ILC BDS Beam-Based Alignment and Tuning

Introduction

This document provides a complete description of the strategy used to simulate the BBA and tuning of the BDS from a post-survey configuration through to collisions at nominal luminosity. The assumed starting point is after a physical survey has taken place, and any initial optics verification. The initial errors reflect this starting assumption, but the strategy is also designed to be re-usable as retuning is necessary on regular intervals and should thus be as efficient as possible. Results of a multi-seed simulation run are included which demonstrate that the stated tolerances used are sufficient to deliver the design luminosity. The simulation is semi-dynamic in that throughout the tuning process the beam orbit is maintained using the BDS 5-Hz feedback system with kicker field errors, finite BPM resolutions and realistic tuning of feedback coefficients. Ground motion and component vibration are omitted here however for the purpose of keeping the simulation time manageable. However, the simulation results should not be affected by the addition of these things, assuming: No significant luminosity degradation happens due to dynamic effects on the timescale of the tuning; and a sufficient number of additional pulses are averaged to mitigate any effects of pulse-pulse jitter from the fast ground motion and component jitter. Given that studies of the 5-Hz feedback system show that luminosity loss is controlled to the ~<2% level under the harshest ground motion models on the few-hour timescale relevant here, it is assumed that the addition of these effects will only affect the results of this study in lengthening the time required for the tuning process depending on the nature of the fast jitter and ground motion.

Initial Error Assumptions and Simulation Notes

All magnets in the BDS are assumed to be individually mounted on movers (apart from the final doublet- see later) with horizontal and vertical degrees of freedom. Sextupoles additionally have roll (rotation about z) capability. All magnets have BPMs physically attached to their upstream edge. All magnets also have their own power supplies. The final cryomodule string (FCMS) of magnets (SF1,QF1,OC1,SD0,OC0,QD0,) are assumed to be fixed in their installation positions, and mounted in simulation on a common girder with 4 degrees of freedom of movement (x,x',y,y'). The reference BPMs for the FCMS are cavity BPMs attached to the upstream face of SF1 and SD0. The final octupoles are actually co-wound on the Sextupoles, and are modeled as thin multipole elements inside SF1 and SD0.

The table in Appendix A shows the set of nominal initial error parameters used. Tolerance studies are performed by changing the error parameters with reference to this initial set.

The simulation package used for these studies was Lucretia, which operates in the Matlab framework. Parallel simulations were carried out with the use of the Matlab Compiler.

The lattice used in this simulation is the ILC2006c lattice, which contains the latest BDS design changes with 2*14mrad IP crossing angles.

Some studies were performed to understand the best way to model the particle beam (see the results section). The Lucretia 6D sparse structure was used for tracking (single bunch

used). This generates a beam in which the rays are distributed according to a grid in the longitudinal phase plane, with 31 grid points in z and 11 grid points in P used here. At each grid point 9 rays are assigned: one at the nominal centroid orbit, and the remaining 8 at +/- sqrt(9/2) sigmas in each of x, px, y, py. This is an emulation of the "slices and macroparticles" beam typically used in linac codes such as LIAR. This allows a much faster tracking-time than a macro-particle representation and was found to give the same results within statistical errors.

All simulations are performed using 100 random seeds, applied to initial error assumptions, beam generation and control system imperfections.

Here, only 1 beam is simulated. The final tuning using sextupole multi-knobs tunes on geometric luminosity, no errors on the luminosity measurement are included, although pulse-pulse fluctuations in the luminosity are present due to the orbit-feedback. Additional averaging for each luminosity measurement needs to be performed if the luminosity measurement error becomes dominant. The luminosity measurement is assumed to arise from a flux measurement of the beamstrahlung e+e- pair radiation. A bunch-by-bunch measurement of this quantity has been put forward in the ILC BCD, so assuming for each pulse the luminosity calculation averages ~3000 independent measurements, statistical errors should be small.

Beam Parameters and Tuning Goals

The beam parameters simulated are those of the nominal parameter set as defined in the ILC BCD. The BCD emittance growth budget for the BDS is 6nm vertical. Hence the initial beam is set up with 1e-5/3.4e-8 normalised horizontal/vertical emittances entering the BDS on-axis. The goal is to all of the simulated seeds to tune to give better than nominal luminosity given the input conditions. Additionally, some luminosity overhead will also be required as GM and component jitter degrade luminosity over time requiring periodic re-tuning. The amount of luminosity overhead achieved thus dictates the required periodicity of the retuning.

Alignment and Tuning Procedure Overview

This is an overview of the simulation steps to tune the BDS; details of individual steps are included below:

1. Track perfect lattice with generated beam and nominal beam parameters which lead to ~ nominal luminosity, use this as a reference beam to compare tuned beam with.
2. Apply all error sources as specified in Appendix A.
3. Switch off Sextupoles and Octupoles.
4. Perform initial BBA using Quad movers and BPMs (see section on Quad BBA) to get beam through to IP.
5. Quadrupole BPM alignment (see section on BPM alignment).
6. Perform Quadrupole BBA.
7. Align Sextupole BPMs (sextupoles left aligned to the beam but switched off at this stage).
8. Move FCMS to minimize FCMS BPM readings.
9. Align Octupole BPMs.
10. Activate sextupole and octupole magnets.

11. Rotate whole BDS about first quadrupole to pass beam through nominal IP position or iteratively move FCMS and re-apply DFS BBA (simulation of finding the other beam at nominal collision point).
12. Zero all BPM readings -> set this as reference orbit for 5 Hz feedback.
13. Apply sextupole multiknobs to tune out IP aberrations and maximise luminosity (see multiknobs section).
14. 5-Hz feedback system used throughout to maintain orbit whilst tuning. See section on 5-Hz feedback for more information.

Quadrupole BBA

This is a procedure used to put the quadrupoles in as straight and dispersion-free position as possible with the beam passing through the magnet centers (within BPM-magnet alignment tolerances). This is achieved using the movers on the magnets with a DFS steering algorithm which also minimises the extent of the moves applied. The algorithm takes the set of Quad BPM values as input and solves for the set of horizontal and vertical quad moves given a coupled response matrix calculated from the ideal lattice. The coupled matrix is required as there are skew-quads in the lattice which have non-zero strengths when used for coupling correction. See Appendix B for definitions of the constraint and correction vectors and the response matrix.

The algorithm used is contained in the 'lscov' Matlab routine. This solves the linear equation $A*c = b$; where c is the correction vector: the desired vector of quad moves plus an additional kick in first quad. (Additional kick provided by further offsetting of that quad with its mover); b is the measurement vector, containing BPM readings (orbit correction) plus 2 sets of dispersive BPM readings (nominal plus and minus a 1% incoming energy change) plus a null vector equal to the size of c to constrain the mover solution; A is the response matrix (see Appendix B). Additionally, a weight vector (w) is supplied to control the solution. The lscov routine then forms a least-squares solution which minimises: *(b- A*c)'*diag(1/w^2)*(b - A*c)*. The vector w is formed from the BPM resolution errors plus expected initial RMS alignment errors for the mover weights. A weight multiplier is added to all the BPM w elements and additionally to just the dispersive BPM w elements. Optimal values for these additional weights were found that are approximately optimal for all seeds simulated. A BPM weight of 6 and a dispersive BPM weight of 105.2 are used. As seen in the results section, this gives a smooth transition of the beam through the magnets and keeps the maximum required mover motion to < 1mm.

The algorithm is applied in 3 stages over the BDS between the first quad and QF1 with a 10 magnet overlap with each stage. Then the FCMS is moved to minimise its BPM readings.

BPM – Magnet Alignment

Quadrupole BPMs

The Quadrupole BPMs are aligned to the field centers of the quadrupoles they are attached to using a nulling Quad-shunting technique. This is done by cycling the Quad strength from 100% -> 80% and back and calculating the quad offset w.r.t. the beam from

the change in downstream BPM responses. A 1-d minimisation routine is then used to move the quad with its mover (the FCMS frame mover used for final doublet) and minimise the downstream BPM responses when the Quad strength is cycled as above. The minimisation routine used is the fminbnd routine included in Matlab, which uses a golden section search and parabolic interpolation method.

Using this in preference to other alignment techniques, such as ballistic alignment, was implemented as it doesn't require completely deactivating the quads. This makes it more robust for frequent application with a limited impact on the lattice. A 20% change in strength is small enough that the magnet can return to full strength to high precision whilst providing enough of a change in downstream BPM readings to perform the alignment.

The number of downstream BPMs used in the routine was optimized for x- and y-independently. 17 BPMs are used for the x-plane and 11 for the y-plane. Lattice error's affect this optimisation, changing the error conditions in principal would require re-optimisation.

The Quad offset is calculated for each BPM reading from:

$x_{Quad} = \Delta x_{BPM} / \left(\Delta R_Q(1,1) * R(1,1) + \Delta R_Q(2,1) * R(1,2) \right)$, and then averaged.

NB: The extraction line Quads are used when there aren't enough remaining Quads in the BDS which requires some extraction line Quads to be instrumented with BPMs.

Sextupole BPMs

The sextupole BPMs are aligned by moving the sextupole through the beam and measuring the response on a downstream BPM. A parabola is fitted to the response and the Sextupole offset is read from the parabola minima. 50 points across +/- 2mm from the original BPM center point are used for the fit.

Octupole BPMs

The tail-folding octupole BPMs are aligned in a similar fashion to the sextupole BPMs. Here the response is cubic, with the alignment read off from the zero-crossing of the 2nd derivative of the fit to the downstream BPM response. It was found that the 6th octupole's strength needed to be increase an order of magnitude to get a large enough downstream BPM response to align its BPM better than the initial 30um value, so was left with the initial alignment error of 30um RMS.

Sextupole Multi-Knobs

Aberrations of the IP beam (dispersion, waist-shift and x-y coupling) exist after initial BBA. These are tuned-out using the strong coupling of FFS sextupole offsets to these aberrations. Dispersion and waist are coupled through sextupole offsets by:

$$\Delta s_{x,y} \sim \Delta x . K_2^s L \beta_{x,y}^s \beta_{x,y}^* \cos(2.\mu)$$

$$\Delta \eta_{x,y}^* \sim \Delta(x,y) . K_2^s L \eta_{x,y}^s \sqrt{\beta_{x,y}^s \beta_{x,y}^*} \sin(\mu)$$

Offsets in the sextupoles also causes <x'y> coupling, which is the dominant coupling term affecting the IP luminosity. The skew quadrupole SQ3FF is used to remove this coupling term during the tuning process.

We have 5 linear knobs required, Waist (x), Waist (y), Dispersion (x), Dispersion (y) and <x'y> coupling. The first 3 sextupoles are used (available on independent movers) plus the movers of the FCMS are also available. Each is moved in x and y, calculating linear slopes for each desired IP aberration. Then, the generated response matrix is inverted using the Matlab lscov routine (see quad bba section for more description of lscov). This calculates orthogonal moves to produce the desired knobs, and does so using the minimal integrated sextupole movements / skew quad dK. The optimised multi-knobs are shown in Appendix C.

Additionally, the four skew quads in the skew-correction system at the front-end of the BDS are cycled through (adjusting magnet strength). Tuning is performed on luminosity to remove the smaller contributions to beam-size growth that arises from the <xy> coupling generated from the accumulated roll errors of all BDS magnets.

When errors are introduced into the simulation, the knobs loose exact orthogonality due to small changes in the lattice mostly due to magnet strength errors. The individual knobs thus need to be applied in an iterative fashion to converge on the optimal luminosity.

Additionally, higher-order aberrations were found to be present after the linear tuning, mainly due to Sextupole roll misalignment. Strength changes of all 5 sextupoles and roll adjustments to the first three provide coupling to the higher-order aberrations. Attempts to generate orthogonal knobs based on second-order IP sigma matrix terms did not prove very successful. However, by simply iterating through the 8 non-linear knobs (re-applying the linear tuning at each stage) and optimising on luminosity, adequate convergence was seen.

So, in total there are 6 linear and 8 higher-order knobs to be applied. The order in which the knobs are applied was found to be important. First, the dispersion, <x'y> coupling and waist knobs are iterated through until convergence. Then the skew quads are iterated through. Then the strength and roll knobs are applied. Different simulation seeds converge at different rates, a sequence of knob applications was developed that produced convergence to a given seeds' maximal luminosity in each case. The exact implementation and sequence of knobs used is described in Appendix D.

Notes on Experimental Knob Verification

In cases with larger lattice errors, the linear knobs don't always converge satisfactorily, or take much longer to converge. It may be possible to experimentally verify and adjust the linear tuning knobs. To do this the IP waist, dispersion and <x'y> coupling terms need to be measured whilst scanning the knobs.

The dispersion can be measured by ramping the incoming beam energy and measuring the transverse deflection of the IP beam position. This can be done either with an IP laserwire system or measuring the relative deflection w.r.t. the other beam using the fast-feedback monitors.

The waist positions are hard to experimentally determine. It may be possible to gain information about how well the waist knobs are performing by applying a waist shift using the knobs and changing the timing of the other beam, giving a known waist shift, and searching for the optimal waist position using the luminosity monitor.

Information on the <x'y> coupling term can be gained with the opposing beam switched off and measuring the IP sigma(1,3) term using a diagonal wirescanner/laserwire at the entrance to the first extraction quad which is at a π phase advance from SQ3FF.

It may also prove possible to access some of these IP terms from topological studies of the beamstrahlung radiation using pair and photon detectors.

5-Hz Feedback System

To keep the beam centered in the sextupole magnet BPMs throughout the multi-knob tuning process and to provide a dynamic modeling environment, a 5-Hz feedback system is applied throughout the alignment and tuning procedure.

The feedback system consists of 6 horizontal correctors, and the 5 horizontal BPM readings in the sextupole BPMs. Also, 6 vertical correctors with the 5 vertical BPM readings in the sextupole BPMs in addition to a further set of vertical BPM readings on 5 FFS quad BPMs. Numbering the magnet BPMs from the start of the BDS, the additional vertical BPM readouts used are: [66 84 100 102 119].

The feedback algorithm uses lscov to invert the corrector -> BPM response matrix with BPM resolutions used as weights (this biases the correction to the sextupole BPMs which have the highest tolerances). A feedback gain of 0.1 is applied to the corrector response vector. The feedback system corrects a step function error within ~50 pulses. This was found to be a good model to deal with ground motion and component vibration effects.

After each tuning step the feedback model is run with 300 pulses tracked to ensure convergence to the desired orbit.

Note that each time a sextupole is moved during the multi-knob tuning phase, the BPM reading offset is adjusted to keep the desired offset in the sextupole magnet whilst the feedback system is running.

Also note that the response matrix needs to be coupled due to the presence of skew quads with potentially non-zero fields. And the response matrix needs to be re-calculated each time a magnet field is adjusted.

Simulation Results

The simulation results shown here are for the nominal error set shown in Appendix A (100 random seeds).

Beam Representation

This is a comparison of a nominal beam tracked through a perfect BDS lattice using different beam representations. The macro-particle representations are generated and tracked 20 times with the mean and RMS geometric luminosities shown. The sparse beam representation is as described in the simulation section above and is uniquely defined. The reason for the particular structure used for the sparse structure is to match that desired for tracking in the linac so that in future simulations the linac and BDS can be co-simulated with the same bunch structure.

	L / % nominal
Macro-particle (10K)	108.7 +/- 3.9
Macro-particle (20K)	105.7 +/- 3.2
Macro-particle (40K)	103.9 +/- 2.8
Macro-particle (80K)	101.4 +/- 1.4
Sparse (31 x 11)	99.93

The sparse structure was chosen for the simulations as it gives a closer match to the calculated design luminosity and can be tracked much faster. The drawback of the sparse representation is that it gives limited information on non-gaussian transverse beam structure. As a consistency check, an 80K macro-particle bunch is tracked through the final tuned lattice for each seed. The mean results with the macro-particle beam are compatible with the sparse beam within statistical errors.

Beam-Based Alignment

After the BBA procedure described in the Quadrupole BBA section above, the mean and RMS absolute magnet positions through the BDS are shown in figure 1. The maximum deviation required by any magnet mover for any seed is ~<1mm.
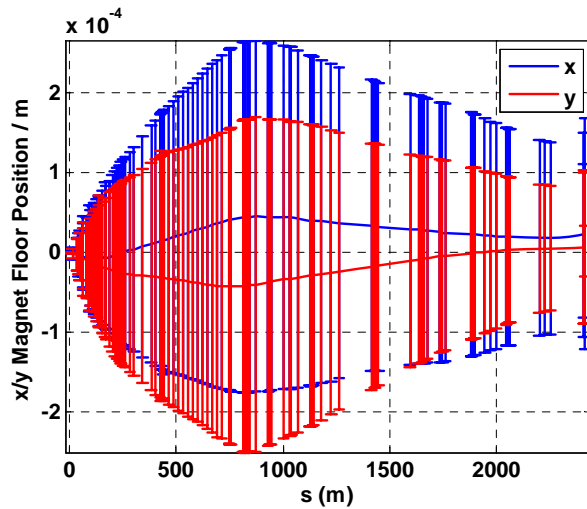
Figure 1: Absolute positions of BDS magnets post-alignment (mean and +/- 1 standard deviation from 100 seeds).

Magnet – BPM Alignment

Figure 2 below shows the RMS alignment of the BDS quadrupole magnets to their respective BPMs after Quad-shunting has been applied. Figure 3 shows the RMS alignment for the 5 FFS Sextupoles and 6 tail-folding Octupoles after the respective alignment techniques have been applied.
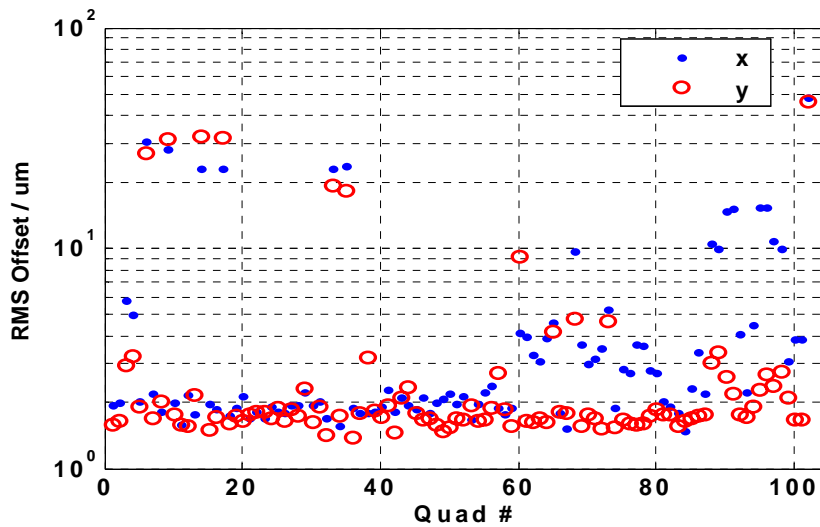

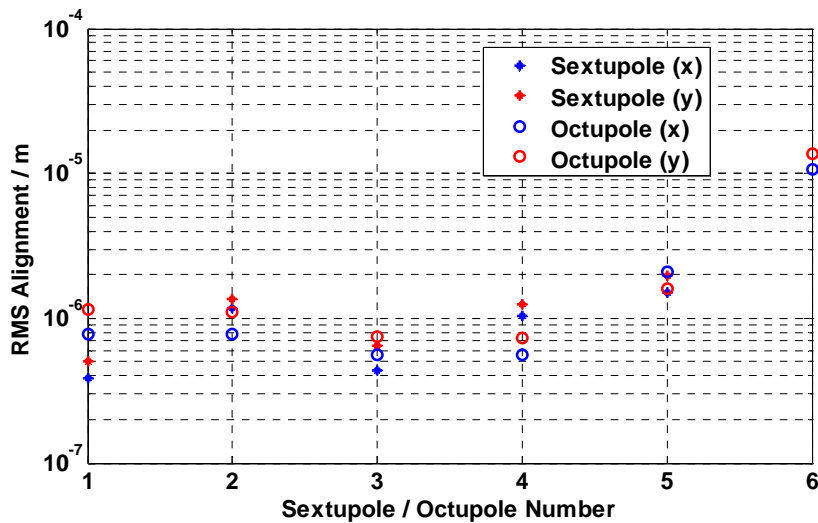**Figure 2: BPM - Quadrupole alignment (RMS from 100 seeds).**

Figure 3: BPM - Sextupole/Octupole (RMS from 100 seeds)

Luminosity Performance

Figure 4 shows the percentage of seeds exceeding a given percentage of nominal luminosity. All simulated seeds exceed nominal luminosity, the median result gives a 7.9% luminosity overhead. The median result for the 80K macro-particle bunch gives a luminosity of 105.9%, which given the 1.4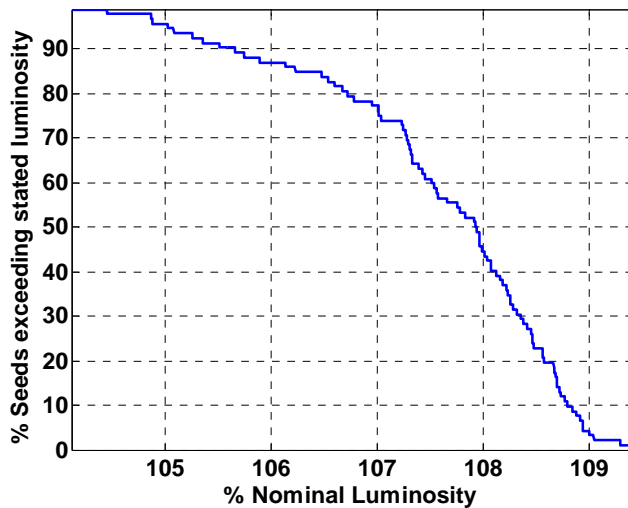% RMS spread of results when tracking a perfect lattice is in good agreement with the sparse bunch structure result. Figure 5 shows the mean and RMS luminosity as a function of multi-knob application. Each x-axis unit is another numbered sequence as described in the section on Sextupole multi-knobs above. It can be seen how the linear knobs are iterated until convergence, then a step is seen where the non-linear knobs are applied and iterated. The dominant effect here comes from the sextupole roll correction. Figure 6 shows the percentage of seeds exceeding design luminosity as a function of knob-application iteration. Figure 7 shows the vertical and horizontal beam spot sizes at the IP.
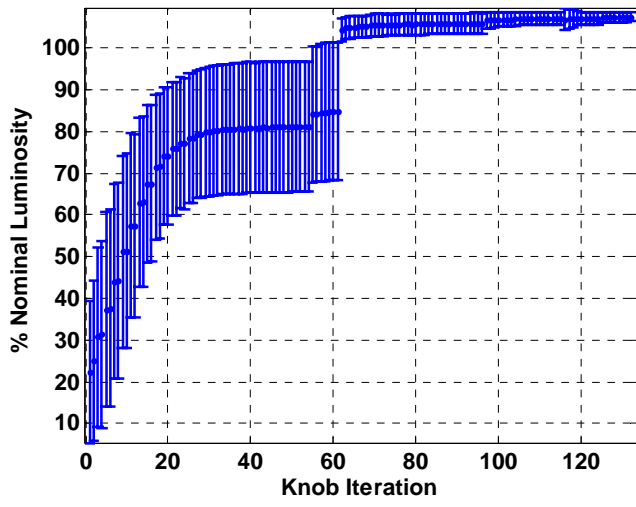

Figure 4: Percentage of seeds exceeding quoted % Nominal Luminosity (100 seeds).

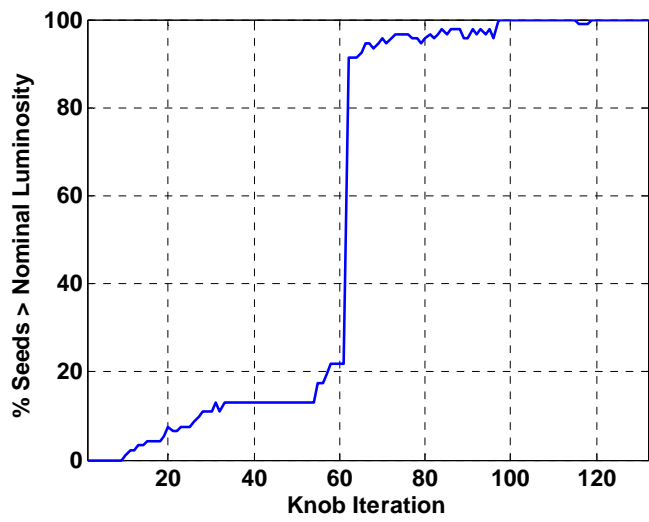**Figure 5: Mean and RMS luminosity vs, multi-knob iteration # (100 seeds).**

Figure 6: Percentage of seeds that exceed nominal luminosity at each multi-knob iteration (100 seeds).
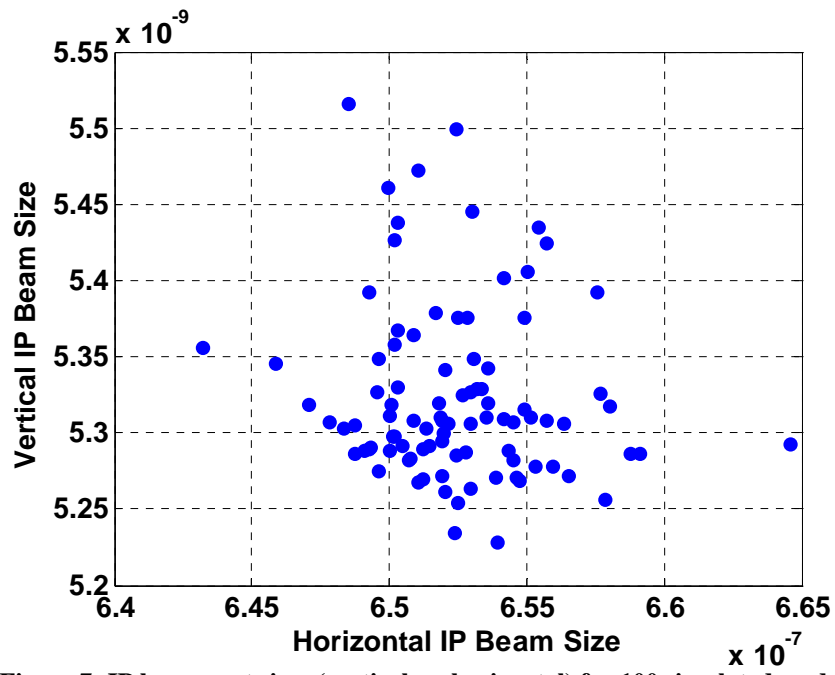
Figure 7: IP beam spot sizes (vertical vs. horizontal) for 100 simulated seeds. Nominal values are 655nm (x) and 5.7nm (y).

Appendix A – Initial Errors

Errors are assumed to have a normal distribution, with RMS values quoted in the table below:

Quad, Sext, Oct x/y transverse alignment	200 um
Quad, Sext, Oct x/y roll alignment	300 urad
Initial BPM-magnet field center alignment	30 um
dB/B for Quad, Sext, Octs	1e-4
Mover resolution (x & y)	50 nm
BPM resolutions (Quads)	1 um
BPM resolutions (Sexts, Octs)	100 nm
Power supply resolution	14 - bit
FCMS: Assembly alignment	200 um / 300urad
FCMS: Relative internal magnet alignment	10um / 100 urad
FCMS: BPM-magnet initial alignment (i.e. BPM-FCMS Sext field centers)	30 um
FCMS: Oct – Sext co-wound field center relative offsets and rotations	10um / 100urad
Corrector magnet field stability (x & y)	0.1 %
Luminosity (pairs measurement or x/y IP sigma measurements)	Perfect

Appendix B – Quad BBA

The quad BBA algorithm solves $A*c=b$ for c given b and the response matrix A and a weight vector w. A, c and b are defined below, see the Quad BBA section for a description of w. So, for the contiguous section of quads 1 -> n being aligned:

$$b = \begin{pmatrix} B_x^0 \\ B_y^0 \\ B_x^- \\ B_y^- \\ B_x^+ \\ B_y^+ \\ c \end{pmatrix}$$, where $B^{0,-,+}$ correspond to (x or y) BPM readings with nominal energy and

nominal +/- 1% incoming energy, c is the corrector vector defined below. B is the vector of BPMs:

$$B = \begin{pmatrix} b_2 \\ b_3 \\ \vdots \\ b_n \end{pmatrix}.$$

$$c = \begin{pmatrix} q_2^x \\ q_3^x \\ \vdots \\ q_{n-1}^x \\ k_1^x \\ q_2^y \\ q_3^y \\ \vdots \\ q_{n-1}^y \\ k_1^y \end{pmatrix}$$, where $q_i^{x,y}$ are quad moves for quad i in the horizontal or vertical plane, k is a

kick generated in quad 1 (produced by additional movement) in the horizontal or vertical plane.

A then, is the response matrix connecting c and b:

$$A = \begin{pmatrix} T^0 \\ T^- \\ T^+ \\ diag(1) \end{pmatrix}$$, where the superscript refers to nominal, plus 1% and minus 1% energy

orbits, the unity diagonal matrix provides the grounds for mover limitation and:

$$T = \begin{pmatrix}
-1 & 0 & 0 & \cdots & \cdots & R_{1,2}(1,2) & 0 & 0 & 0 & \cdots & \cdots & R_{1,2}(1,4) \\
M_{2,3}^{XX} & -1 & 0 & \cdots & \cdots & R_{1,3}(1,2) & M_{2,3}^{XY} & 0 & 0 & \cdots & \cdots & R_{1,3}(1,4) \\
M_{2,4}^{XX} & M_{3,4}^{XX} & -1 & \cdots & \cdots & R_{1,4}(1,2) & M_{2,4}^{XY} & M_{3,4}^{XY} & 0 & \cdots & \cdots & R_{1,4}(1,4) \\
\vdots & \vdots & \ddots & \ddots & \ddots & \vdots & \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\
M_{2,n}^{XX} & M_{3,n}^{XX} & M_{4,n}^{XX} & \cdots & M_{n-1,n}^{XX} & R_{1,n}(1,2) & M_{2,n}^{XY} & M_{3,n}^{XY} & M_{4,n}^{XY} & \cdots & M_{n-1,n}^{XY} & R_{1,n}(1,4) \\
0 & 0 & 0 & \cdots & \cdots & R_{1,2}(3,2) & -1 & 0 & 0 & \cdots & \cdots & R_{1,2}(3,4) \\
M_{2,3}^{YX} & 0 & 0 & \cdots & \cdots & R_{1,3}(3,2) & M_{2,3}^{YY} & -1 & 0 & \cdots & \cdots & R_{1,3}(3,4) \\
M_{2,4}^{YX} & M_{3,4}^{YX} & 0 & \cdots & \cdots & R_{1,4}(3,2) & M_{2,4}^{YY} & M_{3,4}^{YY} & -1 & \cdots & \cdots & R_{1,4}(3,4) \\
\vdots & \vdots & \ddots & \ddots & \ddots & \vdots & \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\
M_{2,n}^{YX} & M_{3,n}^{YX} & M_{4,n}^{YX} & \cdots & M_{n-1,n}^{YX} & R_{1,n}(3,2) & M_{2,n}^{YY} & M_{3,n}^{YY} & M_{4,n}^{YY} & \cdots & M_{n-1,n}^{YY} & R_{1,n}(3,4)
\end{pmatrix}$$

Where:

$$M_{i,j}^{XX} = R_i^q(2,1).R_{i,j}(1,2) + \left(R_i^q(1,1) - 1\right)R_{i,j}(1,1) + R_i^q(3,1).R_{i,j}(1,3) + R_i^q(4,1).R_{i,j}(1,4)$$

$$M_{i,j}^{XY} = R_i^q(2,3).R_{i,j}(1,2) + R_i^q(1,3).R_{i,j}(1,1) + \left(R_i^q(3,3) - 1\right)R_{i,j}(1,3) + R_i^q(4,3).R_{i,j}(1,4)$$

$$M_{i,j}^{YY} = R_i^q(1,3).R_{i,j}(3,1) + R_i^q(2,3).R_{i,j}(3,2) + \left(R_i^q(3,3) - 1\right)R_{i,j}(3,3) + R_i^q(4,3).R_{i,j}(3,4)$$

$$M_{i,j}^{YX} = \left(R_i^q(1,1) - 1\right)R_{i,j}(3,1) + R_i^q(2,1).R_{i,j}(3,2) + R_i^q(3,1).R_{i,j}(3,3) + R_i^q(4,1).R_{i,j}(3,4)$$

and, R is the transport matrix element specified in brackets between the downstream edge of Quad *i* and the BPM on Quad *j*. R^q is the transport matrix element specified through Quad *i*.

Appendix C – Linear Sextupole Multi-Knobs

Moving the sextupoles and FCMS through the beam +/- 20um and SQ3FF through its range, the following were found to be good correction knobs for horizontal and vertical waist and dispersion correction. SQ3FF only used for correction of the x'-y coupling term.

Knob	SF6 (x)	SF6 (y)	SF5 (x)	SD4 (x)	SD4 (y)
Waist (x)	-1	0	-0.2906	0.2401	0
Waist (y)	0.0044	0	1	0.3215	0.0002
Dispersion (x)	-0.2525	0	-1	0	-0.0026
Dispersion (y)	0	1	-0.0012	0	-0.9781

Appendix D – Description and Sequence of Multi-Knob Application

The parameter each knob is tuned on is geometric luminosity. The initial stages will use IP laserwire scanners to calculate this quantity to get the initial tuning. Later with 2 operational beams it will be possible to tune on the IP luminosity monitor signal. Here it is assumed that the other beam is present from the start which gives very good sensitivity to luminosity. This speeds up the simulation and the end results are expected to be the same.

The sequence of tuning knobs applied for all simulated seeds is the following:

[lseq*N [5 lseq 4 lseq 3 lseq]*N]

Where, lseq = [1 2 1 2 1 2], and the codes correspond to the following:

1. x & y dispersion (x and y dispersion knobs applied in parallel) followed by 1d optimisation of <x'y> knob. Initially dispersion is directly minimised through IP dispersion measurement (ramping energy over 1% during pulse train and using FFB beam-beam deflection measurement). Later the knobs are applied to maximise luminosity with 1d optimiser.
2. Consecutively run 1d optimiser for x waist then y waist.
3. Iterate through 5 sextupole magnet strengths. For each sextupole, run 1d optimiser whilst applying lseq after each magnet change attempt.
4. As for above, except for sextupole roll and only over the first 3 sextupoles.
5. Cycle through 4 skew quad strength settings in BDS skew correction system, optimising luminosity for each one. This corrects any residual <xy> coupling terms that aren't picked up during the step 1 procedure.

Every time any one knob is touched, the 5-Hz feedback is iterated until the orbit stabilises.

So, the linear knobs are applied until luminosity stops increasing. Then the coupling and second order knobs are iteratively applied (re-applying linear knobs after each) until this maximises luminosity.

The optimiser used here was the fminbnd algorithm in Matlab. To ensure machine protection during tuning, the boundary conditions need to be carefully set to avoid errant orbits introduced by the optimiser. Greatest efficiency can be met by choosing optimal convergence precision and boundary conditions, which can also be dynamically adjusted as the system gets closer to its optimal tuning.

The number of knob iterations required is different for different error assumptions: higher magnet strength errors, for example, require more iterations.