Analysis of a method to measure the dodecapole component of the LHC triplet magnets with a wobbling closed bump

V. Ziemann

The Svedberg Laboratory, 75121 Uppsala, Sweden

(Dated: May 13, 2002)

Abstract

Due to finite manufacturing tolerances the triplet quadrupoles used in the interaction regions of LHC will have non-vanishing multipole errors which need to be measured and corrected in order to guarantee sufficient dynamic aperture at collision energy. Here we discuss a method to measure the unwanted multipole components by oscillating a closed orbit bump in the interaction region and observing the orbit at pickups outside the bump. The beam's response will contain very weak signals at harmonic frequencies of the sinusoidal excitation. Even though the amplitude of the harmonic signal will be below the resolution of the position monitor system, it can be made visible by adding noise to the original pickup data and subsequent careful filtering and averaging. We use a simple computer model to simulate the oscillating bump that generates pickup signals and then analyze those in a sophisticated signal processing chain in order to retrieve the magnitude of the unwanted multipole components.

PACS numbers: 29.27.Fh

I. INTRODUCTION

In LHC the dominant sources of non-linearities at collision energy – apart from the beam-beam non-linearities – are the quadrupoles nearest to the interaction points (IP) which are very strong in order to provide small beam sizes at the IP and have a large aperture to accomodate the beams' separation due to the finite crossing angle. Another aggravating factor is the high radiation dose from the IP region. Since the quadrupoles are relatively far away from the IP the beta functions are very large, around 4500 m [1]. Furthermore, finite manufacturing tolerances will cause the presence of higher multipoles in the quadrupoles, especially at full excitation. These non-linearities, together with the large beta functions, will be the dominant factors that restrict the dynamic aperture at collision energy. Of the non-linearities, the dodecapole component will be one of the most pronounced [2] and, if left uncorrected, would reduce the dynamic aperture by more than 20 %. This poses the question whether there are beam based methods to determine the magnitude of the dodecapole component and whether the triplet correction magnets – they will be introduced to compensate the dodecapole's adverse effect – are actually compensating rather than enhancing the detrimental non-linearities.

In this report we will discuss a method that is based on driving an oscillating closed local bump across the triplet quadrupole magnets which is similar but simpler to the one presented in Ref. [3]. The x^5 kick from the dodecapole component will cause the local bump to be not closed and harmonics of the wobbling frequency will be visible as very weak signals in the orbit outside the bump. In section 2 we will estimate the magnitude of the harmonic oscillations to be on the order of a few times 10^{-8} m which is well below the sensitivity and also the digitization or quantization threshold of the beam position monitor (BPM) system [4]. We will show, however, that by adding a finite amount of noise – equivalent to the magnitude of the BPM resolution – to the BPM signal, the weak signal can be elevated to observable levels. This can be understood in the sense that the small harmonic oscillations modulate the noise and appear thus above the BPM's quantization threshold. Of course the wanted harmonic signals are completely buried in the noise and we need to employ some advanced digital signal processing (DSP) [5] algorithms in order to extract the weak harmonic signals.

For this purpose we developed a small signal processing toolbox [6] that allows testing

signal processing algorithms in a very simple way on the command line. All programs of the suite handle input streams of ASCII data samples, manipulate them and write the results back to standard output. In this way they act as very much like other UNIX programs that pipe their output into other programs. The program suite contains programs for digital filter design, Finite-Impulse-Response (FIR)-filters, image rejection mixers, adaptive line enhancers (ALE), decimators and many other signal processing tools.

The signal that is analyzed is the beam's position at a fixed BPM which means that it is sampled once per turn at a rate of 11.245 kHz. The noise that is added is spread over the entire spectral range from 0 to the Nyquist frequency of 5.6 kHz. The wobbling signal, on the other hand, is located at very low frequencies of a few Hz. This implies that we can low-pass filter the signal heavily and thereby remove the noise at higher frequencies. Constructing such a filter as a single stage FIR filter is impossible, but building it in three decimating stages works very well. It is then possible to discern some of the wobbling signals by averaging the spectra of the decimated signals. The detection sensitivity can be further increased by using an adaptive line enhancer which is an adaptive filter that efficiently finds coherent signals in noise. The combination of these algorithms makes it possible to detect the presence of the very weak signals that come from the up-mixing of the wobbling frequencies. We will discuss the details in subsequent sections.

The remainder of this report is organized as follows: we first calculate the magnitude of the mixing signals that can be expected using current LHC parameters, and follow up by a description of a simple four-dimensional tracking program that is used to verify the analytical estimates and produce data for the signal processing analysis that is discussed in the next section. We then investigate the effect of imperfections, like non-closure of the bump, on the detection sensitivity and finally discuss further line of work needed in the conclusions.

II. ANALYTICAL ESTIMATES

In order to estimate the magnitude of the mixing signals we will consider a simple model consisting of an ideal closed bump made up of two orbit correctors 180 degree apart kicking the beam by Θ (rad). In LHC four corrector magnets will be used to implement the bump. This does not alter the physics of the following discussion as only the beam's offset in the

dodecapole magnets matters, not the detailed orbit between the steering magnets and the dodecapoles. The two 90 degree transfer matrices between the corrector and the location of the non-linearity and vice versa are given by

$$S = \begin{pmatrix} 0 & \sqrt{\beta_1\beta_2} \\ -1/\sqrt{\beta_1\beta_2} & 0 \end{pmatrix} \tag{1}$$

where β_1 is the beta function at the location of the corrector and β_2 that of the location of the non-linearity. Thus the beam's offset at the location of the non-linearity is

$$x = \sqrt{\beta_1\beta_2}\Theta \tag{2}$$

where it receives a kick given by

$$\Delta x' = f(x) = \frac{K_2 L}{2!}x^2 + \frac{K_3 L}{3!}x^3 + \frac{K_4 L}{4!}x^4 + \frac{K_5 L}{5!}x^5 \ . \tag{3}$$

The $K_n L$ are the integrated non-linear fields defined according to the MAD [7] convention. The small additional kick from the non-linearity will cause a small non-closure of the bump which will be visible as a transverse offset of the beam at the exit of the bump and has a magnitude

$$\begin{pmatrix} \Delta x \\ \Delta x' \end{pmatrix} = \begin{pmatrix} \varepsilon \\ 0 \end{pmatrix} = \begin{pmatrix} \sqrt{\beta_1\beta_2}\, f(\sqrt{\beta_1\beta_2}\Theta) \\ 0 \end{pmatrix} \ . \tag{4}$$

We will consider only very slowly varying bumps such that the bump's non-closure will not excite betatron oscillations but rather cause the closed-orbit to be perturbed. The magnitude of this closed orbit perturbation $(x_{\rm co}, x'_{\rm co})$ at the exit of the bump can be calculated from

$$\begin{pmatrix} \Delta x_{\rm co} \\ \Delta x'_{\rm co} \end{pmatrix} = R_e \begin{pmatrix} \Delta x_{\rm co} \\ \Delta x'_{\rm co} \end{pmatrix} + \begin{pmatrix} \varepsilon \\ 0 \end{pmatrix} \tag{5}$$

where R_e is the full turn transfer matrix starting at the exit of the bump. This equation we can solve for the closed orbit perturbation $(x_{\rm co}, x'_{\rm co})$ and obtain

$$\begin{pmatrix} \Delta x_{\rm co} \\ \Delta x'_{\rm co} \end{pmatrix} = (1 - R_e)^{-1} \begin{pmatrix} \varepsilon \\ 0 \end{pmatrix} \ . \tag{6}$$

The full turn matrix can be parameterized in the following way

$$R_e = \begin{pmatrix} \cos\mu + \alpha_1 \cos\mu & \beta_1 \sin\mu \\ -\frac{1+\alpha_1^2}{\beta_1}\sin\mu & \cos\mu - \alpha_1 \sin\mu \end{pmatrix} \tag{7}$$

where μ is the phase advance per turn related to the tune ν by $\mu = 2\pi\nu$ and α_1, β_1 are the usual Twiss parameters. After some matrix algebra we arrive at the result $x_{\rm co} = \frac{1}{2}(1 + \alpha_1 \cos \mu/2)\varepsilon$ where f and ε are defined in eq. 3 and 4, respectively. In order to simplify the algebra we will assume in the following that α_1 is zero and obtain

$$x_{\rm co} = \frac{1}{2}\sqrt{\beta_1\beta_2}\, f\left(\sqrt{\beta_1\beta_2}\Theta\right) \ . \tag{8}$$

We now assume that the two corrector magnets comprising the bump are excited sinusoidally, namely

$$\Theta(t) = \hat{\theta}\cos\omega t \ . \tag{9}$$

Inserting this into the equation for the closed orbit perturbation we see, that also the closed orbit is perturbed sinusoidally

$$\begin{aligned} x_{\rm co}(t) \ = \ & \frac{\sqrt{\beta_1\beta_2}}{2}\{\frac{K_2L}{2!}\left(\sqrt{\beta_1\beta_2}\hat{\theta}\cos\omega t\right)^2 + \frac{K_3L}{3!}\left(\sqrt{\beta_1\beta_2}\hat{\theta}\cos\omega t\right)^3 \\ & + \frac{K_4L}{4!}\left(\sqrt{\beta_1\beta_2}\hat{\theta}\cos\omega t\right)^4 + \frac{K_5L}{5!}\left(\sqrt{\beta_1\beta_2}\hat{\theta}\cos\omega t\right)^5\} \ . \end{aligned} \tag{10}$$

Introducing the maximum position at the non-linearity $\hat{x} = \sqrt{\beta_1\beta_2}\hat{\theta}$ and rewriting the powers of the cosine in terms of cosines of multiple arguments we get

$$\begin{aligned} x_{\rm co}(t) \ = \ & \frac{\sqrt{\beta_1\beta_2}}{2}\{\frac{K_2L}{2!}\frac{\hat{x}^2}{2}(1 + \cos(2\omega t)) \\ & + \frac{K_3L}{3!}\frac{\hat{x}^3}{4}(3\cos(\omega t) + \cos(3\omega t)) \\ & + \frac{K_4L}{4!}\frac{\hat{x}^4}{8}(\cos(4\omega t) + 4\cos(2\omega t) + 3) \\ & + \frac{K_5L}{5!}\frac{\hat{x}^5}{16}(\cos(5\omega t) + 5\cos(3\omega t) + 10\cos(\omega t))\} \ . \end{aligned} \tag{11}$$

Here we clearly see that the dodecapole term proportional to K_5L will create orbit oscillations at the odd harmonics of the driving frequency. It shows also that the unwanted fundamental is excited ten times stronger than the wanted fifth harmonic. We verified these observations and also that the magnitude of the respective peaks is correct in a simple one-dimensional tracking program and subsequent FFT analysis of the orbit data. Now we have to note that the orbit bump traverses two quadrupole triplets – one on either side of the IP. Simple analysis of the generating hamiltonians ($\propto x^6$) of the dodecapole components separated by 180^o phase advance shows that aberrations due to the dodecapoles will add

and will cause twice the excitation. We thus find that the amplitude of the fifth harmonic $\mathcal{A}(5)$ generated by driving an orbit bump through a non-linearity has the amplitude

$$\mathcal{A}(5) = 2 \times \frac{1}{32} \frac{K_5 L}{5!} \sqrt{\beta_1 \beta_2}\ \hat{x}^5 \ . \tag{12}$$

The magnitude of the $K_5 L$ can be estimated from the the fact that dodecapole error is on the order of $b_6 = 10^{-4}$ times the quadrupole's field at the reference radius of $R = 17\,\mathrm{mm}$. The coefficient b_6 is related to $K_5 L$ by [8]

$$K_n L = \frac{2\pi\, 5!}{1232}\ \frac{b_6}{R^5} \ . \tag{13}$$

Inserting numbers we find that nominal value $b_6 = 10^{-4}$ corresponds to about $K_5 L = 43000/\mathrm{m}^5$. Since we are interested in correcting this magnet error down to a fraction of its nominal value we use

$$K_5 L \approx 7200/\mathrm{m}^5 \ , \tag{14}$$

which corresponds to 1/6 of the nominal value as our reference value throughout the simulations. Inserting this $K_5 L$ into eq. 12 and assuming that the amplitude of the orbit bump is about 3 sigma, or 5 mm in the triplet quadrupoles we arrive at an amplitude of $\mathcal{A}(5) \approx 10 \times 10^{-9}\,\mathrm{m}$. If we assume that we can observe the oscillation in a triplet magnet at another IP where the beta function is large (4500 m) we can expect an amplitude approximately larger by a factor of $\sqrt{4500/150} \approx 5$ which results in an oscillation amplitude of the fifth harmonic of

$$\mathcal{A}(5) \approx 50 \times 10^{-9}\mathrm{m} \ . \tag{15}$$

This is clearly below the resolution and the digitization steps of the beam positioning system. Before we start discussing how to anyhow retrieve these weak signals from the beam we describe the simulation program that will help to verify the estimates and test the signal processing methods mentioned above.

III. SIMULATION

The simulation is based on a four-dimensional tracking program in which the beam is propagated by 4×4 transfer matrices between the dodecapoles which are modeled as thin lenses and the oscillating dipole corrector magnets. The layout is schematically shown in figure 1. The beam enters from the left and after propagating through some linear

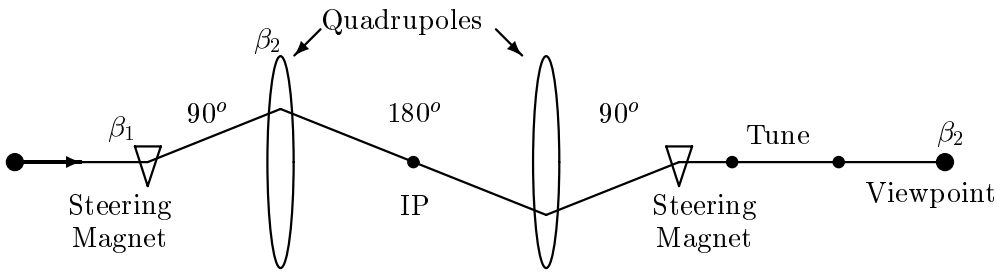

FIG. 1: The schematic layout of the involved elements. The beam propagates from the left to the right and the small triangles depict the wobbling steering magnets.

stretch of beam line encounters an wobbling steering magnet at a position with beta function $\beta_1 = 150\,\mathrm{m}$. Then it traverses a 90 degree section and meets the first triplet with beta function $\beta_2 = 4500\,\mathrm{m}$ that contains a dodecapole error. After the triplet quadrupoles the beam traverses the IP section. Since there is a very pronounced waist at the IP the phase advance between the triplets is very close to 180^o. The beta function at the outgoing triplet is also $\beta_2 = 4500\,\mathrm{m}$. After another 90^o section the wobbling steering magnet on the outgoing side is traversed. After the steering magnet a section is inserted that allows to adjust the tunes to the standard LHC values of $\nu_x = n.31$ and $\nu_y = n.28$. The final section is a piece of linear beam line that brings the beam to the viewpoint in another IP's triplet magnet which is both the starting and end-point of this small LHC model.

The sections with linear phase advance between two locations are modeled by a matrix containing the beta and alpha functions that maps real space into normalized phase space, the part that maps the horizontal coordinates is given by

$$\begin{pmatrix} 1/\sqrt{\beta_x} & 0 \\ -\alpha_x/\sqrt{\beta_x} & \sqrt{\beta_x} \end{pmatrix} \tag{16}$$

and a corresponding matrix for the vertical coordinates. The beta function in eq. 16 correspond to the first location. This matrix is followed by a rotation matrix containing the phase advances and a matrix that maps normalized phase space back into real space. The latter matrix is the inverse matrix of the one given in eq. 16 that contains, however, the beta and alpha function at the second location. The dodecapole is described by a single non-linear thin-lens kick. The wobbling steering magnets are modeled as horizontal kicks that are modulated by a very slow sinusoidal excitation at 0.912×10^{-4} times the revolution or sampling frequency that is 11.245 kHz in LHC which results in a wobbling frequency of

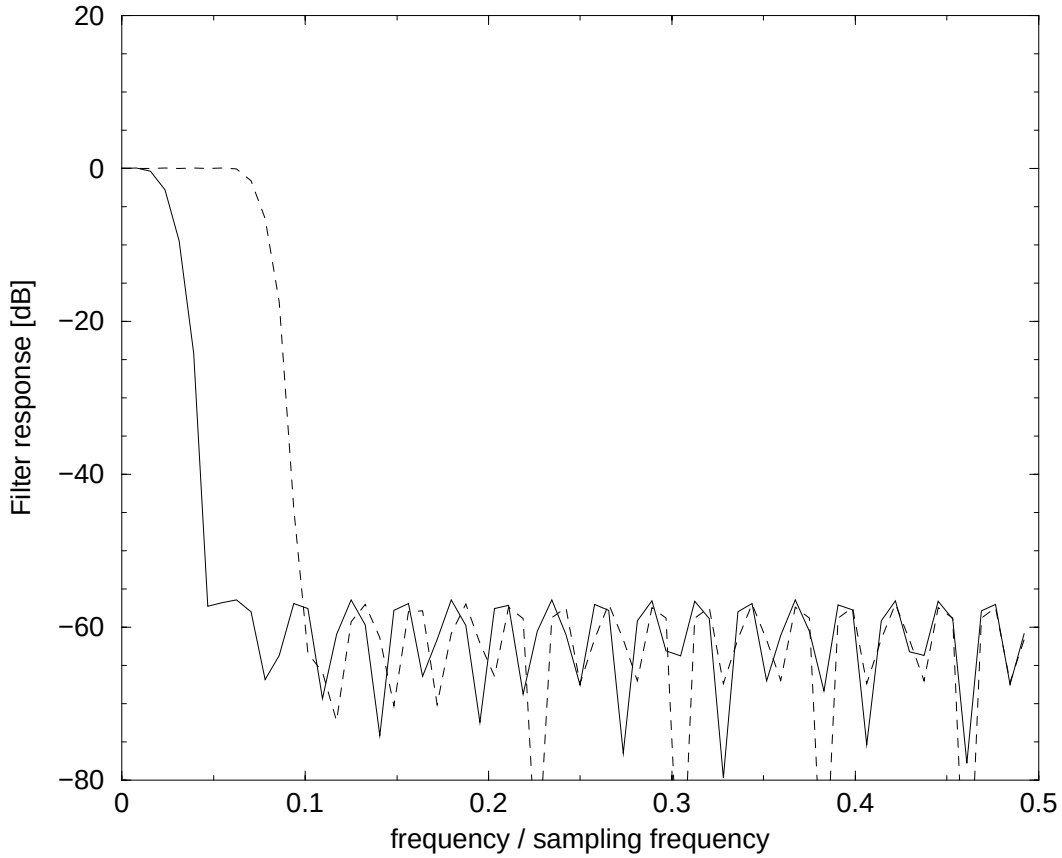

FIG. 2: The filters used in the decimation stages. The solid line shows the filter used in the by-10 decimation stage and the dashed line that one of the by-5 decimation stage. Note that the stopband attenuation is nearly 60 dB (factor 1000) for both filters.

about 1 Hz, which is the maximum frequency with which the wobbling steering magnets can be excited at collision energy. We are thus faced with the problem of accurately detecting a signal and its harmonics which takes on the order of 10 000 revolutions to perform a single oscillation. This task will require to track up to 10^9 revolutions, which takes about a full day of CPU time on a high end PC. The tracking program is written in FORTRAN and writes the beam's position at the viewpoint out to standard output. This output stream is then easily piped into other programs for further processing.

In order to test the simulation we ran about 10^6 turns and analyze the data to verify the estimates of the amplitudes given in the previous section. Since the oscillating frequency is on the order of 10^{-4} times the sampling frequency it is impossible to directly Fourier transform the entire data set and resolve the low frequencies with some accuracy. Instead we low pass filter the entire data set by a FIR-filter that cuts everything above 0.05 times

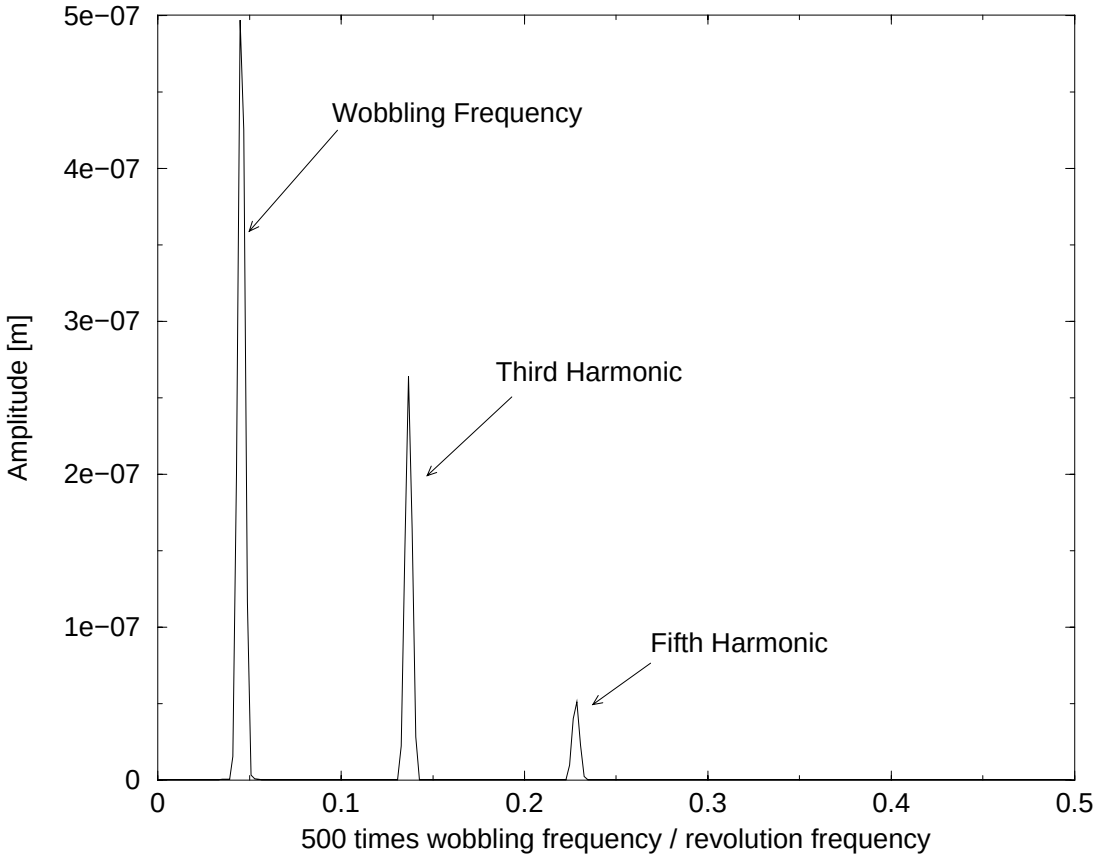

FIG. 3: The low frequency spectrum of the horizontal beam position sampled at the viewpoint after decimating the raw data stream by a factor of 500.

the sampling frequency and only keep every tenth sample, we thus decimate the low-pass filtered data by a factor 10. The filtering makes sure that no high frequency noise is aliased into the base band which contains the desired wobbling signals. This procedure expands the frequency axis near zero frequency by a factor 10. We apply two such by-10 decimation stages and one by-5 decimation stage. The filters used in the decimation stages are shown in Fig. 2. In total we thus have three decimation stages which expand the frequency axis near low frequencies by a factor 500. Fourier transforming the decimated signal then reveals the low frequencies present in the beam.

Figure 3 shows the results of the analysis of the horizontal beam position signal sampled at the viewpoint shown in Fig. 1. Note that the horizontal axis is stretched by a factor of 500 due to the decimation process. We clearly see three peaks at one, three, and five times the wobbling frequency which roughly have peak ratios of 10 to 5 to 1 as is expected from eq. 11. Moreover the peak amplitude of the fifth harmonic is about 5×10^{-8} m and agrees

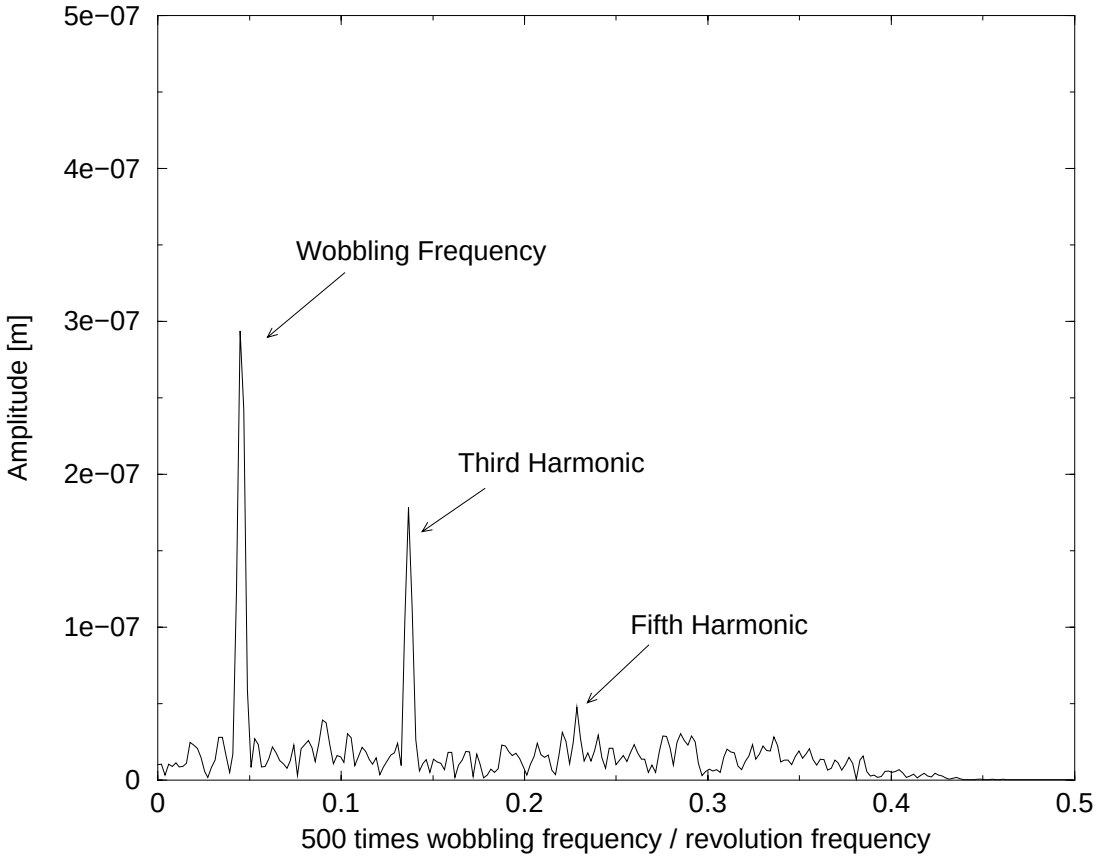

FIG. 4: The spectrum corresponding to Fig. 3 but with adding noise of 5 μm and digitizing at the 5 μm level.

well with the analytical prediction given in eq. 15.

Having found analytical estimates and the computer simulation in agreement but the magnitude of the signal still below the detection threshold of the beam positioning system we will now proceed to investigate ways to make these very weak signals visible.

IV. SIGNAL PROCESSING

The tracking simulation program writes the beam's horizontal position to standard output. In order to simulate noise we add a random number with gaussian distribution truncated a three sigmas and a sigma (rms) of 5×10^{-6}m to standard output of the noise-adding program and then quantize it with a granularity of 5×10^{-6}m. The resulting data stream is then passed through the same three decimating stages discussed above. The resulting spectrum is shown in Fig. 4 where we see that the peaks are now a little bit smaller than

those of the unperturbed system and they are placed on top of a noise floor which has a magnitude of about 2 to 3 times 10^{-8} m.

This noise floor is considerably less than the $5\,\mu$m noise we added initially. This considerable reduction comes from the fact that the original noise is spread evenly over the full spectral range from 0 Hz to the Nyquist frequency at half the sampling rate which is about 5.6 kHz in our case. By decimating by a factor 500 we thus remove a considerable amount of noise. In order to quantify the reduction we realize that low pass filtering is equivalent to averaging and that averaging reduces the noise level proportional to the inverse square root of the number of averages. We can thus expect a reduction of about $1/\sqrt{500} \approx 23$ through the decimation process. This estimate is probably a little bit too conservative, because the low-pass filter cuts at a lower frequency than 1/10 of the sampling frequency and therefore reduces the noise power by more than 1/10. A further reduction of the noise level can be expected from the Fourier transformation which calculates the spectral power in many bins of small bandwidth. The noise power per bin is then reduced in much the same way as discussed above and reduces the rms noise level by a factor $2/\sqrt{N}$ where N is the number of samples used to calculate the Fourier transformation. The derivation of this result is deferred to the appendix. In the calculations presented in this report we use $N = 512$ and can thus expect a reduction of $2/\sqrt{512} = 1/11$. The total reduction of the noise level due to decimation and Fourier transforming is then about 1/250. With an initial noise level of $5\,\mu$m and a reduction by 250 we expect the noise level to be on the order of 2×10^{-8} m which is consistent with what we observe in Fig. 4. In summary we can state that a reduction of the bandwidth that we observe is accompanied by a reduction of the noise floor, provided that the out-of-band noise is removed by filtering in order to avoid aliasing into the observation bandwidth.

Another feature that is clear from looking at Fig. 4 is that the fifth harmonic signal is barely elevated above the noise floor. We can, however, display the spectra as a function of time in the form of a spectrogram as can be seen in Fig. 5 where the the vertical axis from Fig. 4 is translated into a grey-scale value and higher values are displayed as darker spots. The horizontal axis covers the same frequency range as is shown in Fig. 4 and the vertical axis in corresponds to time as one spectra after another is translated into grey-scale and then displayed. Time runs from top to bottom and we have only printed every fifth spectrum in order to make the display easily comparable to those displayed later. The ticks

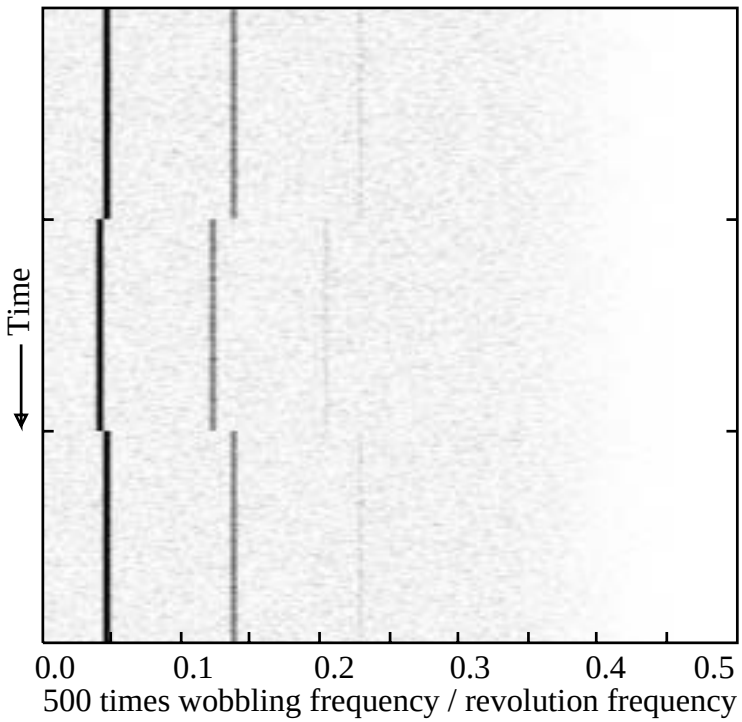

FIG. 5: Spectrogram of the fifth harmonic when the wobbling frequency was temporarily lowered. The horizontal axis corresponds to that in Fig. 4 and time increases from top to bottom and the total time corresponds to 3×10^8 turns.

on the vertical axis mark the time it takes to complete 10^8 turns, Also take notice that we show all spectrograms on a linear scale with auto-scale enabled such that the highest peak (which is shown in black) is always the maximum and the smallest value (shown in white) is zero. In Fig. 5 we can clearly discern the fundamental, third and fifth harmonic as vertical bands running from top to bottom, albeit the fifth harmonic is barely visible. During the center third of the observation time we have changed the wobbling frequency of the steering magnets from 0.912×10^{-4} times the revolution frequency to 0.812×10^{-4} and observe a step in the observed spectra of the wobbling harmonics. Note that the step scales with the harmonic number and therefore is five times bigger for the fifth harmonic visible as the right most band in Fig. 5.

As pointed out in the previous paragraph the fifth harmonic, which is the signature of the dodecapole component, is rather weak and we will show that we can zoom in on the fifth harmonic and decimate the frequency range again by a factor of 5 with the corresponding reduction of the noise floor. In order to achieve the frequency zoom we note that the

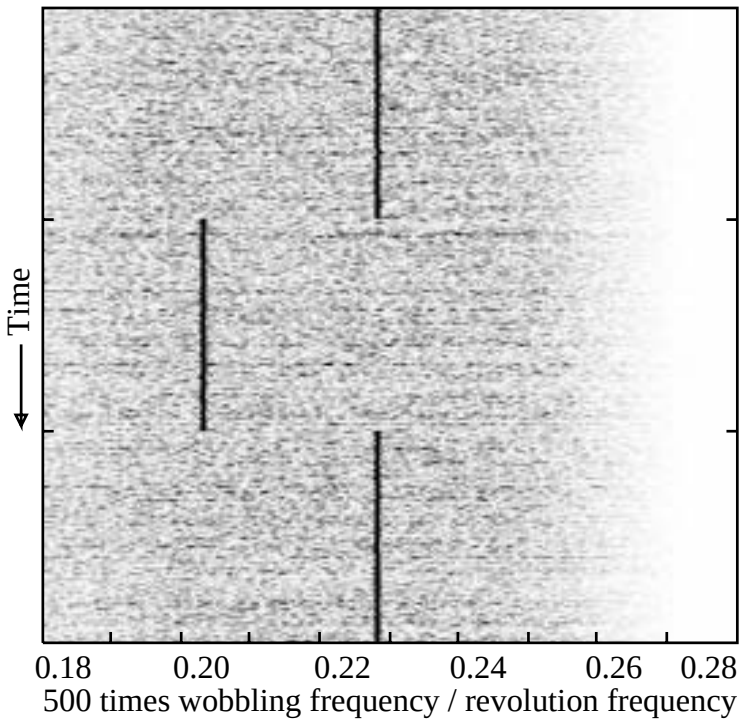

FIG. 6: Spectrogram of the fifth harmonic after using an image rejection mixer and a decimation factor of 5 using the same raw data that led to Fig. 5. The horizontal frequency scale now runs between 0.18 and 0.28.

frequency of the fifth wobbling harmonic lies between 0.20 and 0.23 after the by-500 decimation ($0.912 \times 10^{-4} \times 5 \times 500 = 0.228$) as can also be verified from Fig. 4. We will use an image rejection mixer and mix the data stream with a frequency of 0.18, then low-pass filter with the filter shown as dashed line in Fig. 2 which has a cut off frequency of 0.1. This will restrict the range to lie between 0.18 and 0.28. This range can then be expanded to span the full range of 0 to 0.5 by adding a by-5 decimation step. We use an image rejection mixer to avoid the noise from the lower sideband to alias into the observation range. For a discussion of the image rejection mixer see appendix B. The result of frequency zooming can be seen in Fig. 6 where the weak fifth harmonic and the frequency step is now clearly visible above the noise floor.

The signal to noise ratio of the fifth harmonic can be further improved by passing the signal through an adaptive line enhancer (ALE) [9] before displaying it as a spectrogram. The ALE, which is described more closely in appendix C, is an adaptive filter that very efficiently picks up coherent signals in a noisy background. We show the result of applying

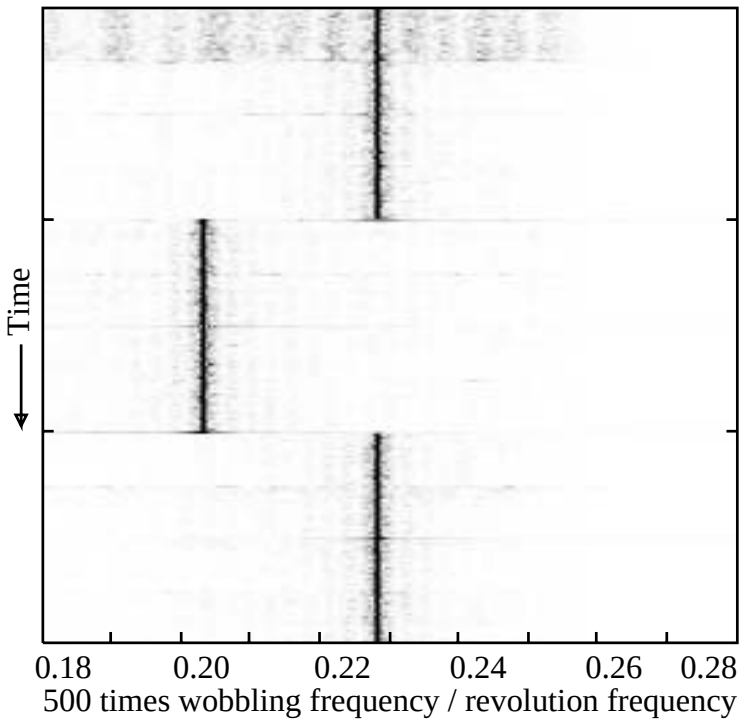

FIG. 7: Spectrogram of the fifth harmonic after image rejection mixer, decimation factor of 5, and adaptive line enhancer using the same raw data that led to Fig. 5. The scales are identical to that of Fig. 6.

the ALE after the image rejection mixer and obtain the spectrogram shown in Fig. 7 with considerably increased signal to noise ratio. We also observe equally spaced horizontal lines that are the product of re-initializing the ALE in order to ease finding the new frequency after the jump. This works because the convergence time of the adaptive filter is much shorter than the time between re-initializations.

We now turn to the investigation of what magnitude of dodecapole component will be visible using the wobbling method. To do this we run the simulation starting with the magnitude of $K_5 L = 7200\,\mathrm{m}^5$ and then progressively reduce it by a tenth of its magnitude such that the dodecapole component is turned off at the end of the simulation. We run 10^8 turns at each of the eleven different settings. Figure 8 shows the resulting spectrum after the by-500 decimation stage with the fundamental clearly visible on the left hand side until the last stage when the non-linearity is turned off. The third harmonic is visible almost as long but the desired fifth harmonic signal is barely visible as a weak band near the center of the plot. Note that we display the spectra in linear scale with auto-scaling turned on.

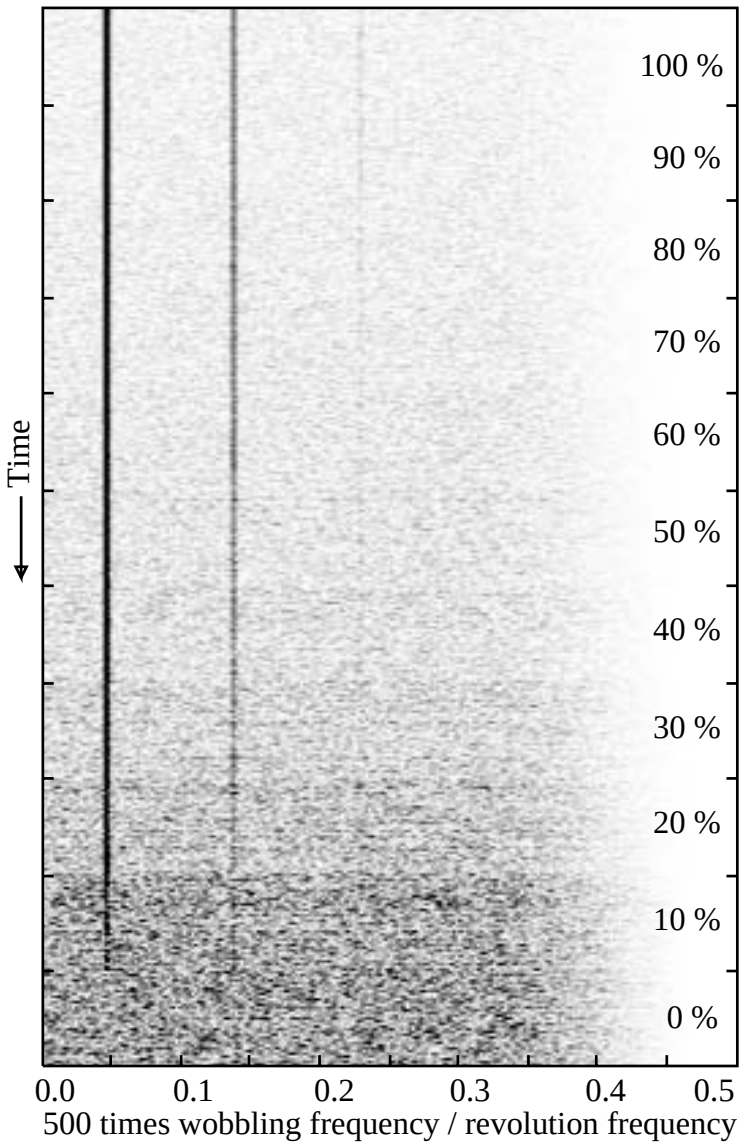

FIG. 8: A spectrogram after the by-500 decimation while the dodecapole excitation is reduced from reference excitation to zero in 10 steps. Since the display is linear with auto-scale the noise floor appears to grow on the signal until it is completely swamped. The time corresponds to 11×10^8 turns.

In order to see down to which dodecapole excitation the fifth harmonic is visible when using the image rejection mixer and the ALE we pass the same data that led to Fig. 8 through both signal processing tools and zoom in on the frequency range between 0.18 and 0.28. The result is shown in Fig. 9 where the ticks on the vertical axis mark the instances when the dodecapole component changes magnitude. We observe that the signal is picked up by the ALE down to dodecapole excitations of about a third of the reference value. Keeping

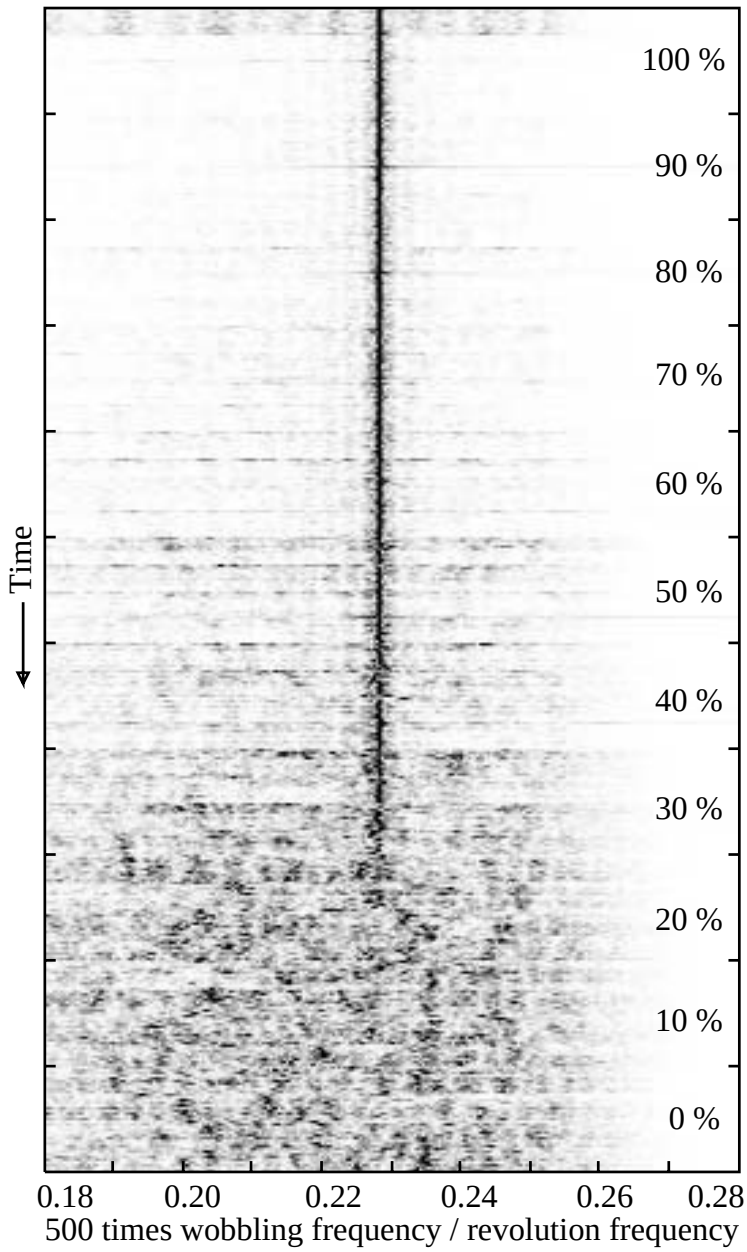

FIG. 9: Spectrogram of the same data as in Fig. 8 after zooming in on the fifth wobbling harmonic with the help of an image rejection mixer and an ALE. Note that the ALE is able to latch on the signal until the dodecapole component reaches about a third of the reference value.

in mind that the reference value corresponds to $b_6 = 1/6 \times 10^{-4}$ we find that under optimal conditions should make it possible to measure the correction down to the 5 to 10 % level of the nominal dodecapole error with $b_6 = 10^{-4}$.

V. IMPERFECTIONS

In the previous sections we consider the ideal system with exactly closed bumps and no other errors involved. These restrictions we will now drop and investigate the robustness of it with respect to parametric and other variations. To start with we look at errors in the wobbling steering magnet that lead to a non-closed bump.

In the simulation we simulated this effect by changing the excitation of the downstream steering magnet by up to 10 % and found that the peak of the fundamental increases dramatically and that also the wobbling harmonics increased in amplitude, while the noise floor stayed largely unchanged. Since the image rejection mixer efficiently suppresses the unwanted harmonics this does not create any additional problem. We found, however, that the fifth harmonic is actually slightly better visible than in the closed bump case. We conclude that the matching of the amplitude of the steering magnets and thus the closure of the bump will be un-critical.

A related effect is a mismatch in the phase advance inside the bump. If the bump is closed but the phase advance inside the bump is not distributed 90-180-90 degrees but 92-178-90 degrees no significant change of the spectrum to that one of the original layout is visible. We conclude that the detailed phase advances between the elements are un-critical, if the bump is closed.

If the dodecapoles are transversely misaligned we find that even wobbling harmonics, that are absent in the original spectrum in Fig. 3, appear. This will, however, not pose a problem, because the filtering can reliably extract the interesting fifth harmonic.

Having the bump not closed, even in the linear approximation, will cause non-linearities lying outside the bump to contribute to the wobbling spectrum. We do not expect this to be a problem, because of the dependence of the wobbling harmonics on the amplitude is of high power. A potential problem may come from the interaction of the chromaticity correction sextupoles which are strong and can generate dodecapole-like terms in third order. The magnitude of this effect needs to be investigated later more carefully in a more full-blown simulation.

Noise in the steering magnets will also cause a temporary non-closure of the bump and will mostly cause the wobbling fundamental to grow. This is, however, not a problem because the mixture of image rejection mixer and ALE will blank out these un-interesting harmonics

and zoom in on the fifth wobbling harmonic. On the other hand, steering magnet jitter will have to be carefully controlled in any case, because it will cause emittance growth. In the simulation we added a noisy steerer observed the full spectrum without decimation. The tune line was constantly growing until the beam's amplitude grew very large, at which point we aborted the simulation.

Another source of noise is transverse position jitter of the triplet quadrupoles due to ground vibrations that will affect beam's position. The steering error due to a quadrupole offset from the central axis by Δx is $\Delta x/f$ where f is the focal length of the quadrupole and is on the order of 20 m in our case. If we neglect the phase advance between the jittering quadrupole and the position where we observe the beam position we can estimate the position jittering somewhat pessimistically to be $\beta\Delta x/f$. Inserting $\beta = 4500$ m and $\Delta x = 0.1\,\mu$m we obtain a position jittering of $22\,\mu$m. This is a little bit larger than the $5\,\mu$m assumed above for the noise that is actually needed to elevate the weak wobbling harmonics above the digitization threshold. We conclude that weak quadrupole position jitter somewhat below the $0.1\,\mu$m level may be advantageous.

Other multipoles inside the closed bump will generate mixing products of the wobbling frequencies that will contribute to the observable low frequency spectrum. The dominating contributions will come from the triplet quadrupoles, because there the beta functions are so large. Its is easy to see that the odd multipole components (sextupole, decapole, ...) of the triplet quadrupoles will largely cancel because the phase advance of nearly 180 degrees between the triplets. The octupole component with its x^3-like kick will excite third order wobbling harmonics which are clearly discernible from fifth order harmonics generated by the dodecapole. The higher multipoles, beyond the dodecapole, should be rather weak. Furthermore, the dependence of the wobbling spectrum on the bump amplitude is even more disadvantageous than that of the dodecapole.

In this report we only investigate the effect of two localized dodecapoles and neglected the correction dodecapoles that are foreseen in the interaction regions for the cancellation of the dodecapole components of the triplet quadrupoles. If, on the other hand, the phase advance between the triplet and the correction magnets is small, this non-locality will have a small impact. The system of triplet and correction magnets will almost act as a single dodecapole and the discussion of the magnitude of the dodecapole from the previous section remains valid and implies that we will be able to detect dodecapole components on the order

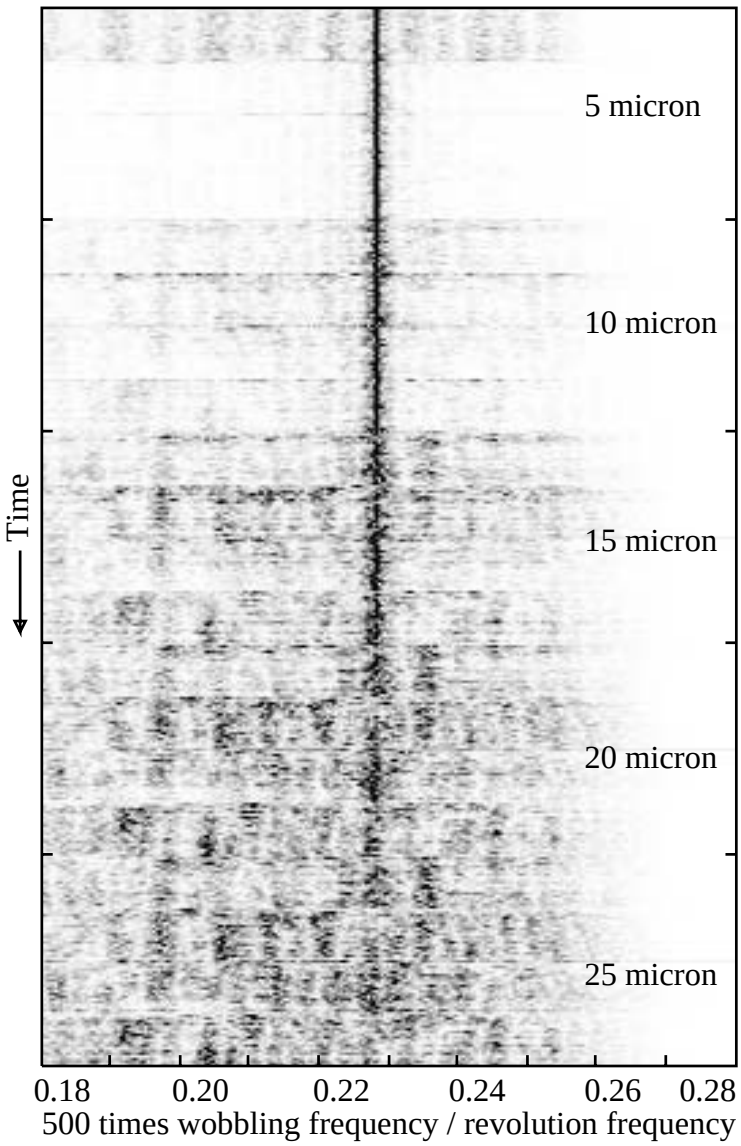

FIG. 10: Spectrogram of the fifth harmonic while increasing the added noise from 5 μm to 25 μm. The total time corresponds to 5×10^8 turns.

of the 5 to 10 % level of the nominal errors.

As we have pointed out before, noise is needed to elevate the wobbling signals above the digitization threshold. So, no noise is bad, a little noise helps, and too much is obviously also bad, because it swamps the wobbling harmonics. In order to get some feeling for the beneficial noise range we run simulations with noise levels of $0, 5, 10, 15, 20, 25\,\mu$m rms noise level (truncated at 3 sigmas) and observe how well the final fifth wobbling harmonic can be picked up. We find that without noise no signal is visible. The results of the simulation with non-zero noise level are shown in Fig. 10 where we increase the noise from 5 to 25 μm

in 5 steps of 5 μm from top to bottom. We see that the signal to noise ratio progressively deteriorates and that the ALE is able to reliably pick up the signal up to about halfway down Fig. 10 which corresponds to 15 μm. We thus need a "noise window" in which the presented method works between about 5 and 15 μm.

As a final and rather crucial point is *how* to apply noise to the measurements. We saw that applying noise directly to the beam will cause the emittance to grow and is thus not advisable. The noise must, however, be applied before the analog-to-digital converters (ADC) employed in the beam positioning system. A detailed investigation of how this can be done is beyond the scope of this report but needs to be addressed later.

VI. CONCLUSIONS

We present a method to detect the dodecapole errors from the triplet quadrupole magnets in LHC at collision energy. The method is based on driving a slow oscillating or wobbling closed orbit bump through the interaction region. The non-linearities cause the bump to be non-closed and we can observe their effect on the orbit outside the bump by investigating the low-frequency Fourier spectrum which shows harmonics of the wobbling frequency. We estimate the amplitude of these harmonics to be around 50×10^{-9} m in a simple analytical model and verified this estimate in a simulation program. Since the amplitude of the spectrum is below the digitization threshold we add noise to the signal to lift it above the threshold where it can be digitized. The added noise is spread out over the spectral range of up to half the sampling or revolution frequency while the wobbling signals only occupy a very small frequency band at low frequency where careful filtering can make them detectable. We show that a combination of filters, decimators, image rejection mixers and adaptive line enhancers makes it possible to detect signals coming from dodecapole components of a magnitude down to 5 to 10 % times the nominal errors expected for LHC. We then briefly discussed limitations of the method which need to be addressed further.

A point of rather principal nature is the question of how to actually feed the noise into the system. A detailed investigation of the beam positioning system is required to settle this point. Another point that requires further attention are simulations of a more complete, and more realistic model that goes beyond the simple toy model used in this report.

Finally we wish to thank O. Brüning and F. Ruggiero, CERN for an invitation to CERN,

the suggestion to work out a method to determine the dodecapole components of LHC's triplet quadrupoles and interesting discussions. O. Brüning's valuable comments are highly appreciated. Discussions with T. Lofnes, TSL regarding signal processing methods were very helpful. The permission of TSL's directorate to spend time on this research is gratefully acknowledged.

---

[1] The LHC Study group, "The Large Hadron Collider", CERN/AC/95-05 (LHC), 1995.

[2] Frank Schmidt at the 18th LCC meeting, 13th February, 2002.

[3] V. Ziemann, "Measuring Hamiltonian Coefficients with a Wobbling Method," Particle Accelerators **55,** 419, 1996.

[4] J. Bosser, C. Bovet, C. Fischer, R. Jung, H. Koziol, H. Schmickler, L. Vos, "LHC beam instrumentation conceptual design report," CERN LHC Project Report 370, February 2000.

[5] E. Ifeachor, B. Jervis, "Digital Signal Processing," Addison Wesley, 1993.

[6] V. Ziemann, "F_Tools, Signal Processing on the Command Line", TSL-Note, 2002/53, April 2002, unpublished.

[7] H. Grote, C. Iselin, "The Mad Program, Version 8.19", CERN/SL/90-13(AP) (Rev. 5), 1996.

[8] O. Brüning, S. Fartoukh, "Field Quality Specification for the LHC Main Dipole Magnets", LHC-Project-Report-501, October 2001.

[9] J. Zeidler, "Performance Analysis of LMS Adaptive Prediction Filters", Proceedings of the IEEE, Vol. 78, No. 12, p. 1781, 1990.

[10] The American Radio Relay League, D. Reed, ed., "The ARRL Handbook for Radio Amateurs 2002", 79th edition, 2001.

APPENDIX A: SPECTRAL NOISE FLOOR

As discussed in the main body of this report the very weak harmonic signal can be detected because the noise floor of the added random noise is rather low. In this appendix we derive the magnitude of the noise floor that appears when Fourier transforming random numbers x_k with rms magnitude σ. We assume that the random numbers are completely

uncorrelated and have mean zero. They obey the following statistics

$$< x_k x_l >= \sigma^2 \delta_{kl} \tag{A1}$$

where the angle brackets denote ensemble average over many realizations of the random numbers and δ_{kl} is the usual Kronecker delta. The spectral power s_k of harmonic k of the random sequence is given by the Fourier sum

$$s_k = \frac{2}{N} \sum_{k=1}^{N} x_k e^{2\pi i k/N} \ . \tag{A2}$$

Calculating the square of the absolute value of the power and taking the ensemble average we get

$$\begin{aligned} < s_k {s_l}^* > &= \left(\frac{2}{N}\right)^2 \sum_{k=1}^{N} \sum_{l=1}^{N} \sigma^2 \delta_{kl} e^{2\pi i (k-l)/N} \\ &= \left(\frac{2}{N}\right)^2 \sum_{k=1}^{N} \sigma^2 = \frac{4}{N} \sigma^2 \end{aligned} \tag{A3}$$

where we have used the definition of the statistics of the random numbers given in eq. A1. The amplitude of the noise floor is then given by the relation stated in the main body of the text.

APPENDIX B: IMAGE REJECTION MIXER

An image rejection mixer is well-known from RF design and is also known as a single sideband mixer [10]. It addresses the problem that both lower and upper side bands are mixed into the base band when mixing a RF signal with a local oscillator. To visualize the problem assume a signal around 1.3 MHz. In order to observe it in base-band near zero frequency we mix it with a 1.2 MHz local oscillator signal and obtain a 2.5 MHz signal and a 0.1 MHz signal. The former is easily removed by low-pass filtering. The 0.1 MHz signal is the wanted base-band signal. Unfortunately, that is not the only signal that is mixed to 0.1 Mhz. If there is some signal at 1.1 MHz, for example noise or some other perturbation, originally present, it will also mix to 0.1 MHz and will show up at base-band. An image rejection mixer is a device that will mix only frequencies higher (or lower) than the local oscillator to base-band.

To understand the operation of an image rejection mixer we consider the basic schematic diagram shown in Fig. 11. The sampled input signal enters from the left and is split in an

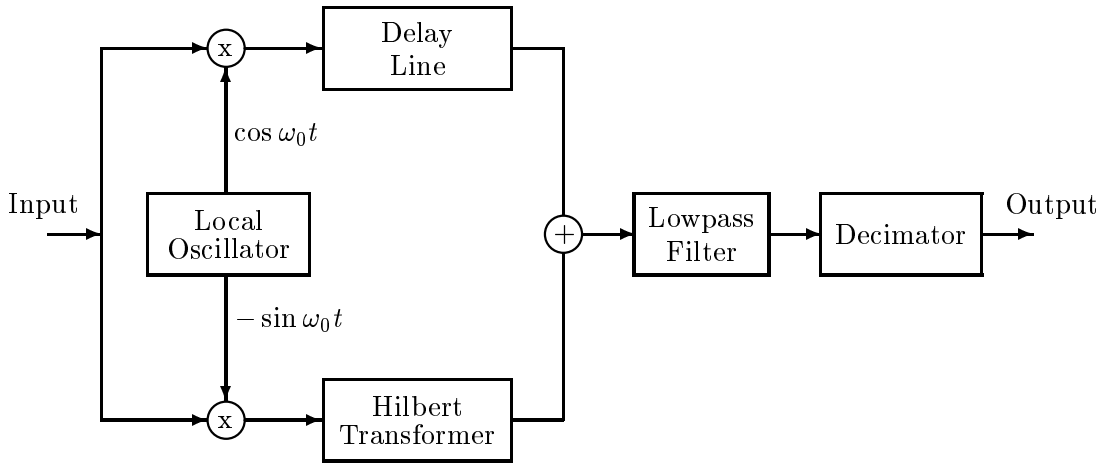

FIG. 11: Schematic block diagram of the image rejection mixer used in the data analysis.

upper and a lower branch which are mixed or multiplied by a cosine or sine signal from a local oscillator, respectively. This results in two signals that are 90 degrees out of phase with each other. The signal in the lower branch is passed through a Hilbert transformer which causes a 90 degree phase shift for all frequencies and a delay [10]. Only the delay is applied to the signal in the upper branch. It can be shown that the two 90 degree phase shifts that the lower branch is subjected to compared to the signal in the upper branch makes sure that the lower side band is canceled once the two signals are summed up. The signal with only the upper sideband is then low-pass filtered and optionally decimated. The schematic shown in Fig. 11 is implemented in one of the programs of the f_tools [6] suite.

APPENDIX C: ADAPTIVE LINE ENHANCER

An ALE is a signal processing algorithm that helps to increase the signal to noise ratio of coherent signals embedded in a noisy background such as weak RF signals used in cellular phones. It's working principle can be easily explained by discussing the schematic block diagram shown in Fig. 12. The noisy sampled data stream enters at the lower left corner where it is split in two parts. One half is directed to a delay line that is typically a few samples long and is split again. One half is passed through a FIR filter whose coefficients are determined by a LMS algorithm [5] and the other half is directly fed to the LMS algorithm. The other input to the LMS algorithm is e the difference signal between the original data stream and the delayed and filtered data stream. The FIR filter weights w are then adjusted

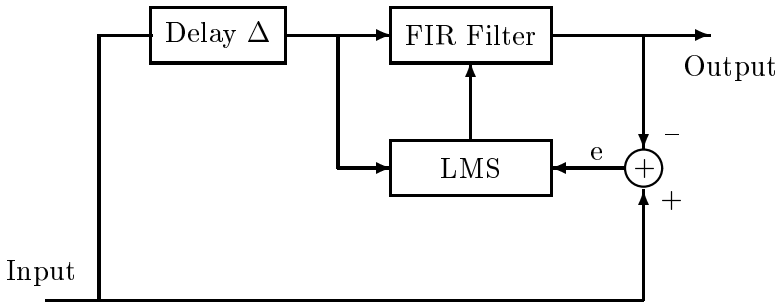

FIG. 12: Schematic block diagram of an adaptive line enhancer.

in the LMS according to

$$w_l(k+1) = w_l(k) + 2\mu e(k)x(k-l-\Delta) \tag{C1}$$

where w_l denotes the $l-$th filter coefficient, k labels the data sample just processed, Δ is the delay and μ is the LMS learning parameter [5].

So, why does this algorithm produces reduced noise or enhanced coherent signals at the output? To understand this consider for the moment a simple system in which only a noisy signal – assumed uncorrelated – is delayed and subtracted from the original data stream. Since we assume that the noise is uncorrelated this difference signal will be zero on average. This situation corresponds to Fig. 12 without the FIR and LMS box. But even with the FIR the noise will average out at the output. So, all the FIR and LMS block have to work on is building a filter that minimizes coherent signals in the error signal e. But this is done by building a FIR filter that picks up the coherent part. Since we are not using e as the output signal but the output of the FIR filter that is made to enhance the coherent part of the signal we get just that – the coherent part of the signal.

APPENDIX D: USING F_TOOLS

In order to illustrate how the data analysis is performed we show the f_tools commands that generate the plots shown in the main body of this report. A more detailed description of the f_tools program suite can be found in Ref. [6].

The original data stream is generated by a FORTRAN program named `b6` that prints the beam position out to standard output. It is interfaced to the f_tools by UNIX-pipes.

The command line, that was actually put in a script file, that led to Fig. 4 is shown here

```
b6 B6.LAT 1e8\
     | f_noise -am 5e-6 -seed 3\
     | f_digitize -am 5e-6\
     | f_dec -ra 10 -ff lp-0.05.dat\
     | f_dec -ra 10 -ff lp-0.05.dat\
     | f_dec -ra 5 -ff lp-0.1.dat\
 > data.dat
```

where `B6.LAT` is the input file that contains the beam line definition and `1e8` is the number of turns that are to be tracked. The backslash command at the end of the line denotes a continuation line under UNIX. Note the pipe symbol `|` at the beginning of every line that means that the output of the program on the previous line is piped to standard input of the program following the pipe symbol. The first program of the original data stream that `b6` generates is `f_noise` which adds noise of an amplitude 5×10^{-6} to the data stream followed by the digitizer `f_digitize` that quantizes the data stream in steps of 5×10^{-6}. The three next lines starting with `f_dec` perform the decimation stages with rates of two times 10 and 5 once. The parameter following the `-ff` command line parameter is the filter file to be used before decimation. Finally the output of the last decimation stage is redirected by the UNIX redirector (`>`) into a file named `data.dat` which can then be further processed and inspected later.

The filters used in the decimation stages were generated by a call to `f_remez_lowpass` which generates a Parks-McClellan-Remez optimal low-pass filter with the parameters specified on the command line.

```
f_remez_lowpass -N 91 -bw 0.06 -df 0.035 > lp-0.1.dat
```

Here we specify that the filter is 91 coefficients long, has a cutoff frequency of 0.06 times the sampling rate and a transition band width of 0.035 times the sampling rate. The calculated filter coefficients are then just redirected into the file `lp-0.1.dat`.

The spectrum of the data stream can be inspected with the `f_scope` program. The result of the command line

```
cat data.dat | f_scope --
```

leads to Fig. 4. The `--` command line option specifies that all the default spectrum analyzer configuration with Blackman windowed FFT is to be used. The `f_scope` program generates both the instantaneous spectrum (or waveform) display and a spectrogram that continuously displays the most recent spectra in a grey-scale spectrogram as shown for example in Fig. 5.

The following f_tools command generates the spectrogram shown in Fig. 6

```
cat data.dat | f_irm -fr 0.18 -dec 5 -ff lp-0.1.dat\
             | f_scope --
```

which rather similar to the previous example, except that the image rejection mixer `f_irm` is sandwiched between the data stream in file `data.dat` and the spectrum analyzer `f_scope`. `f_irm` takes the frequency (`-fr 0.18`), the decimation rate (`-dec 5`), and the low-pass filter file (`-ff lp-0.1.dat`) as command line arguments. The image rejection mixer that is part of the f_tools suite uses a Hilbert transformer with a filter length of 129 taps. The resulting rejection efficiency of the lower side-band is around 60 dB.

The inclusion of the ALE that is shown in Fig. 7 is the result of

```
cat data.dat | f_irm -fr 0.18 -dec 5 -ff lp-0.1.dat\
             | f_ale -nf 71 -de 3 -init 1e4\
             | f_scope --
```

where we use the same image rejection mixing and decimation as in the previous example, but add a step in which the `f_ale` program uses a delay $\Delta = 3$ samples to dynamically generate a FIR filter with 71 taps. The `-init 1e4` command line argument causes `f_ale` to initialize all working arrays after having read 10 000 samples which leads to the horizontal stripes in Fig. 7, but is necessary if one works with varying frequencies like the step shown in Fig. 7. If running without re-initializing the arrays, the ALE has difficulties latching on to the new frequency once a filter for the initial frequency is built.