

FIRST APPLICATION OF THE LGC2 SOFTWARE TO SuperKEKB MAIN-RING SURVEY DATA: A SECTION-TO-SECTION CONSISTENCY CHECK

T. K. Tran^{*1}, M. Masuzawa¹

¹High Energy Accelerator Research Organization (KEK), Tsukuba, Japan

Abstract

Alignment of SuperKEKB, whose nano-beam optics demand sub-millimetre magnet positioning, relies on laser-tracker surveys that have proven highly accurate: the surveyed ring circumference was confirmed directly by the beam during commissioning. The survey is, however, performed intermittently, while the tunnel and the interaction region move on daily, seasonal and long-term scales, and the continuous monitoring (hydrostatic levelling, gap sensors and a GPS network) is handled by separate systems rather than in a single adjustment. This motivates evaluating a rigorous, unified three-dimensional framework. We report the first application of LGC2 (Logiciel Général de Compensation), the least-squares network-adjustment software developed and maintained at CERN, to SuperKEKB main-ring survey data. For the present analysis the ring is divided into nine sections following the existing naming convention: the Tsukuba, Oho, Fuji and Nikko regions, each split into left and right halves, plus the interaction region; these boundaries are introduced solely for the analysis and do not correspond to any change in the survey. The full data set is brought into a common machine frame with LGC2. The reconstruction is validated through comparisons of reference monuments measured across each section boundary. The good agreement obtained between adjacent sections confirms the consistency of the reconstructed network and demonstrates that LGC2 is a suitable platform for integrating periodic survey with continuous HLS and future frequency-scanning-interferometry (FSI) monitoring at SuperKEKB.

INTRODUCTION

SuperKEKB is a double-ring electron-positron collider of 3016 m circumference, built in the tunnel that was excavated for TRISTAN and later used by KEKB. Its nano-beam scheme squeezes the vertical beam size at the interaction point (IP) to a few tens of nanometres, which makes the machine considerably more sensitive to magnet misalignment than its predecessor [1, 2]. The main-ring magnets were surveyed with laser trackers in the summer of 2015, after the 2011 earthquake had destroyed the KEKB survey network, and were then aligned to a smoothed reference curve [1]. That survey was successful: the circumference predicted from it agreed with the beam to within 1 mm (0.3 ppm), and the two rings were found to differ in circumference by only 0.2 mm [1, 3].

The tunnel, however, keeps moving. It has subsided by more than 40 mm with respect to the IP since 1998, its floor level follows the seasonal temperature swing by more than 1 mm, and the interaction region shows daily and long-term drifts of its own [2, 4]. These motions are followed by several independent systems: hydrostatic levelling (HLS) near the IP, capacitive gap sensors on the final-focus cryostats, a surface GPS network, and the periodic laser-tracker survey. They are not combined in a single adjustment. The network adjustment of the 2015 survey was carried out with the geodetic packages PANDA and PAG-U [1].

LGC (Logiciel Général de Compensation) is the least-squares network adjustment software developed at CERN since the 1990s; LGC2 is its current C++ implementation, available as open source with documentation [5]. It computes rigorously in three dimensions, accepts laser-tracker, total-station, levelling, inclination and wire observations in one adjustment, and returns the full statistics of the solution: the a posteriori variance factor with its χ^2 test, the residual of every observation with its normalised value, reliability numbers, and the standard deviation of every coordinate. At the LHC it already integrates hydrostatic and wire-positioning sensors with the survey network. That is precisely the kind of unified framework SuperKEKB lacks today, and it is the reason we decided to test LGC2 on the existing 2015 data before considering it for future campaigns.

This paper addresses a practical question: what the raw laser-tracker data can tell us about the survey on their own, and what LGC2 adds. We first describe the data and the precision information they carry. We then adjust each of the nine sections independently and read the quality of the survey out of the adjustment statistics. Finally, the nine sections are joined in one adjustment of the whole ring, built on the tunnel monuments only, and the same statistics are used to locate the parts of the survey that need attention.

THE 2015 SURVEY DATA

How the Ring Was Measured

The survey strategy is described in Ref. [1]. Faro laser trackers (Ion and Vantage) were used by three teams working in parallel. The network was built as successive unit rectangles, about 8 m long and 2 m wide, and one tracker set-up covered roughly four such rectangles, i.e. about 30 m of tunnel. Adjacent set-ups share targets so that the results can be cross-checked. More than 1200 network control points on the tunnel wall and floor and more than 4000 magnet reference points, over 5000 points in total, were measured

^{*} tuyet@post.kek.jp

in more than 400 set-ups. Long distances were compared against a total station. Each magnet carries reference holders on which the tracker target is placed; magnet pitch and roll were measured separately with digital inclinometers.

The monuments are of three kinds. Two rows, named IT and IB (inner top, inner bottom), are drilled into the inner tunnel wall at two heights, about 1.2 m apart (Fig. 1, left). A third row, named OV or OH (outer vertical, outer horizontal), is on the opposite wall of the tunnel, across the beam line from the IT and IB rows, where the monuments are mounted on steel posts among the pipe and cable-tray supports (Fig. 1, centre and right). This distinction turns out to matter for the result.

The data set we received consists of 387 report files, one file per tracker set-up, organised in nine folders that follow the KEK naming convention: Tsukuba left and right (TL, TR), Oho (OL, OR), Fuji (FL, FR), Nikko (NL, NR), and the interaction region (IP). Each file is named after the nearest IT monument, although the instrument does not stand on that monument; it stands on a tripod in the middle of the tunnel, about 1.7 m away and 0.3 m higher. Altogether the files contain 18 190 sightings to 4786 distinct points: 3614 magnet targets, 1091 wall monuments of the IT, IB and OV/OH rows, and about 80 floor and hall monuments (Fig. 3). A median set-up sees 40 points, the median sight distance is 11 m and the longest 70 m. From the report dates in the file headers the campaign ran from 27 July to 9 September 2015, i.e. 44 days. Table 1 lists the nine sections.

Figure 2 shows what the survey looks like on the ring. Each square is one place where the tracker was set up; the dots around it are the points it measured, magnet targets and monuments alike, so the chain of squares and dots is the survey itself, walked once around the ring. Where two sections meet, the same points were measured from set-ups belonging to both, and the red circles mark these nine joints with the number of points measured from both sides. The interaction region is enlarged in the inset: it is only about 40 m long, but its three set-ups see 47 points in common with Tsukuba L and 48 with Tsukuba R, more than any other joint, so it is what connects the two halves of the ring at 12 o'clock. Every design target but one was measured.

Table 1: The Nine Sections of the 2015 Data Set

Section	Set-ups	Targets	Arc [deg]
TL (Tsukuba left)	51	443	46.8
IP	3	56	4.7
TR (Tsukuba right)	51	445	47.3
OL (Oho left)	44	518	48.1
OR (Oho right)	48	537	46.6
FL (Fuji left)	47	384	47.0
FR (Fuji right)	45	381	46.1
NL (Nikko left)	50	520	46.1
NR (Nikko right)	48	522	50.0
Total	387	3806	378

What One File Contains

Every report file gives, for each sighted point, Cartesian coordinates X, Y, Z in millimetres, a per-point precision figure, and the number of shots. The essential point for everything that follows is that these coordinates are expressed in the local frame of that set-up: the origin is the optical centre of the instrument and the azimuth of the axes is arbitrary. The same wall monument therefore appears in nine different files with X values ranging over more than three metres. The raw data are 387 small maps that cannot be compared with each other directly.

Two properties of the files turned out to be decisive. First, all 387 files store a @LEVEL plane: the tracker was levelled with its inclinometer at every set-up, and the local Z axis is the true vertical. Second, the polar observations of the instrument (horizontal angle, zenith angle, slope distance) are not in the export; they have to be recovered from the Cartesian triplet, which is an exact transformation but loses any correlation the tracker software introduced between the three coordinates. The nominal columns of the reports are empty. The design coordinates of the 3624 magnet targets come from a separate table in the machine frame; it contains no wall monuments, so their approximate coordinates had to be propagated from the magnet targets before the adjustment.

What the Raw Data Alone Can Show

The only quality information in the raw files is the per-point precision column (Fig. 4). Over all 18 769 rows its median is 0.039 mm and its 95th percentile 0.129 mm; one point exceeds 0.4 mm. By point type, the magnet targets are the best (median 0.037 mm), the IT and IB monuments are close behind (0.045 mm to 0.047 mm), and the OV/OH monuments are the widest (median 0.052 mm, upper quartile above 0.08 mm). Its exact definition is not documented in the file; it is the internal repeatability of the shots to one target from one set-up, as computed by the export software.

Two things follow. The instrument itself was performing well: a repeatability of a few hundredths of a millimetre is what a modern tracker delivers at 10–30 m. And the number is a within-set-up quantity. It says nothing about whether two set-ups agree on the position of the monument they both saw, nothing about the orientation of one set-up relative to the next, and nothing about any error that is smooth over more than one set-up, which is exactly the kind of error that matters for the closed orbit. Those questions can only be answered once all set-ups are expressed in one frame with redundancy, and that is what the adjustment is for.

ADJUSTMENT OF THE NINE SECTIONS WITH LGC2

The Input File

A script converts each section into one LGC2 input file. Its structure is the standard one of LGC2 [5]:

- *TITR, *OLOC, *PREC: title, local Cartesian frame, output precision;




Figure 1: Left: a survey monument drilled into the tunnel wall, with its identification label. Centre: the tunnel seen from the magnet side; the red ellipse marks an OV monument on the opposite wall, across the beam line; the arrow points to its close-up view. Right: the same OV monument, mounted on a steel post among the pipe and cable-tray supports.

Laser-tracker setups and the 9 section-to-section joints, whole-ring plan view




Figure 2: Plan view of the 2015 survey. Grey band: the 3624 design targets. Small dots: the measured points, coloured by section. Squares: the 387 tracker set-ups. Red circles: the nine joints between neighbouring sections, with the number of points measured from both sides. The inset enlarges the interaction region, where three set-ups (IPL08, IPCENTER, IPR08) connect Tsukuba L and Tsukuba R. The star marks the one design target that was not measured. Set-up positions are approximate and serve only to show where the instrument stood.

- ***INSTR / *POLAR:** the instrument error model, one group of a priori standard deviations for horizontal angle, zenith angle and distance (constant plus proportional part). The instrument model is not recorded in the export files, so the group name is a placeholder; the sigmas themselves are estimated from the data (see below);
- ***CONSI LIBR:** a free network with the datum defined by inner constraints on the approximate coordinates, six parameters (three translations, three rotations) in the first runs and four (three translations and the rotation about the vertical) once the stations were declared levelled;




Figure 3: Content of the raw data set. Left: number and type of measured points in each section, with the number of set-up files. Right: magnet targets per section, split between the electron ring (HER) and the positron ring (LER).




Figure 4: The per-point precision column of the raw files. Left: distribution over all 18 769 rows. Right: the same value by point type.

- ***FAUT, *PUNC, *JSON**: outlier detection with the data-snooping test, error ellipsoids, and the interactive HTML report;
- one ***TSTN** block per set-up, followed by ***VO** (unknown orientation) and one ***PLR3D** line per sighting, i.e. horizontal angle, zenith angle and distance recomputed from the Cartesian export; feeding the Cartesian coordinates directly as ***OBSXYZ** was tried first and did not converge;
- ***POIN**: approximate coordinates of every point, taken from the design table for the magnet targets and propagated to the monuments.

The station keyword deserves a comment. With **R0T3D** each station is free to rotate about three axes; with **IHFIX** it is free to rotate only about the vertical, i.e. the inclinometer levelling stored in every file is used. Because the plumb lines converge by about $149 \mu\text{rad}$ across the 948 m diameter of the ring, the zenith angles were corrected for the curvature of the plumb-line field with a scale factor k per section, determined by running the adjustment at three values of k and taking the minimum of σ_0^2 . On section TL the switch from **R0T3D** to **IHFIX** with a four-parameter datum changed σ_0 from 0.980 to 0.992 after re-estimating the zenith sigma, and reduced the vertical standard deviation of the points from 0.074 mm

to 0.026 mm: the levelling information costs nothing and is worth having.

The a priori sigmas were refined by variance-component estimation: after each run the sigma of each observation type was rescaled by the ratio of its a posteriori to a priori variance, until σ_0 entered its critical interval. Every section needed two to four such passes.

One Section Read in Detail: TL

Figure 5 shows the adjusted section TL as LGC2 returns it: 51 stations in the middle of the tunnel, the magnet targets in a band along the beam line, the IT and IB rows on one side, the OV/OH row on the other, and the sight lines that tie each station to its neighbours through the monuments. Figures 6 and 7 are taken from the HTML report that LGC2 writes for every run. Since the rest of the paper is read from these plots, we first say what they show.

The *residual* of a sight line is the difference between the observed value (angle or distance) and the value recomputed from the adjusted coordinates. It is the disagreement between that sight line and all the other sight lines that saw the same points. A histogram of residuals (Fig. 6) answers two questions: is it centred on zero, i.e. is there a systematic error left in the observations of that type, and how wide is it, i.e. how well do the set-ups agree with each other. Its

standard deviation is the number quoted as the a posteriori sigma of that observation type.

The *normalised residual*, written RES/SIG in the report, is the residual divided by the standard deviation declared for that observation in the error model. A residual of 0.3 mm on a 2 m distance is a bad one; on a 20 m distance it is ordinary. Dividing by the sigma makes all sight lines comparable. If the error model describes the measurements correctly, the normalised residuals follow a standard normal distribution: centred on zero, standard deviation one, 99.7 % inside ± 3 . A value beyond three means that the sight line disagrees with the rest of the network by more than the error model allows, and LGC2 uses this quantity for its outlier test.

The *Q-Q plot* (Fig. 7, left) compares the sorted residuals with the values they would have if they were exactly normal. Points on the diagonal mean a normal distribution; a middle on the diagonal with tails bending away from it means that most sight lines behave as the error model says, but a small group is further out than a normal distribution would produce. The density plots to its right show the same thing as a histogram with the fitted Gaussian drawn on top, once in original units and once normalised. The skewness and kurtosis printed under them summarise the shape: skewness zero and kurtosis three for a normal distribution.

Read together, these plots tell what the survey is worth. A histogram centred on zero with tails on the diagonal of the Q-Q plot means clean measurements with a correct error model; a shifted histogram means a systematic error in one set-up or one observation type; heavy tails mean that a minority of sight lines, and usually a specific kind of target, is less stable than the rest.

The residual histograms of the three observation types are centred on zero, with standard deviations of 2.02 cc (3.2 μ rad) for the horizontal angle, 1.99 cc (3.1 μ rad) for the zenith angle and 0.064 mm for the distance. No station stands out in the stacked histogram, so there is no set-up with a systematic offset. The normality check gives a more differentiated picture: the bulk of the residuals follows the normal distribution, but both tails bend away from it and the kurtosis is 47 against 3 for a normal distribution. A small number of sight lines are further out than the error model allows. Ten of them, with normalised residuals beyond the data-snooping threshold, were switched off between the third and the fourth run of TL, and this is what brought σ_0 from 0.980 to 0.988.

This is the first thing the adjustment adds to the raw data. The precision column of the file could not have found these ten sight lines, because each of them is perfectly repeatable within its own set-up; they are only wrong with respect to the other set-ups that saw the same point.

All Nine Sections

Table 2 summarises the passing run of each section, and Table 3 the a posteriori error model, i.e. the standard deviations the observations turned out to have. All nine adjustments pass the χ^2 test on σ_0 . The interaction region deserves a caveat: with three set-ups and 309 degrees of

freedom its critical interval is four times wider than that of a full section, so it passes a weaker test, and it also contains the final-focus magnets that are known to have moved during the campaign.

Table 2: LGC2 Adjustment of Each Section (Passing Runs)

Sec.	Set-ups	Points	Rays	d.o.f.	σ_0
TL	51	665	2425	5133	0.988
IP	3	79	183	309	1.056
TR	51	673	2344	4866	0.995
OL	44	714	2156	4200	1.000
OR	48	742	2362	4722	0.995
FL	47	580	2099	4422	0.990
FR	45	566	2019	4230	0.981
NL	50	724	2303	4593	0.989
NR	48	720	2260	4482	0.980
All	387	5463	18151	36957	9/9 pass

Table 3: A Posteriori Standard Deviation of One Observation, per Section

Sec.	Hz angle [μ rad]	Zenith [μ rad]	Distance [mm]
TL	3.2	3.1	0.06
IP	5.3	5.1	0.08
TR	2.8	6.3	0.13
OL	3.2	4.3	0.11
OR	2.5	5.9	0.04
FL	2.8	6.6	0.09
FR	6.3	2.4	0.08
NL	3.6	2.3	0.06
NR	3.6	2.3	0.04
Whole ring	1.9	4.4	0.07

We want to be explicit about what $\sigma_0 \approx 1$ proves. The first run of TR gave $\sigma_0 = 2.26$ and the first run of NR gave 0.79 with the same observations; only the declared sigmas differed. With three variance components per section, σ_0 can always be driven to one. What can be judged is whether the resulting error model is physically sensible and consistent from one section to the next, and this is what Table 3 is for.

The angular sigmas, 2.5 μ rad to 6.3 μ rad, are consistent with a tracker used over 10–70 m. The distance sigmas, 0.04 mm to 0.13 mm, are several times larger than the raw precision column and several times larger than the ranging accuracy of the instrument. That excess is not instrument noise; it is the disagreement between set-ups on the position of the same target, which means target mounting and, as shown below, mainly one monument type. The table also shows that the nine sections are not alike. The zenith sigma varies by a factor of three (2.3 to 6.6 μ rad) and the distance sigma by a factor of three as well; FR has the worst horizontal angles and TR the worst distances. Nine independently adjusted sections therefore mean nine error models, nine free datums and nine sets of coordinate standard deviations that cannot be compared with each other.




Figure 5: Section TL after adjustment, drawn from the LGC2 output file. Left: the whole section, 51 stations, 665 points, 2425 sight lines. Right: a 90 m stretch rotated along the tunnel axis, showing the stations, the three monument rows, the magnet targets and the sight lines.




Figure 6: Section TL: histograms of the residuals of the 2425 sight lines for horizontal angle (cc, 1 cc = 1.57 μ rad), zenith angle and distance, from the LGC2 report. Colours distinguish the stations.




Figure 7: Section TL, horizontal-angle residuals: Q-Q plot against the normal distribution (left), and density of the residuals with the fitted Gaussian, in original units and normalised by the a priori sigma (right).

Consistency Across Section Boundaries

The nine sections overlap by about 18° of azimuth in total and form a closed chain TL–IP–TR–OL–OR–FL–FR–NL–NR–TL. The number of points common to neighbouring sections, counted from the point names in the raw report files, is 24 to 48 (Fig. 2); the two joints at the IP carry 47 and 48 common points, almost all of them hall and wall monuments. The IP region is small in extent but it is what closes the ring at 12 o'clock: without it, TL and TR share only ten points.

Two independently adjusted sections can be compared on these common monuments, and we did so. But this comparison has a limit that we think is worth stating plainly, because it is what motivated the last part of this work. Each section adjusted on its own carries its own free datum. Placing the sections next to each other does not make a ring: nothing forces the loop to close, and the orientation of each station, known to a few microradians, accumulates along a chain of some 140 stations per quarter ring. The vertical compo-

nent of that accumulation disappears once the stations are declared levelled, because every station is then tied to the same plumb line. The azimuth component does not, because nothing in a section-by-section treatment refers the azimuth of one section to that of another. The only way to remove it is a simultaneous adjustment in which the loop closure is a constraint inside the least squares, not a check afterwards. This is the reason for the network of the next section.

What the Section Adjustments Say About the Survey

Three findings about the survey itself came out of the section adjustments.

First, the survey is internally consistent at a level that supports the alignment tolerances used at KEK (0.2 mm transverse for quadrupoles and sextupoles [1]): all nine sections pass the global test, the residuals are centred on zero for all three observation types, and no set-up shows a systematic offset.

Second, the outliers that the data-snooping test finds are individual sight lines, not stations. In the nine sections, between 2.3 % and 3.5 % of the observations have a normalised residual beyond 2.576, against 1 % expected for a normal distribution; this is the heavy tail seen in the Q–Q plot of TL, and it is concentrated on the OV/OH monuments.

Third, 505 of the 3614 magnet targets (14 %) were sighted from a single set-up. They have no redundancy, no residual, and an error on them cannot be detected. They are concentrated in OL, NL, NR and OR, the sections with the most magnets per set-up, and fall on about ten set-ups. A second sighting of these targets would be the cheapest single improvement to the next survey.

ONE NETWORK FOR THE WHOLE RING

Data Used

The network of the whole ring is built on the tunnel monuments only: the 1147 points of the IT, IB and OV/OH rows and the hall and floor monuments, sighted from the 386 set-ups (one set-up whose export was found to be in the wrong frame was dropped). The 3614 magnet targets are not in the network. There are two reasons. The practical one is size: with all points included the computation of the full variance–covariance matrix, which LGC2 performs after the solution has converged, exceeded the memory of a 16 GB workstation. The more important one is principle: the magnets are the object being measured, and a reference network should rest on objects that do not move. The monuments also fix the geometry better, because they sit in three rows at three heights, whereas the magnet targets lie in a band 0.19 m high. Sixteen targets that were found, after a first run, to sit on magnet bodies rather than on the wall were removed. The magnet targets are computed afterwards inside the finished network; that step, and what it shows about individual magnets, is outside the scope of this paper.

Input File

The input file has the same blocks as the section files. The differences are: *CONSI LIBR TX TY TZ RZ, a four-parameter datum; every station declared IHFIX with the zenith angles corrected for the plumb-line convergence with the k of its section; and a *POLAR block with one group of sigmas per section, carried over from the section adjustments, so that the error model of each part of the ring is the one estimated from its own data. The file has 386 *TSTN blocks, 5288 *PLR3D lines, i.e. 15 864 observations, and 4989 unknowns.

Result

The adjustment converges in three iterations with $\sigma_0 = 0.991$ inside the critical interval (0.987, 1.013), no warnings, and no observation with a normalised residual above eight. Figure 9 shows the network as it comes out of the adjustment: the 386 set-ups and the 1147 monuments, drawn directly from the coordinates in the LGC2 output file, with nothing added. It is the counterpart of Fig. 2 after the adjustment,

without the magnet targets, and it closes. The standard deviations of the monument coordinates are 0.82 mm in X , 0.84 mm in Y and 0.070 mm in Z . The horizontal values are not the quality of a point: with three translations and one rotation free over a 948 m ring, every point carries the same datum uncertainty in phase. The vertical value is the real measurement quality, and it is locked by the levelled stations. The angular sigma of the whole-ring solution, $1.9 \mu\text{rad}$, is tighter than that of any section (Table 3), as expected from the far larger redundancy (10 879) and a single consistent frame.

What the Whole-Ring Adjustment Says About the Survey

The value of the whole-ring solution for judging the survey is that all 15 864 observations are now tested against one error model in one frame, so a problem anywhere on the ring shows up in the same statistics. The distribution of the normalised residuals (residual divided by its a priori sigma) of all observations is centred on zero and essentially contained in ± 3 . Broken down by observation type, 0.42 % of the horizontal angles and 0.49 % of the zenith angles exceed three sigma, close to the 0.27 % expected for a normal distribution. The distances are different: 2.04 % exceed three sigma, sevenfold the expectation, and LGC2 flags the distance residuals as non-normal (Fig. 8).

The report lets us trace that tail to its source. Sorting the sight lines with a normalised distance residual beyond three by the type of monument they point at gives Table 4: none of the 1291 sight lines to IB monuments, 0.4 % of those to IT monuments, but 6.7 % of those to OV/OH monuments. All eight worst distance residuals of the whole ring point at OV monuments. The raw precision column had already hinted at this (Fig. 4, right), but only weakly, because within one set-up an OV target repeats almost as well as any other. It is between set-ups that its position is less consistent. Two differences from the IT and IB rows are visible in Fig. 1: the OV/OH monuments are sighted across the beam line, i.e. over the longest distances of each set-up, and they are mounted on steel posts among the pipe and cable-tray supports rather than drilled into the concrete wall. Which of the two is responsible, or whether the holders themselves are the cause, cannot be decided from the survey data alone; it is a question for the next campaign.

Table 4: Sight Lines With Normalised Distance Residual Above 3, by Monument Type, Whole-Ring Adjustment

Monument row	Flagged / total	Rate
IB (inner bottom)	0 / 1291	0 %
IT (inner top)	8 / 1918	0.4 %
OV / OH (opposite wall)	126 / 1871	6.7 %
Floor and hall	9 / 208	4.3 %

The practical conclusion is not to drop the OV/OH row: its 1871 sight lines are what connect the two sides of the tunnel and they carry a third of the network. It is to give




Figure 8: Whole ring, distance residuals of the 5288 sight lines: Q-Q plot (left) and density with fitted Gaussian, original and normalised (right). Both tails leave the normal distribution; the kurtosis is 36 in original units.




Figure 9: The whole-ring network as adjusted by LGC2, drawn from the coordinates of the output file: 386 set-ups (squares) and 1147 tunnel monuments (dots), coloured by section. No magnet target is part of this network.

that row its own, larger distance sigma in the error model, to check the OV holders before the next campaign, and to prefer the IT and IB rows when a monument has to serve as a reference between epochs.

CONCLUSION

The 2015 laser-tracker survey of the SuperKEKB main ring has been adjusted with LGC2, section by section and then as one network of the whole ring. The raw files carry a per-point precision of 0.04 mm median, which shows that the instrument was performing well but says nothing about the agreement between set-ups. That agreement is what LGC2 measures. All nine sections pass the global test with a physically sensible error model, 2.5 μ m to 6.3 μ m for the angles and 0.04 mm to 0.13 mm for the distances, and adjacent sections share 24 to 48 common points through which they can be compared. The section-by-section treatment has one limit: it leaves each section with its own datum and lets the azimuth error accumulate around the ring. The whole-ring network, built on the 1147 tunnel monuments with the stations tied to the plumb line, removes it: $\sigma_0 = 0.991$, vertical coordinate sigma 0.07 mm, and one error model for the entire ring.

The statistics of that adjustment locate the weak part of the survey. The distance residuals are non-normal, and the tail belongs almost entirely to the OV/OH monument row on the wall opposite the IT and IB rows: 6.7 % of its sight lines are flagged, against 0.4 % for the IT row and none for the IB row. The IT and IB monuments drilled into the wall are the ones to rely on; the OV/OH monuments and their mounting should be inspected before the next campaign, and the 505 magnet targets that were sighted from one set-up only should be sighted from a second one.

Taken together, these results show that LGC2 can take the SuperKEKB survey data as they are, reproduce a consistent network of the whole ring in one adjustment, and return a diagnosis of the survey that the raw data alone cannot give. The next step is the one that motivated this work: adding the HLS readings near the IP, and later FSI observations, to the same adjustment.

ACKNOWLEDGEMENTS

We are grateful to Guillaume Kautzmann and Francis Klumb of the Geodetic Metrology group, Beams Department, CERN, for introducing us to LGC and for their guidance in using it. We also thank the SuperKEKB magnet group and the survey teams of the 2015 campaign for the data.

REFERENCES

- [1] M. Masuzawa, T. Adachi, T. Kawamoto, and Y. Ohsawa, “Alignment of SuperKEKB main ring magnet system”, in *Proc. IWAA’16*, Grenoble, France, Oct. 2016.
- [2] M. Masuzawa, T. Nagasaki, S. Nakamura, Y. Ohsawa, and R. Ueki, “Long-term alignment stability of the SuperKEKB tunnel”, in *Proc. IPAC’25*, Taipei, Taiwan, Jun. 2025, pp. 418–420. doi:10.18429/JACoW-IPAC2025-MOPM044
- [3] Y. Funakoshi *et al.*, “Beam commissioning of SuperKEKB”, in *Proc. IPAC’16*, Busan, Korea, May 2016, pp. 1019–1021. doi:10.18429/JACoW-IPAC2016-TUOBA01
- [4] M. Masuzawa, “Experiences with SuperKEKB IR installation & alignment”, presented at eeFACT’25, Tsukuba, Japan, Mar. 2025, unpublished.
- [5] LGC2 – Logiciel Général de Compensation, CERN, <https://lgc2.docs.cern.ch/>