic representation of a transmitter at the instant $t1$ when the first fanned beam L1 intersects the sensor in position P and at the instant $t2$ when the second fanned beam L2 passes through P is shown in Figure 2, where two values for the slant angles are also shown. Ideally, the shape of each of the fanned beams should be adjustable to adapt to the characteristics of the measurement tasks within a workspace. Although two beams are usually mounted on a rotating head, configurations with four beams per head have also been reported (Hedges *et al*. 2003, column 5). To differentiate between the two fanned beams on a transmitter, their time position relative to the strobe signal is considered (see Figure 2).

The fanned beams are often reported as planar (Liu *et al*. 2008; Maisano *et al*. 2008), as depicted in Figure 2. Yet, the same beams when emitted from the source typically have a conical shape that is first deformed into a column via a collimating lens and then into a fan-shape via a fanning lens (Hedges *et al*. 2003, column 6). It is believed that only an ideal chain of deformations would transform completely and perfectly the initial conical shape into a plane. For these reasons, the final shape of the beam is believed to preserve traces of the initial shape and to be more accurately modelled with a portion of a conical surface, rather than a plane. Each of the two conical surfaces is then represented by a vector, called a cone vector, that is directed from the apex to the centre of the circular directrix of the cone. The angle between a cone vector and any of the generatrices on the cone surface is called the cone central angle. This angle is

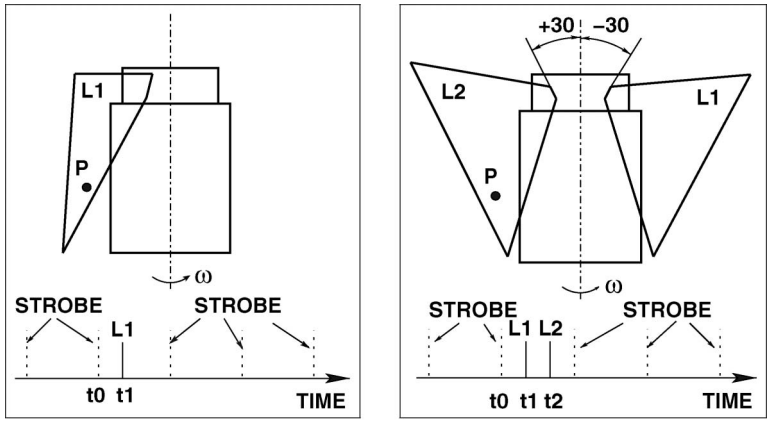

(a) Instant $t1$. (b) Instant $t2$.

Figure 2. Schema of a transmitter at the instants $t1$ and $t2$ when the first and second fanned beams, respectively, intersect the position (P) of a sensor. -30 and $+30^\circ$ are two arbitrary values of the slant angle.

designated by α_1 and α_2 for the first and the second beams, respectively. The apex of each cone lies on the axis of rotation of the spinning head. In Figure 3, a schema of the portion of the conical surface representing a rotating laser beam is displayed. In this figure, two portions of conical surfaces are shown to illustrate α_2 and ϕ_2 ($\phi_2 > 0$, having established counterclockwise angle measurements around the x-axis as positive).

The angular separation between the optical axes of the two laser modules in the rotating head is denoted by θ_{off}, when observed from the direction of the rotational axis of the spinning head. The rotation of the head causes each of the cone surfaces, and therefore their cone vectors, to revolve around the same axis. The angular position of the cone vector at a generic instant is denoted by $\theta_1(t)$ and $\theta_2(t)$ for the first and second fanned beams, respectively. These angles are also referred to as scan angles and are defined relative to the strobe LED synchronisation signal, as illustrated below.

Wireless sensors are made of one or more photodetectors and a wireless connection to the controlling unit for the transmission of the positional information to the central controlling unit. The use of the photodetectors enables the conversion of a received signal (stroboscopic LED, first fanned laser, second fanned laser) into the instant of time of its arrival ($t0$, $t1$ and $t2$ in Figure 2). The time intervals between these instants can then be converted into measurements of scan angles from the knowledge of the angular velocity of the head for each transmitter (ω in Figure 2). It is expected that $\theta_1 = \omega \times (t1 - t0)$ and that $\theta_2 = \omega \times (t2 - t0)$. At the instant $t0$ when the LED signal reaches the generic position P, the same LED signal also flashes in any direction. Therefore, at the very same instant $t0$, the LED fires also in the reference direction where the angles in the plane of rotation are measured from (i.e. $\theta_1 = \theta_2 = 0$).

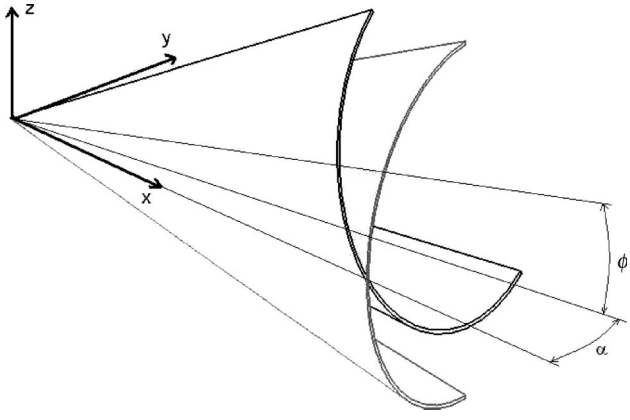

Figure 3. Schema of a shaped laser beam with two portions of conical surfaces to show the central angle α_2 and the slant angle ϕ_2.

In this study, any plane orthogonal to the axis of rotation is referred to as a plane of rotation. For any spherical coordinate system having the rotational axis of the transmitter as the z-axis and the apex common to the aforementioned conical surfaces as the origin, the angle θ_1 swept by the cone vector of the first fanned beam in the time interval $t1 - t0$ is connected with the azimuth of P measured from any possible reference direction x established in the xy-plane, which is the plane of rotation passing through the common apex of the conical surfaces.

From a qualitative point of view, the elevation (or the zenith) of P can be related to the quantity $\omega \times (t2 - t1)$. By analogy with Figure 2, it is argued that, also in the case of conical fanned shaped beams, when the elevation (or zenith) of P is increasing (decreasing), the time interval $t2 - t1$ is also increasing. Vice versa, the reason why a time interval $t2 - t1$ is larger than another can only be found in the fact that the position of the sensor in the first case has a higher elevation than in the second.

In the most typical configuration, two receivers are mounted on a wand or a bar in calibrated positions. A tip of the wand constitutes the point for which the location is calculated based on the signals received by the two sensors. When the receivers are mounted on a bar, the bar is then often referred to as vector bar. If such a receivers-mounted bar is short, say with a length between 100 and 200 mm, it is then called a mini vector bar. These devices are equipped with firmware providing processing capabilities. The firmware enables the computation of azimuth and elevation of the wand or bar tip for each of the spherical reference systems associated with each of the transmitters in the system. This firmware is called a position computation engine (PCE).

A vector bar therefore acts as a mobile instrument for probing points as shown in the schema of Figure 1(b). More recently, receiving instruments with four sensors have been developed, enabling the user to identify both the position of the tip and the orientation of the receiving instrument itself.

3. The role of the bundle adjustment algorithms in the indoor GPS

The computation of the azimuth and elevation of the generic position P in the spherical reference system of a generic transmitter enables the direction of the oriented straight line l from the origin (the apex of the cones) to P to be identified. However, it is not possible to determine the location of P on l. In other words, it is

not possible to determine the distance of P from the origin. Therefore, at least a second transmitter is necessary to estimate the position of P in a user arbitrarily predefined reference system $\{Uref\}$. In fact, assuming that the position and orientation of the ith and jth transmitters in $\{Uref\}$ are known, then the coordinates of the generic point on l_i and on l_j can be transformed from the spherical reference system of the transmitters to the common reference system $\{Uref\}$ (cf. Section 2.3 in Craig 1986). Then, P can be estimated with some nonlinear least squares procedure, which minimises the sum of the squared distances between the estimates of the coordinates of P in $\{Uref\}$ and the generic point on l_i and l_j. Only in an ideal situation would l_i and l_j intersect. In any measurement result, the azimuth and elevation are only known with uncertainty (cf. Sections 2.2 and 3.1 in JCGM 2008). Very little likelihood exists that these measured values for l_i and l_j coincide with the 'true' unknown measurands. The same very little likelihood applies therefore to the existence of an intersection between l_i and l_j. When adding a third kth transmitter, qualitative geometrical intuition supports the idea that the distances of the optimal P from each of the lines l_i, l_j and l_k are likely to be less variable until approaching and stabilising around a limit that can be considered typical for the measurement technology under investigation. Increasing the number of transmitters is therefore expected to reduce the variability of the residuals. The estimation of the coordinates of P, when the position of the transmitters is known, is often referred to as a triangulation problem (Hartley and Sturm 1997; Savvides *et al*. 2001).

If the position and orientation of the transmitters in $\{Uref\}$ are not known, then they need to be determined before the actual usage of the measurement system. To identify the position and orientation of a transmitter in $\{Uref\}$, six additional parameters need to be estimated (cf. Section 2.2 in Craig 1986). This more general engineering problem is often referred to as three-dimensional (3D) reconstruction and occurs in areas as diverse as surveying networks (Wolf and Ghilani 1997), photogrammetry and computer vision (Triggs *et al*. 2000; Lourakis and Argyros 2009). The estimation of three-dimensional point coordinates together with transmitter positions and orientations to obtain a reconstruction which is optimal under a prespecified objective function and an assumed errors structure is called bundle adjustment (BA). The objective or cost function describes the fitting of a mathematical model for measurement procedure to the experimental measurement data. Most often, but not necessarily, this results in minimising the sum of the squares of the deviations of the measurement data from their values predicted with nonlinear functions of the unknown parameters (Triggs *et al*. 2000; Lourakis and Argyros 2009). A range of general purpose optimisation algorithms, such as for instance those of Gauss–Netwon and Levenberg–Marquardt, can be used to minimise the nonlinear objective function. Alternatively, significantly increased efficiency can be gained if these algorithms are adjusted to account for the sparsity of the matrices arising in the mathematical description of 3D reconstruction problems (Lourakis and Argyros 2009).

In the measurement system investigated, a BA algorithm is run in a set-up phase whereby the position and orientation of each transmitter in $\{Uref\}$ are determined. Therefore, during the subsequent deployment of the system (measuring phase), the coordinates of a point are calculated using the triangulation methods mentioned above.

However, as is typically encountered in commercial measurement systems, the BA algorithms implemented in the system are not disclosed completely to the users. This makes it difficult for both users and researchers to devise analytical methods to assess the effects of these algorithms on the measuring system. In this investigation, consideration is given to experimental design and statistical techniques to estimate the effect that decisions taken when running the built-in BA algorithm exert on measurement results.

4. Experimental set-up

Four transmitters were mounted on tripods and placed at a height of about two metres from floor level. The direction of the rotational axis of each transmitter spinning head was approximately vertical. Each of the four transmitters was placed at the corners of an approximate square of side about eight metres.

A series of six different targets fields labelled I, II, III, IV, V and VI and respectively consisting of 8, 9, 10, 11, 12 and 13 targets was considered during the BA procedure. Each of these fields was obtained by adding one target to the previous field, so that the first eight targets are common to all the fields, the first nine targets are common to the last five fields and so on. A schema of this experimental configuration is shown in Figure 4.

All the fields were about 1.2 m above floor level. The target positions were identified using an isostatic support mounted on a tripod which was moved across the workspace. A set of the same isostatic supports was also available on a carbon-fibre bar that was used to provide the BA algorithm built in the system with a requested measurement of length (i.e. to scale the system). A distance of 1750 mm between two isostatic supports on the carbon-fibre bar was measured on a coordinate-measuring machine (CMM). The carbon-fibre bar was then placed in the central region of

the workspace. The coordinates of the two targets 1750 mm apart were measured with iGPS and their 1750 mm distance was used to scale the system in all the target fields considered. In this way, the scaling procedure is not expected to contribute to the variability of the measurement results even when different target fields are used in the BA procedure. Figure 5 shows an end of the vector bar used in this set-up (the large sphere in the figure), while coupled with an isostatic support (the three small spheres) during the measurement of a target position on the carbon-fibre bar.

The BA algorithm was run on each of these six targets fields so that six different numerical descriptions of the same physical positions and orientations of the transmitters were obtained.

Six new targets locations were then identified using the isostatic supports on the carbon-fibre bar mentioned above. Using the output of the BA executions, the spatial coordinates of these new target locations were measured. The approximate position of the six targets relative to the transmitters is shown in the schema of Figure 6.

Each target measurement consisted in placing the vector bar in the corresponding isostatic support and holding it for about 30 s. This enabled the measurement system to collect and store about 1200 records of target coordinates in $\{Uref\}$ for each of the six targets. In this way, however, the number of records for each target is different, owing to the human impossibility of manually performing the measurement procedure with a degree of time control sufficient to prevent this situation occurring.

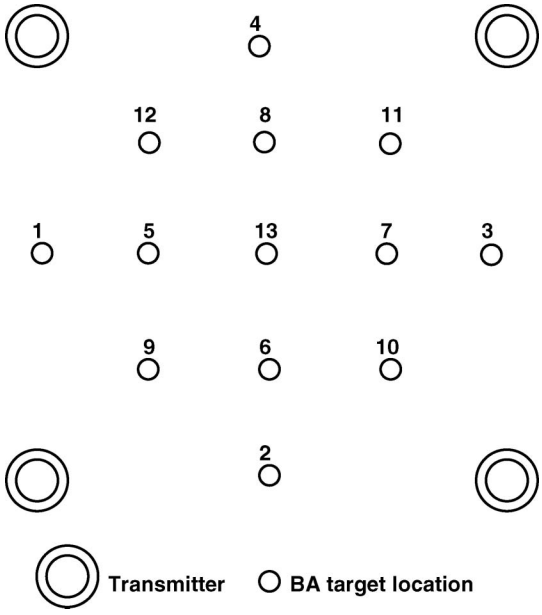

Figure 4. Target fields I, II, III, IV, V and VI.

5. Results

Each of the six target positions displayed in Figure 6 and labelled 1, 2, 3, 4, 5, 6 was measured using each of the six BA set-ups I, II, ..., VI, giving rise to a

Figure 5. Isostatic support identifying a target.

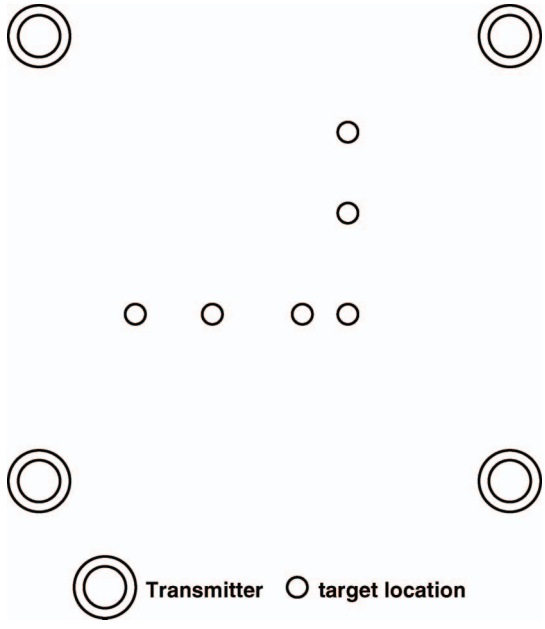

Figure 6. Target field when running the instrument.

grouping structure of 36 measurement conditions (cells).

When measuring a target location, its three Cartesian coordinates in $\{Uref\}$ are obtained. To reduce the complexity of the analysis from three-dimensional to mono-dimensional, instead of these coordinates the distance of the targets from the origin of $\{Uref\}$ is considered. Central to this investigation is the estimation of the effect on the target–origin distance due to the choice of a different number of target points when running the BA algorithm. The target locations 1, 2, ... , 6 do not identify points on a spherical surface, so they are at different distances from the origin of $\{Uref\}$, regardless of any possible choice of such a reference system. These target locations therefore contribute to the variability of the measurements of the target–origin distance whereby the detection of a potential contribution of the BA set-ups to the same variability can be hindered. To counteract this masking effect, the experiment was carried out by selecting first a target location and then randomly assigning all the BA set-ups for that location to the sequence of tests. This was repeated for all the six target positions. Such an experimental strategy introduces a constraint to a completely random assignment of the 36 measurement conditions to the the run order. In the literature (cf. Chapters 27, 16 and 8 in Neter *et al*. 1996, Faraway 2005 and Faraway 2006, respectively), this strategy is referred to as randomised complete block design (RCBD). The positions of the targets 1, 2, ... , 6 constitutes a blocking factor identifying an experimental unit or block, within which the BA set-ups are tested. The BA set-ups I, II, ... , VI constitute a random sample of all the possible set-ups that differ only in the choice of the location and number of points selected when running the BA algorithm during the system set-up phase. On the other hand, the analysis of the obvious contribution to the variability of the origin–target distance when changing the location of the targets would not add any interesting information to this investigation. These considerations lead to describing the experimental data of the RCBD with a linear mixed-effects statistical model, which is first defined and then fitted to the experimental data.

5.1. *Mixed-effects models*

The distance d_{ij} of the i th ($i = 1, \ldots, 6$) target from the origin measured when using the jth ($j = \mathrm{I}, \ldots, \mathrm{VI}$) BA procedure is modelled as the sum of four contributions: a general mean μ, a fixed effect τ_i due to the selection of the i th target point, a random effect b_j due to the assignment of the jth BA set-up and a random error e_{ij} due to all those sources of variability inherent in any experimental investigation that is not possible or convenient to control. This is described by the equation

$$d_{ij} = \mu + \tau_i + b_j + e_{ij}. \tag{1}$$

In Equation (1) and hereafter, the Greek symbols are parameters to be estimated and the Latin symbols are random variables. In particular, the b_j's have zero mean and standard deviation σ_b; the e_{ij}'s have zero mean and standard deviation σ. The e_{ij}'s are assumed to be made of independent random variables normally distributed, i.e. $e_{ij} \sim N(0,\sigma^2)$. The same applies to the b_j's, namely $b_j \sim N(0,\sigma_b^2)$. The e_{ij}'s and the b_j's are also assumed to be independent of each other. Under these assumptions, the variance of d_{ij}, namely σ_d^2, is given by the equation

$$\sigma_d^2 = \sigma_b^2 + \sigma^2. \tag{2}$$

Using the terminology of the 'Guide to the expression of uncertainty in measurement' (cf. Definition 2.3 in JCGM 2008), σ_d is the standard uncertainty of the result of the measurement of the origin–target distance.

As pointed out in the previous section, the number of the determinations of the target–origin distance that have been recorded is different for each of the 36 measurement conditions. For simplicity of the analysis, the number of samples gathered in each of these conditions has been made equal by neglecting the samples in excess of the original minimum sample size over all the cells. This resulted in considering 970 observations in each cell. The measurement result provided by the instrument in each of these conditions and used as a realisation of the response variable d_{ij} in Equation (1) is then defined as the sample mean of these 970 observations. There is a single measurement result in each of the 36 cells. The parameters of the model, i.e. μ, τ_i, σ_b and σ, have been estimated by the restricted maximum likelihood (REML) method as implemented in the `lme()` function of the package nlme of the free software environment for statistical computing and graphics called `R` (cf. R Development Core Team 2009). More details about the REML method and the package `nlme` are presented in Pinheiro and Bates (2000). The RCBD assumes that there is no interaction between the block factor (target locations) and the treatment (BA set-up). This hypothesis is necessary so that the variability within a cell represented by the variance σ^2 of the random errors can be estimated when only one experimental result is present in one cell. In principle, such an estimation is enabled by considering the variation of the deviations of the data from their predicted values across all the cells. This would estimate the variability

of an interaction effect, if it were present. If an interaction between target locations and BA set-ups actually exists, the estimate $\hat{\sigma}$ of σ provided in this study would account for both interaction and error variability in a joint way and it would not be possible to separate the two components. Therefore, from a practical point of view, the more the hypothesis of no interaction is violated, the more $\hat{\sigma}$ overestimates σ.

After fitting the model, an assessment of the assumptions on the errors has been performed on the realised residuals, i.e. the deviation of the experimental results from the results predicted by the fitted model for corresponding cells ($\hat{e}_{ij} = d_{ij} - \hat{d}_{ij}$). The realised residuals plotted against the positions of the targets do not appear consistent with the hypothesis of constant variance of the errors. In fact, as shown in Figure 7(a), the variability of the realised residuals standardised by $\hat{\sigma}$, namely $\hat{e}_{ij} = (d_{ij} - \hat{d}_{ij})/\hat{\sigma}$ seems different in different target locations.

For this reason, an alternative model of the data has been considered which accounts for the variance structure of the errors. This alternative model is defined as the initial model (see Equation 1), bar the variance of the errors which is modelled as different in different target locations, namely:

$$\sigma_i = \sigma_{\text{new}} \times \delta_i, \quad \delta_1 = 1. \tag{3}$$

From Equation (3) it follows that σ_{new} is the unknown parameter describing the error standard deviation in the target position 1, whereas the δ_i's ($i = 2, \ldots, 6$) are the ratios of the error standard deviation in the ith target position and the first.

The alternative model has been fitted using one of the class variance functions provided in the package `nlme` and the function `lme()` so that also σ_{new} and the δ_i's are optimised jointly with the other model parameters (μ, τ_i and σ_b) by the application of the REML method (Section 5.2 in Pinheiro and Bates 2000).

For the alternative model, diagnostic analyses of the realised residuals were not in denial of its underlying assumptions. The standardised realisations of the residuals, i.e. $\hat{e}_{ij} = (d_{ij} - \hat{d}_{ij})/\hat{\sigma}_i$, when plotted against the target locations (Figure 7(b)) do not appear any longer to exhibit different variances in different target locations, as was the case in the initial model (Figure 7(a)). The same standardised realisations were also found not to exhibit any significant departure from normality.

The fact that all the target fields have more than 50% of the targets in common together with the fact that each field has been obtained by recursively adding a single target to the current field may cause the experimenters to expect that the measurement results obtained when different target fields have been used in

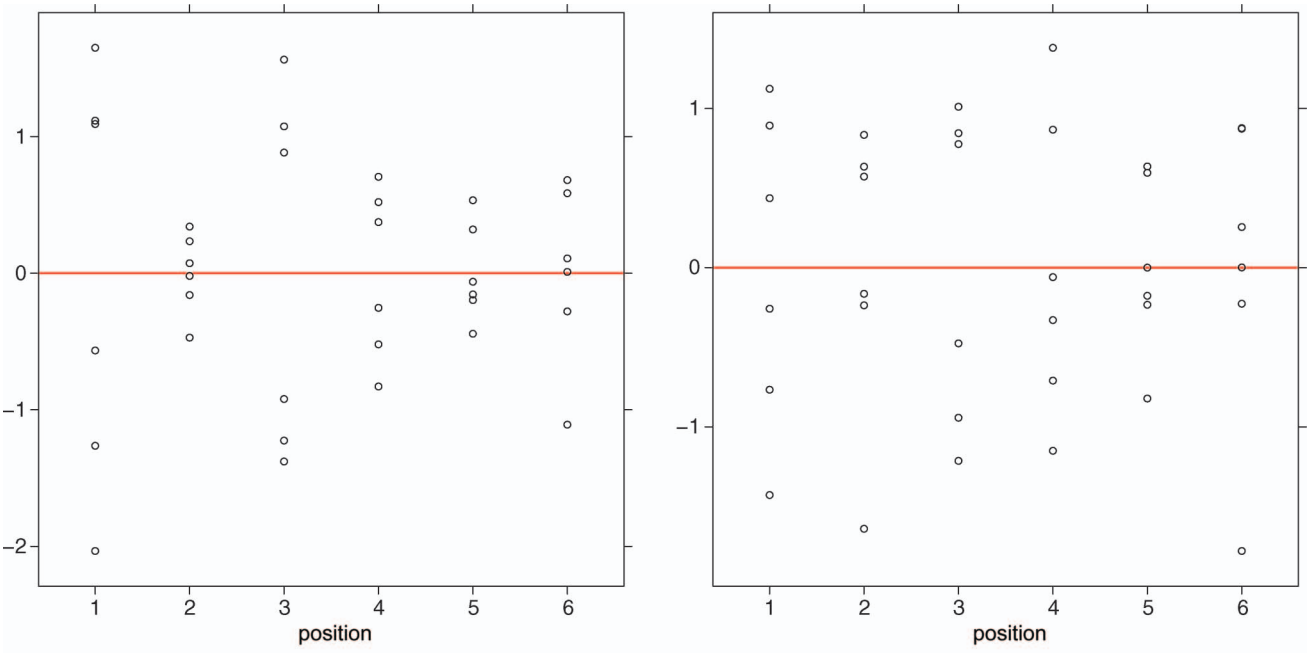

(a) Initial model. (b) Alternative model.

Figure 7. Realisations of the standardised residuals (dimensionless) grouped by target positions for the initial and the alternative mixed-effects models.

the BA procedure have some degree of correlation. If that were the case, then the experimental results should be in denial of the assumed independence of the random effects b_j's. The random effects, like the errors, are unobservable random variables. Yet, algorithms have been developed to predict the realisations of these unobservable random effects on the basis of the experimental results and their assumed model (Equations 1, 2 and 3 with the pertinent description above). The predictor used in this investigation is referred to as the best linear unbiased predictor (BLUP). It has been implemented in `nlme` and it is described, for instance, in Pinheiro and Bates (2000). The predicted random effects $\hat{b_j}$'s for the model and the measurement results under investigation are displayed in Figure 8(a). To highlight a potential correlation between predicted random effects relative to target fields that differ by only one target, the $\hat{b_{j+1}}$'s have been plotted against the $\hat{b_j}$ in Figure 8(b) ($j = 1, \ldots, 5$). From a graphical examination of the diagrams of Figure 8 it can be concluded that, in contrast with what the procedure for establishing the targets fields may lead the experimenter to expect, the measurement results do not appear to support a violation of the hypothesis of independence of the random effects. Similar values for the BLUPs and therefore similar conclusions can be drawn also for the initial mixed-effect model (the BLUPs for the initial model have not been reported for brevity).

As suggested in Pinheiro and Bates (2000) (Section 5.2, in particular), to support the selection between the initial and the alternative model, a likelihood ratio test (LRT) has been run using the generic function `anova()` implemented in `R`. A p-value of 0.84% led to the rejection of the simpler initial model (8 parameters to be estimated) when compared with the more complex alternative model (8 + 5 parameters to be estimated). The same conclusion would hold if the selection decision is made on the basis of the Akaike information criterion (AIC) also provided in the output of `anova()` (read more about AIC in Chapters 1 and 2 of Pinheiro and Bates 2000).

This model selection bears significant practical implications. From a practitioner's point of view, in fact, selection of the alternative model means that the random errors have significantly different variances when measuring targets in different locations of the workspace. The workspace is not homogeneous: there are regions where the variability of the random errors is significantly lower than in others. This also means that a measurement task can therefore be potentially designed so that this measuring system can perform it satisfactorily in some regions of its workspace but not in others.

REML estimates of the parameters that have practical implications are as follows:

$$\hat{\sigma}_b = 160.7\,\mu\text{m}, \tag{4}$$

$$\hat{\sigma} = 14.28\,\mu\text{m}, \quad \hat{\delta}_2 = 0.2625, \quad \hat{\delta}_3 = 0.8599,$$
$$\hat{\delta}_4 = 0.3706, \quad \hat{\delta}_5 = 0.1260, \quad \hat{\delta}_6 = 0.5446. \tag{5}$$

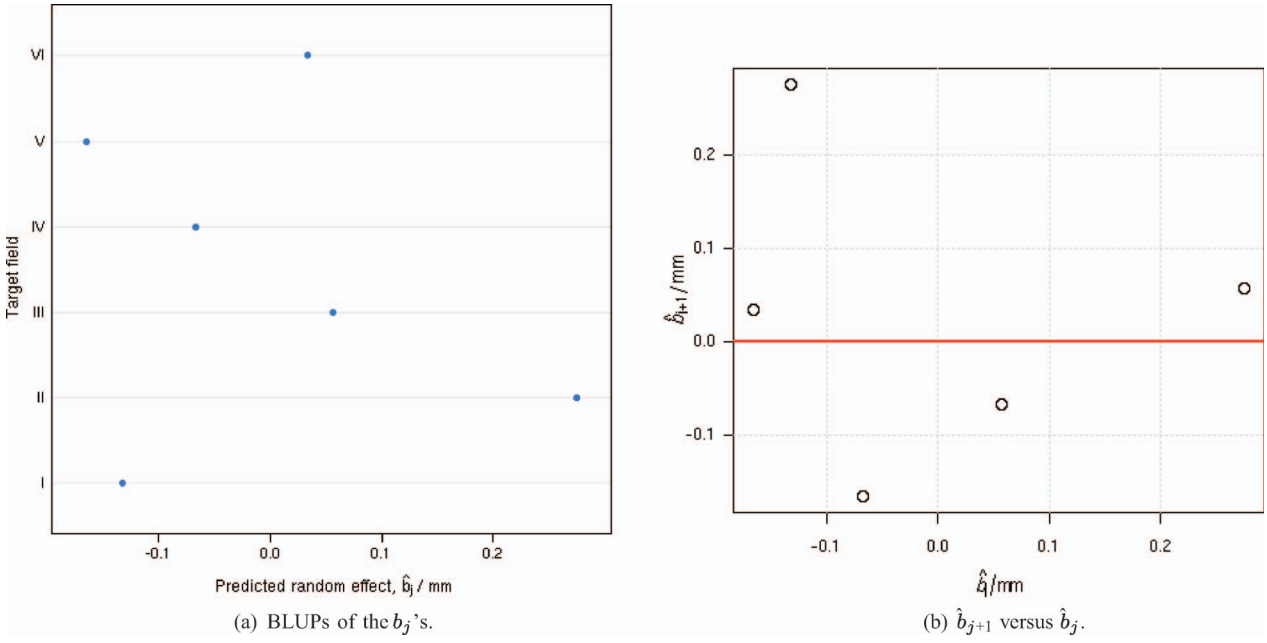

(a) BLUPs of the b_j's. (b) $\hat{b}_{j+1}$ versus $\hat{b}_j$.

Figure 8. BLUPs of the random effects for each targets field and the graphically insignificant autocorrelation between BLUPs of random effects associated with consecutive targets fields.

Estimates $\hat{\tau}_i$ confirm the tautological significance of the location of the targets or block factor, whereas $\hat{\mu}$, depending on the the parametrisation of the model, can for instance be the centre of mass of the point locations or can also be associated with a particular target location (cf. Chapters 13 and 14 in Faraway 2005). All these estimates do not convey any practical information. They are therefore not reported.

The significance of the random effect associated with the BA set-up procedure has been tested using a likelihood ratio approach, where the alternative model has been compared with a null model characterised by an identical variance structure of the errors but without any random effect (i.e. $\sigma_b = 0$). The p-value was less than 10^{-32} under the assumption of a chi-squared distributed likelihood ratio. In reality, as explained in Section 8.2 of Faraway (2006), such an approach is quite conservative, i.e. it tends not to reject the null hypothesis by overestimating the p-value. However, given the extremely low p-value ($<10^{-32}$), there is strong evidence supporting the rejection of the null hypothesis of an insignificant random effect (H_0: $\sigma_b = 0$).

From a practical point of view, this indicates that caution should be exerted when selecting the target locations for running the BA algorithm during the set-up phase: when repeating the BA procedure during the set-up with identical positions of the transmitters, the consideration of a different number of targets significantly inflates the variability of the final measurement results.

Substituting the estimates of Equations (4) and (5) in the adaptation of Equation (2) to the alternative model, it is derived after a few passages that the choice of a different number of targets when running the BA algorithm during the set-up phase accounts for 99.22, 99.94, 99.42, 99.89, 99.99 and 99.77% of the variance of the measured origin–target distance when the target is in locations 1, 2, 3, 4, 5 and 6, respectively. If there was no discretion left to the operator when selecting the number of targets and their locations during the BA procedure, then the overall variability of the final results in each of the location tested could have been reduced by the large percentages reported above.

It may be worth pointing out that the designed experiment considered in this investigation could be replicated K times, on the same or in different days. The obtained measuring results could then be modelled with the following equation:

$$d_{ijk} = \mu + \tau_i + b_j + c_k + e_{ijk} \qquad (6)$$

with $c_k \sim N\left(0, \sigma_c^2\right)$, $k = 1, 2, \ldots, K$, being the random effect associated with the kth repetition of the experiment. The significance of the random effects c_k's could then be tested in a similar way as the significance of the b_j's has been tested above. The practical use of the model of Equation (6) is twofold. First, it enables the experimenter to detect if a significant source of varibility can be associated with the replication of the whole experiment. For instance, if each replication takes place in slightly different natural and/or artificial light conditions, then testing the significance of the c_k would tell if these enviromental conditions had significant effects on the measurement results (d_{ijk}). The estimate $\hat{\sigma}_c$ would quantify the increased variability of the response varible attributable to them. Second, the increased number of measurements taken would raise the confidence of the experimenter in the estimates of $\hat{\sigma}_b$, $\hat{\sigma}_c$ and $\hat{\sigma}$. For instance, it would dissipate (or confirm) the suspicion that the experimenter may have that the random effects attributed in Equation (1) to the different setups, namely the b_j's, may be contributed to by the natural variability due to repetition which was estimated in Equations (4) and (5). This further study can be considered as future work.

5.2. *Transient definition and analysis*

In the above analysis, the average of all the 970 experimental data in a cell has been considered. The variability of each of these 970 determinations of distance, say σ_t, is significantly larger than that of their average (σ_d). If these determinations were mutually independent, then it would be $\sigma_d = \sigma_t/\sqrt{(970)}$. But the determinations are instead highly correlated, owing to the fact that they are taken at varying sampling intervals of the order of milliseconds. Identifying the correlation structure of these determinations is beyond the scope of this investigation. In this study, when the instrument is measuring the tth determination, say $d_{t,ij}$, a running average of all the determinations measured until that instant, say $\bar{d}_{t,ij}$, is considered. An interesting question that arises is: 'How many determinations are sufficient for the instrument to provide a measurement $\bar{d}_{t,ij}$ that does not differ much from the measurement result d_{ij}?'. A 2μm maximum deviation from d_{ij} has been considered for differentiating the steady and the transient states of $\bar{d}_{t,ij}$. The value $t^\star$ has been used to identify the end of the transient. In other words, for any index $t > t^\star$ it holds that $|d_{t,ij} - d_{ij}| < 1\mu$m.

In Figure 9, for each of the 36 experimental conditions, two continuous horizontal lines 1 μm apart from the measurement result d_{ij} delimit the steady-state region, whereas a single vertical dashed line indicates the transition index $t^\star$ from the transient to the steady state as defined above.

From Figure 9, it is observed that for the same target location (panels in the same column) the

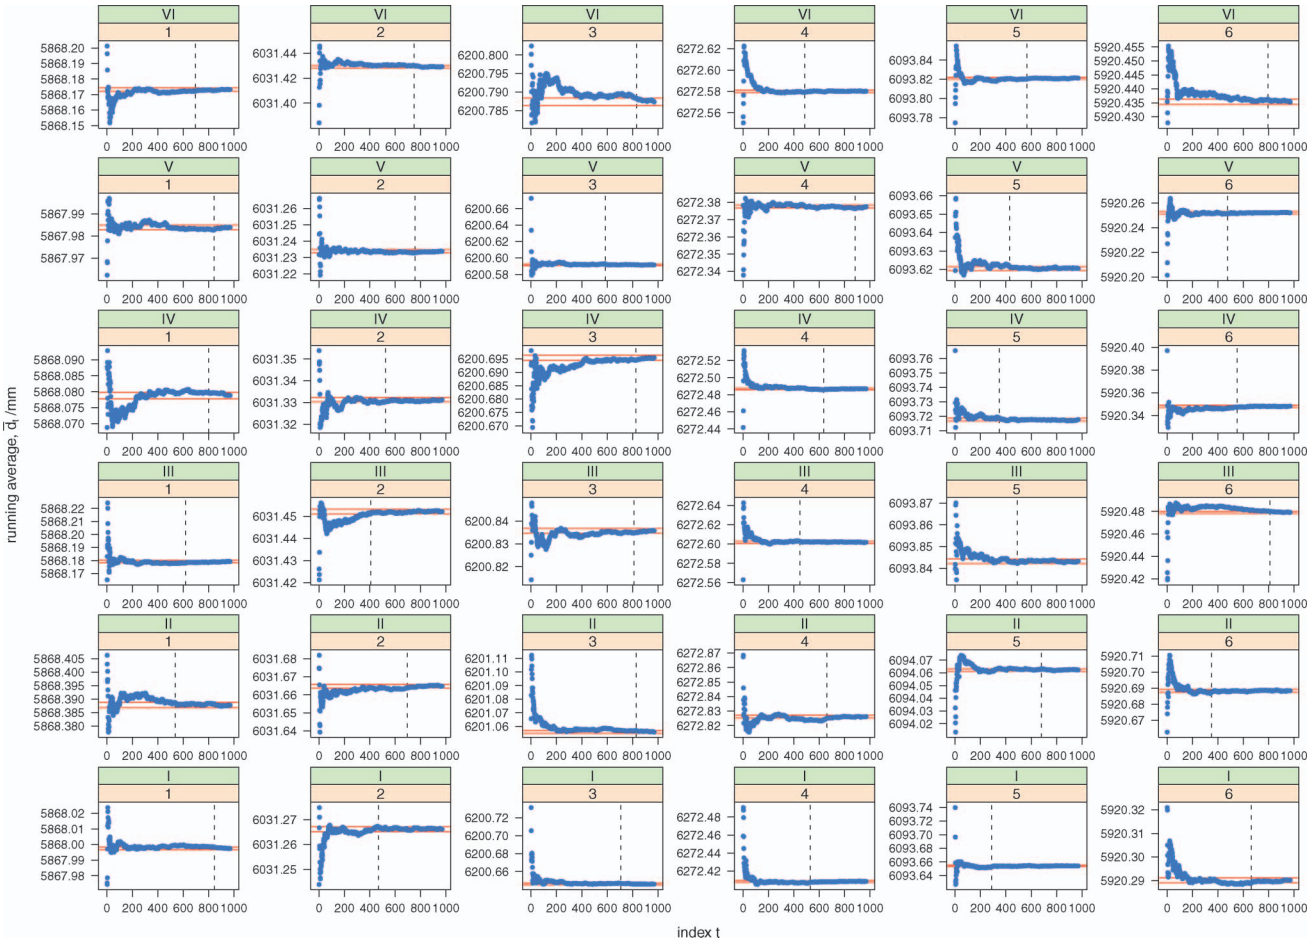

Figure 9. Transition from the transient and the steady state.

transition from transient to steady state may occur at different $t^{\star}$'s for different BA set-ups (different vertical dashed lines in each panel). This suspicion is even stronger when considering $t^{\star}$ for the same BA set-up but for different target locations (panels on a row in Figure 9).

To ascertain whether the variation of $t^{\star}$ with the BA set-ups and with the target locations examined is significant or is only the result of uncontrolled or uncontrollable random causes, the experimental values of $t^{\star}$ calculated starting from the RCBD already discussed have been analysed with a fixed-effects ANOVA model (cf. Section 16.1 in Faraway 2005). The values of $t^{\star}$ have been computed by an ad hoc function implemented in R by one the authors. The $t^{\star}$'s are assumed as though they have been generated by the following equation:

$$t^{\star}_{ij} = \mu + \beta_i + \gamma_j + e_{ij}, \qquad (7)$$

where the β_i's and the γ_j's are the effects of the blocking factor (the target locations) and of the BA set-ups, respectively, whereas the e_{ij}'s are the random errors, assumed independent, normally distributed with constant variance and zero mean. The parameters have been estimated using the ordinary least squares method as implemented in the function `lm()` in `R` (cf. R Development Core Team 2009). The assumptions underlying the models have been checked on the realised residuals and nothing amiss was found. To test the potential presence of interaction between the two factors in the form of the product of their two effects, a Tukey test for additivity was also performed (cf. Section 27.4 in Neter *et al.* 1996). This test returned a p-value of 30.43%. It is therefore concluded that the experimental data do not support the rejection of the hypothesis of an additive model in favour of this particular type of interaction effect of target locations and BA set-ups on $t_{ij}^{\star}$.

The effect of the target positions on $t^{\star}$ was significant, i.e. $H_0 : \beta_i = 0 (i = 1, 2, \ldots, 6)$ gives rise to a p-value of 3.88% (under the hypotheses of the model). However, the effect of the BA set-ups did not appear to be significant, i.e. $H_0 : \gamma_i = 0 (i = 1, 2, \ldots, 6)$

gives rise to a p-value of 84.96% (under the hypotheses of the model).

From a practical point of view, there are two main implications of these findings. First, the selection of a different number of targets when running BA algorithms during the set-up phase does not appear to have significant consequences on the duration of the transient for obtaining a measurement. Second, the duration of the transient appears to be significantly different for different target locations within the workspace. This may be stated as follows: there are regions of the workspace that require longer transient periods than others before a measurement result stabilises, and this is expected to have consequences for the accuracy and precision of the determination of the position of moving objects (tracking).

In fact, if a target point is in motion at a speed sufficient for a number of determinations greater than $t^{\star}$ to be recorded in each measured point of its trajectory, then all the measurements results will be representative of a steady state. But this may hold for only some portions of the target trajectory; for others, characterised by a larger $t^{\star}$, such a condition may not be satisfied with a consequent inflation of the variability of those estimated positions, which may also be biased.

6. Conclusions

The main characteristics of the Metris Indoor GPS system have been reviewed on the basis of information in the public domain. In particular, the working principles of the system have been presented in terms of a cone-based mathematical model.

The overall description of the system has been instrumental in highlighting the key role of bundle adjustment procedures during the set-up of the system. The selection of the number and location of target points that are used when running the bundle adjustment procedure during the set-up phase can be affected by discretionary judgements exerted by the operators.

To investigate the statistical significance of the effects of this selection, a randomised complete block design has been run on the distance between the origin of the reference system and the measured positions of target locations different from those used during the bundle adjustment in the set-up phase. This design enhances the possibility for the potential effects of different set-ups on the origin–target distances to be detected by discriminating them from the obvious effects of the target positions. The set-ups considered were different only in the number of the targets used when executing the bundle adjustment procedure.

A mixed-effects and a fixed-effects linear statistical model were fitted to the measurements results using the restricted maximum likelihood method and the ordinary least squares technique, respectively.

The measurement results defined as the sample average of the 970 determinations of distance recorded in each target location for each set-up have been analysed with the mixed-effects model. By analysing the realisations of the residuals, statistically different standard deviations of the random errors were identified for different target positions. The work envelope of the instrument do not therefore appear homogeneous: in some areas the variability of the random error is greater than in others, when performing measurements of the distance of a target from the origin. Owing to this heterogeneity, the punctual estimates of the standard uncertainty of the measured distances (σ_d) were different for different target position and lay in a range between 160.8 and 161.4 μm. The different set-ups, tested to be statistically significant, always accounted for more than 99.2% of the estimated standard uncertainty (the percentage varies for different target positions). This quantitative evidence suggests that the selection of points when running the bundle adjustment algorithms in the set-up phase should not be overlooked. Performing this selection in a consistent way according to some rule that ideally leads to chosing the same points when the transmitters are in the same positions may be a course of action worth considering. Also, for replication and comparison purposes, it may be advisable to quote the locations of the targets used in setting up the system when reporting the results of a measurement task.

The duration of the transient, i.e. the number of determinations of distance needed for their current average to be within ± 1 μm from the measurement result (the average of the 970 determinations), has been analysed with the fixed-effects model. The different set-up configurations considered did not have any significant effect on the duration of the transient. However, this duration was significantly different in different target locations. It can therefore be concluded that the working space of the instrument is heterogeneous also for the characteristics of the transient of measurement. It is expected that this conclusion has negative implications on the precision and unbiasedness of the measurements obtained when using the instrument for tracking moving points or moving objects that the target points (or vector bars) are attached to. Given a pre-specified configuration of iGPS transmitters without any zone partitions having been pre-established among them, if an object is moving within an area of the working space of such an iGPS, say area A, its position may be tracked correctly, because the transient is sufficiently short there. But if the same movement of the same object is tracked by the same iGPS in another area of the same

iGPS working area, say area B, the system may not be able to track its location correctly because the transient may not yet have finished.

Acknowledgements

This study is part of the research initiatives of The Bath Innovative Design and Manufacturing Research Centre (IdMRC), which is based in the Department of Mechanical Engineering of the University of Bath and which is supported by the United Kingdom Engineering and Physical Sciences Research Council (EPSRC). In particular, this investigation was carried out within the scope of the IdMRC research theme 'Metrology, Assembly Systems and Technologies' (MAST), which is coordinated by Professor Paul Maropoulos.

References

BS ISO 5725-1, 1994. *Accuracy (trueness and precision) of measurement methods and results – Part 1: General principles and definitions* [online]. London: British Standards Institution. Available from: http://shop.bsigroup.com/ [Accessed 14 March 2010].

Craig, J.J., 1986. *Introduction to robotics mechanics & control*. Reading, MA, USA: Addison-Wesley.

Estler, W.T., *et al*., 2002. Large-scale metrology – an update. *Annals of the CIRP*, 51 (2), 587–609.

Faraway, J., 2005. *Linear models with R*. Boca Raton, FL, USA: Chapman & Hall/CRC.

Faraway, J., 2006. *Extending the linear model with R: generalized linear, mixed effects and nonparametric regression models*. Boca Raton, FL, USA: Chapman & Hall/CRC.

Hartley, R.I. and Sturm, P., 1997. Triangulation. *Computer Vision and Image Understanding*, 68 (2), 146–157.

Hedges, T.M., *et al*., 2003. Position measurement system and method using cone math calibration. United States Patent US6535282B2, March.

JCGM (Joint Committee for Guides in Metrology), 2008. *JCGM:100:2008 (GUM:1995 with minor corrections) Evaluation of measurement data – guide to the expression of uncertainty in measurement* [online]. Sèvres, France: BIPM – Bureau International des Poids et Mesures. Available from: http://www.bipm.org [Accessed 7 June 2009].

Liu, H., *et al*., 2007. Survey of wireless indoor positioning techniques and systems. *IEEE Transactions on Systems, Man, and Cybernetics – Part C: Applications and Reviews*, 37 (6), 1067–1080.

Liu, Z., Liu, Z., and Lu, B., 2008. Error compensation of indoor GPS measurement. *In*: C. Xiong *et al*., eds., *Proceedings of the first international conference on intelligent robotics and applications (ICIRA2008), Part II*, 15–23 October 2008. Wuhan, People's Republic of China. Lecture Notes in Computer Science 5315. Berlin: Springer-Verlag, 612–619.

Lourakis, M.I.A. and Argyros, A.A., 2009. SBA: a software package for generic sparse bundle adjustment. *ACM Transactions on Mathematical Software*, 36 (1), 1–30.

Maisano, D., *et al*., 2008. Indoor GPS: system functionality and initial performance evaluation. *International Journal of Manufacturing Research*, 3 (3), 335–349.

Neter, J., *et al*., 1996. *Applied linear statistical models*. 4th ed. New York: Irwin.

Pinheiro, J.C. and Bates, D.M., 2000. *Mixed-effects models in S and S-plus*. New York: Springer-Verlag.

Piontek, H., Seyffer, M., and Kaiser, J., 2007. Improving the accuracy of ultrasound-based localisation systems. *Personal and Ubiquitous Computing*, 11 (6), 439–449.

Priyantha, N.B., Chakraborty, A., and Balakrishnan, H., 2000. The Cricket location-support system. *In*: *Proceedings of the 6th annual international conference on mobile computing and networking (MobiCom'00)*, 6–11 August. Boston, MA. New York: ACM, 32–43.

R Development Core Team, 2009. *R: a language and environment for statistical computing* [online]. Vienna: R Foundation for Statistical Computing. Available from: http://www.R-project.org. ISBN 3-900051-07-0.

Sae-Hau, C., 2003. Multi-vehicle rover testbed using a new indoor positioning sensor. Thesis (Master's). Massachusetts Institute of Technology.

Savvides, A., Han, C.C., and Strivastava, M.B., 2001. Dynamic fine-grained localization in ad-hoc networks of sensors. *In*: *Proceedings of the 7th annual international conference on mobile computing and networking (MobiCom'01)*, 16–21 July. Rome. New York: ACM, 166–179.

Triggs, B., *et al*., 2000. Bundle adjustment – a modern synthesis. *In*: B. Triggs, A. Zisserman, and R. Szeliski, eds. *Proceedings of the international workshop on vision algorithms (in conjunction with ICCV'99)*, 21–22 September 1999, Kerkyra, Corfu, Greece. Lecture Notes on Computer Science 1883. Berlin: Springer-Verlag, 298–372.

Welch, G., *et al*., 2001. High-performance wide-area optical tracking – the HiBall tracking system. *Presence: Teleoperators and Virtual Environments*, 10 (1), 1–21.

Wolf, P. and Ghilani, C., 1997. *Adjustment computations: statistics and least squares in surveying and GIS*. New York: Wiley.