

Chapter 3

Statistical analysis of RHIC beam position monitors performance

3.1 Introduction

BPMs are used in accelerators to record the average orbit and transverse turn-by-turn displacements of the bunch centroid motion. RHIC consists of two six-fold symmetric rings with six interaction regions. There are 160 BPMs per plane per ring (yellow & blue): 72 dual-plane BPMs distributed through the IRs, and 176 single-plane BPMs distributed in the arcs [35]. Each BPM channel acquires 1024 consecutive turn-by-turn positions.

It is imperative to understand the functioning of each BPM to obtain reliable and consistent data for any beam dynamics analysis. In this chapter we perform a first detailed statistical analysis to evaluate the performance of the RHIC BPM system with the aid of numerical tools. A large set of data files recorded after applying a horizontal kick during RUN 2002-03 were analyzed using SVD and FFT techniques [25, 36]. Both techniques were independently employed to identify malfunctioning BPMs from available data sets. Statistical behavior of BPM performance was extracted to characterize each pick-up and make a comparison between the two methods. A comparison between two different operation years is also presented and improvements in the BPM system is evident.

3.2 FFT TECHNIQUE

The Fourier spectrum of turn-by-turn data has already been used to determine faulty BPM signals from the Super Proton Synchrotron of CERN [36]. This technique relies on the fact that the Fourier spectrum of an ideal signal has well localized peaks while noisy or faulty signals show a randomly popu-

lated Fourier spectrum. A priori two observables seem to provide information on this regard: the average and the rms of the background of the Fourier spectrum. The average background depends on other parameters apart from the noise of the signal, therefore it is discarded for our purposes. The rms of the background of the Fourier spectrum is larger for noisier BPMs and its dependence with other parameters is negligible. This observable is therefore used to identify the noisy pick-ups. It is estimated by computing the rms of the amplitudes of the spectral lines within one or several spectral windows. These windows are chosen in such a way that an ideal pick-up would not show any peak within them. It is important to avoid including the zero frequency and the betatron tunes in any window. An illustration of a typical configuration for RHIC is shown in Fig. 3.1, using a good signal.

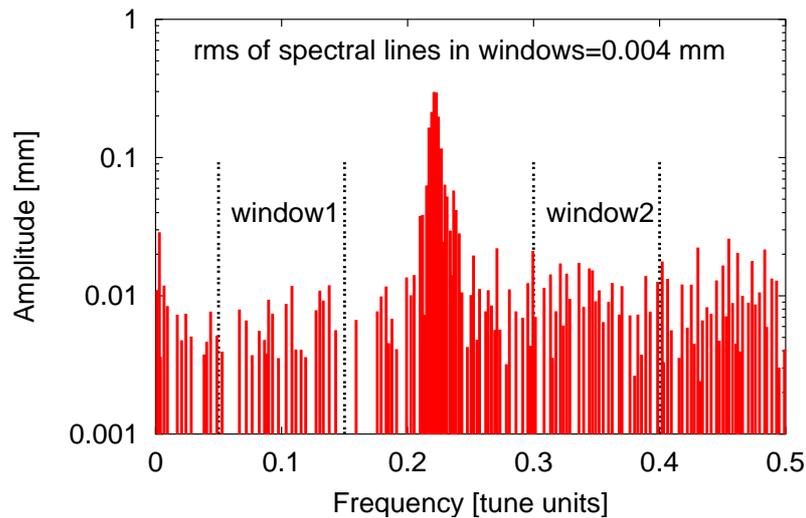


Figure 3.1: Fourier spectrum of a good RHIC BPM signal. The spectral lines within the windows are used to determine the rms noise observable.

A signal is considered faulty if its rms noise observable is larger than a certain threshold. The value of the threshold is extracted from statistics over a large number of signals. A histogram of rms observables from all signals is constructed. Typically a large peak containing the largest percentage of the data is observed in the low rms values of the histogram. This peak contains the set of physical signals, while its long tail with larger rms values contains the faulty signals. A suitable threshold is chosen towards the end of the tail. It will be shown in section 3.4 that particular choices of the cut do not give qualitatively different results.

3.3 SVD Technique

SVD has proved to be a powerful numerical tool in a wide variety of disciplines and has been recently applied to beam dynamics under the name “Model Independent Analysis” [25]. SVD is used to identify the principle components by maximizing the cross-covariance between time-dependent data. The number of non-zero singular values gives the effective rank of the matrix. For a matrix containing turn-by-turn data from several BPMs, the dominant singular vectors represent the temporal and spatial variation of physical modes characterizing the motion of the beam.

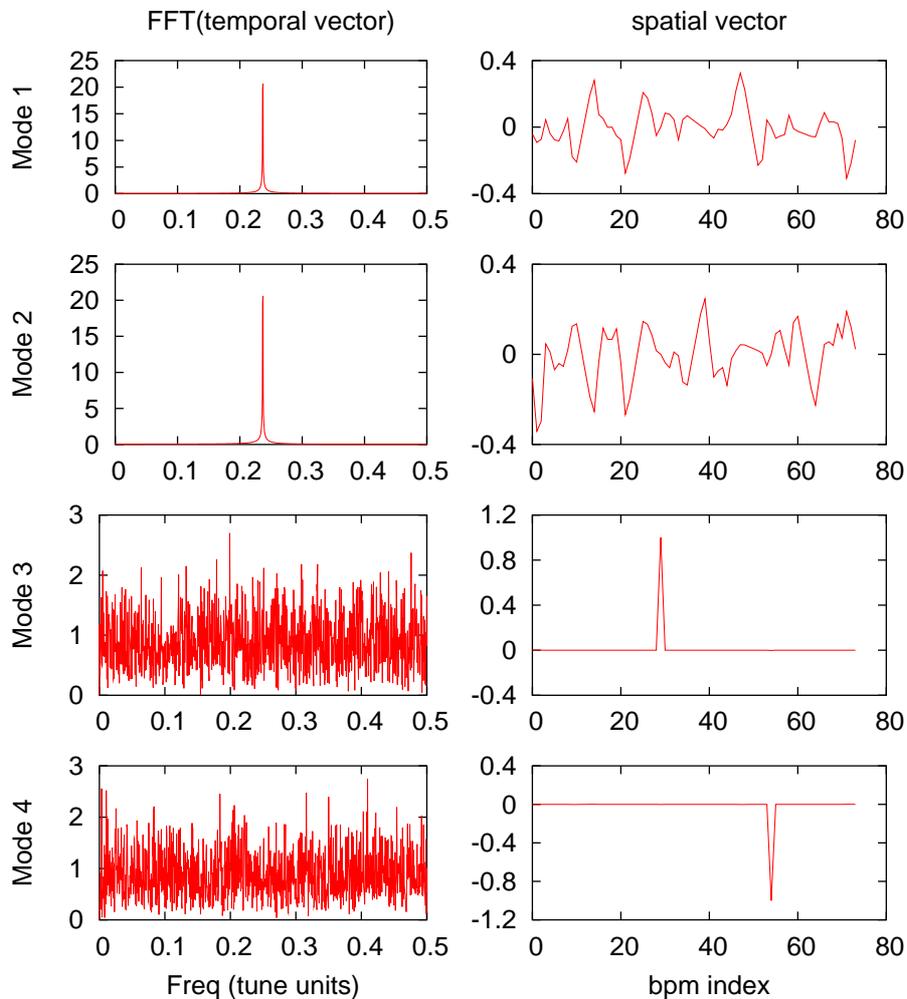


Figure 3.2: Spatial vectors and FFT of temporal vectors of the 4 dominant modes from simulation data. Note that singular vectors are normalized.

The potential scope of SVD analysis in accelerators runs deep. This chapter attempts to exploit one facet that aids in identifying malfunctioning BPMs. Fig. 3.2 shows typical plots from the application of SVD to simulated BPM data (after subtracting the average orbit), with only linear elements and two noisy BPMs at arbitrary locations. The singular value spectrum shows two dominant modes (“Mode 1” and “Mode 2”) corresponding to the betatron oscillation in the plane of observation. The Fourier transform of the temporal mode yields the betatron tune. “Mode 3” and “Mode 4” show sharp spikes in their spatial vectors. The signal manifested in this mode is localized to a particular BPM location indicating a potentially noisy BPM. The Fourier transform of their temporal modes yields all frequencies, confirming that these are noisy BPMs. In this situation of relatively ideal conditions, we find distinct peaks localized at corresponding noisy BPM locations and a flat signal elsewhere. However, from real data, we observe multiple peaks in the spatial vectors, due to random correlations between the noisy BPMs. One such possibility was simulated by setting the noise amplitudes in the two pickups to be approximately equal. We observe two dominant modes as above, but the spatial vectors now contain two peaks in each mode, as shown in Fig. 3.3.

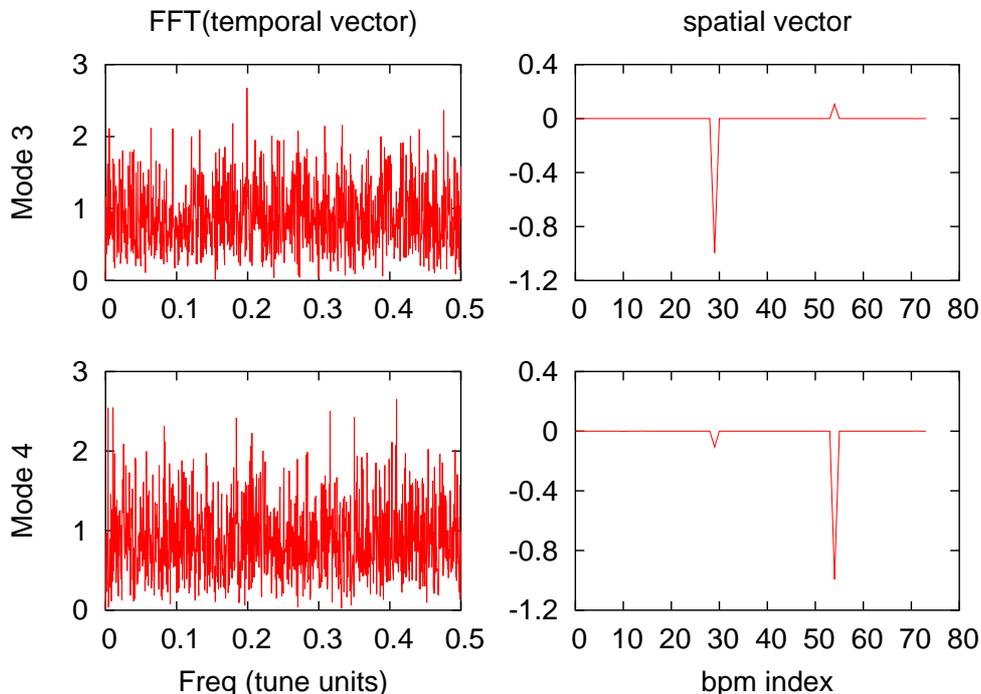


Figure 3.3: Spatial and temporal vectors of modes corresponding to two noisy BPMs with correlation.

A simple approach to identify faulty BPMs would entail finding modes with spatial vectors consisting only of localized peaks. A localized peak threshold value of 0.7 or greater might be sufficient to tag them as noisy BPMs [37]. However, we explore a statistical approach to understand the characteristics of BPM signals from which we choose the appropriate thresholds. A histogram of the largest peaks in each spatial mode for a large set of data is constructed to determine these thresholds. It will also be shown from statistics that an alternate approach using the norm of n largest peaks in each mode is a more accurate procedure to determine threshold values. One should note that the SVD method is insensitive to flat signals. Some preprocessing of BPM data, using peak-to-peak signal information, will be effective before applying such techniques. This is discussed below.

3.4 Analysis

About 2000 data sets (1000 for each ring) taken during Run 2003 of deuteron-gold collisions were used in this statistical approach. Statistical cuts were applied during data preprocessing before using the two techniques. This data is then used to determine independent thresholds for each technique, to identify noisy BPMs. We discuss each statistical cut in detail to explore the advantages and limitations of such an approach.

3.4.1 Hardware Cut

RHIC BPM hardware internally determines the status for each turn-by-turn measurement in a data set. Status bit information acquired in this way allows us to identify obvious hardware failures. We make a simple histogram of all BPMs that fail (status = 0), as shown in Fig. 3.4. BPMs failing this cut are removed from the data and are not included in further analysis. This histogram also helps us identify any consistent hardware problems. We also find a number of files in which fewer than 30 BPMs were present, and these files were excluded for this analysis separately for each plane.

3.4.2 Peak-to-Peak Thresholds

A peak-to-peak cut is necessary because both techniques become less sensitive as signal oscillations become small. Peak-to-peak signal values for all BPMs in all data sets were calculated, and plotted in a histogram in Fig. 3.5. One has to be careful to choose an appropriate cut. If optics and machine

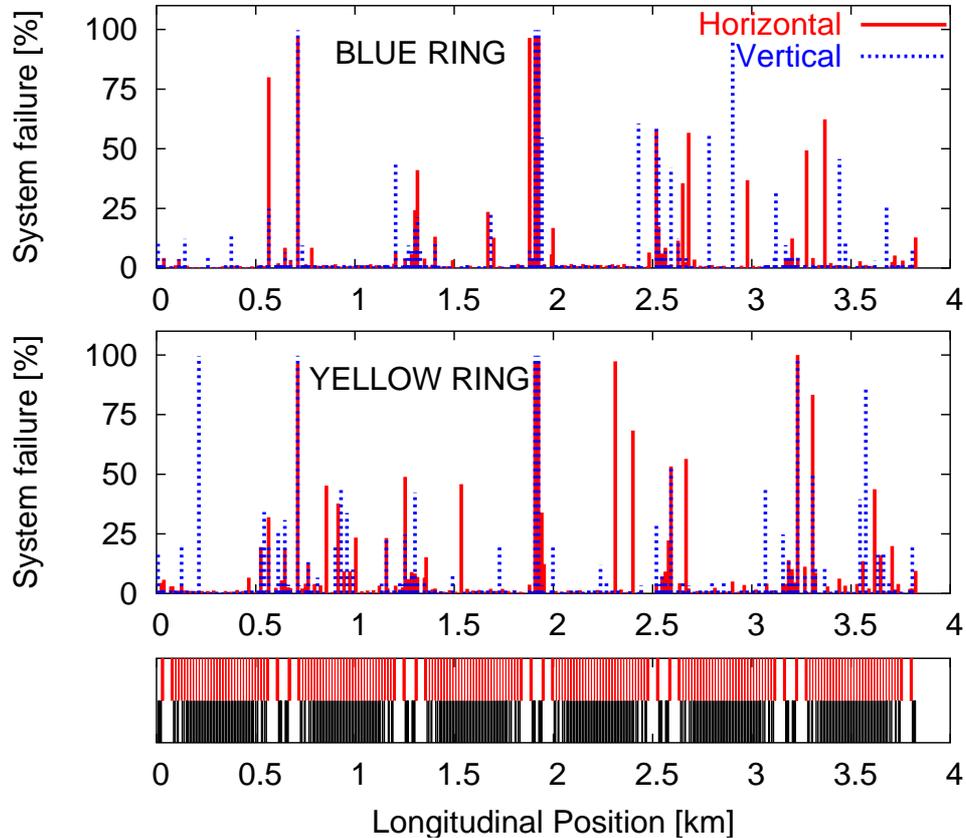


Figure 3.4: Percentage of occurrences of system failure per BPM versus longitudinal location. A representation of the lattice (dipoles in black and quadrupoles in red) is shown in the bottom graph.

conditions were static, it would be easier to determine the best threshold. However, this is not true in regular operation. A very low cut might help retain good BPMs with small signal, but it will also retain BPMs that are faulty. A large cut removes faulty BPMs, but will also identify many good BPMs as faulty.

The data sets being analyzed are mostly horizontally kicked data, hence we will discuss horizontal plane features in detail. There is a distinct minimum at 0.3 mm and a peak below that value which we believe are BPMs that do not respond to beam current. If a 0.3 mm cut removes good BPMs for a specific set of unexcited data, they will appear as background in the final identification, and will not be tagged as faulty BPMs. We observe a large peak around 1 mm

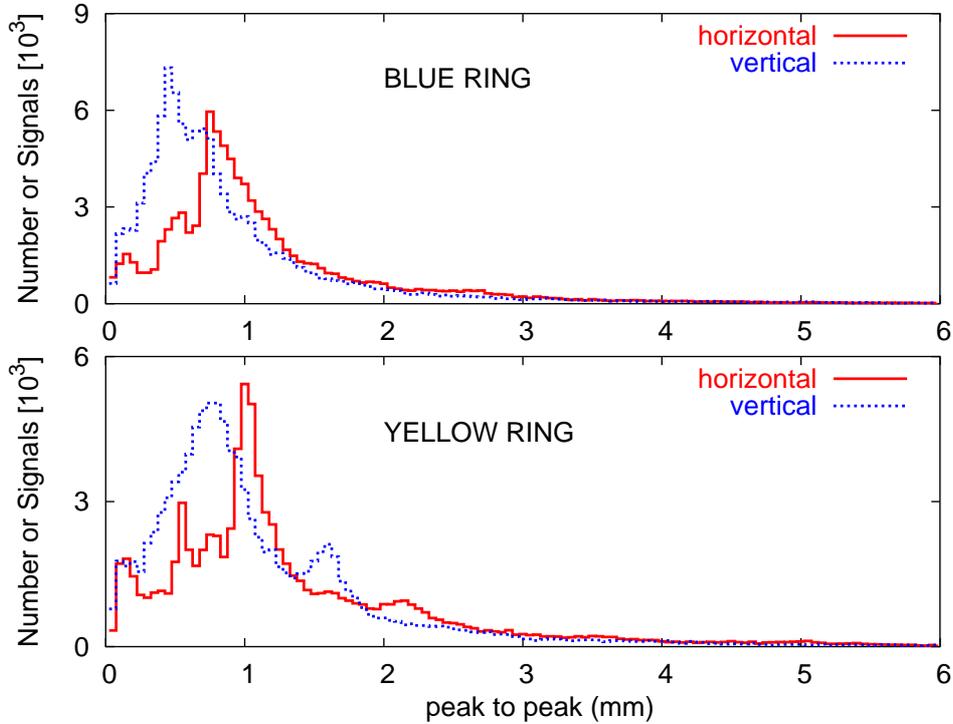


Figure 3.5: Peak-to-peak values for all BPMs in all data files.

indicating the typical oscillation amplitude for most pick-ups. Two smaller peaks are also observed between 1.5 mm and 2.2 mm, and two more between 0.5 mm and 1 mm, possibly indicating signals from different sets of optics or different kick amplitudes.

The vertical plane signals appear at smaller amplitudes because the data analyzed are mostly kicked in the horizontal plane. In such cases smaller peak-to-peak thresholds must be chosen. A summary of thresholds and files analyzed are given in Table 3.1.

Table 3.1: Thresholds for peak-to-peak values

Plane	Ring	peak-to-peak	# of files analyzed
H	yellow	0.3	678
V	yellow	0.15	815
H	blue	0.3	708
V	blue	0.15	833

3.4.3 FFT Analysis

The histograms of the FFT rms observable for the two rings and two planes are shown in Fig. 3.6. It is remarkable that the four peaks show almost identical features. This confirms the fact that the hardware system of the pick-ups is very similar for the two rings and planes. From these plots it is concluded that suitable cuts lie in the range between $1.5 \mu\text{m}$ and $3 \mu\text{m}$. Signals with rms noise above the cut are labeled faulty.

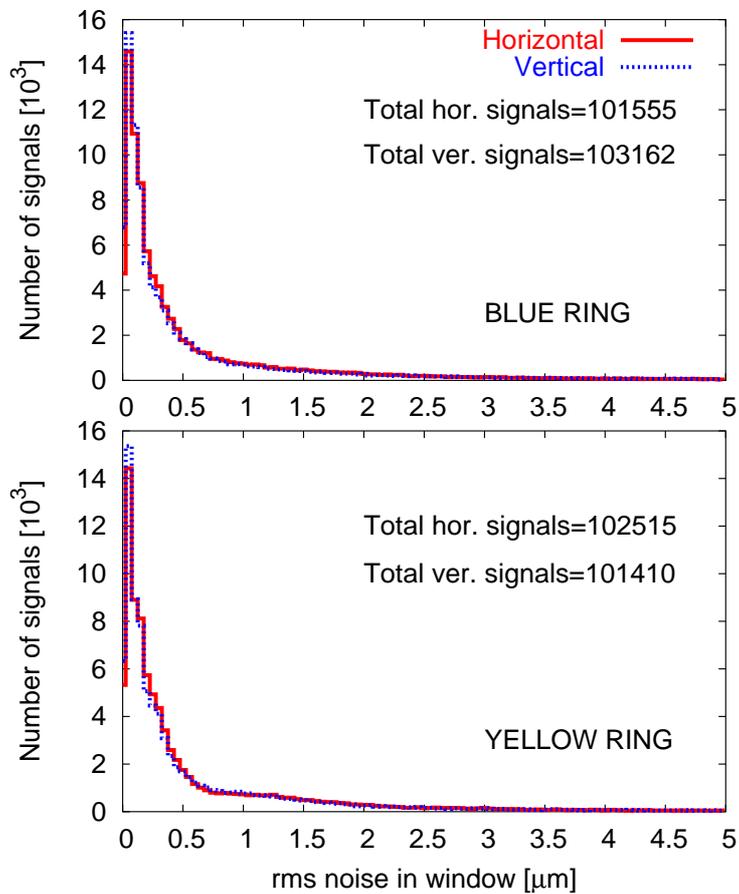


Figure 3.6: Histograms of the rms observable for the two rings and two planes. The total number of signals used in the histograms are shown.

To study the performance of particular BPMs, we record the number of faulty signals provided by each pick-up. Fig. 3.7 shows these occurrences plotted versus longitudinal location for two different cuts ($1.5 \mu\text{m}$ and 2.7

μm). Pick-ups that perform worse than the others are clearly identified from either cut.

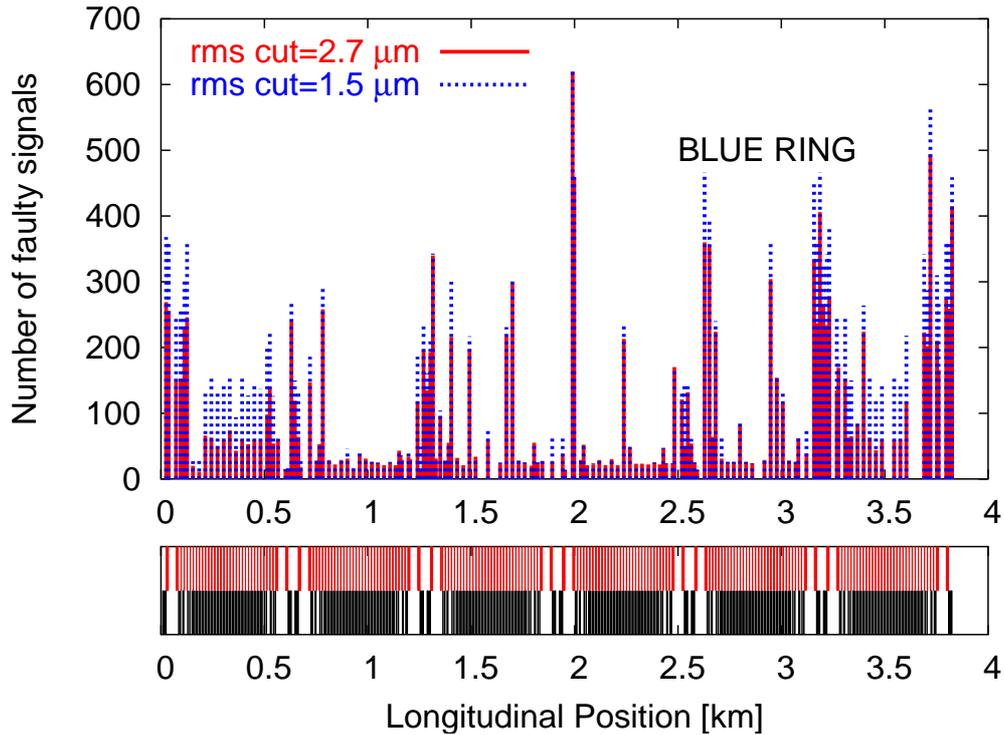


Figure 3.7: Comparison between two different rms cuts showing qualitatively similar results.

To obtain more information on the performance of particular BPMs, the average of the rms noise observables coming from each pick-up is computed and plotted in Fig. 3.8 versus the longitudinal location of the pick-up. The picture shows few spikes that correspond to those systematically bad pick-ups. These spikes happen to be dense in the interaction regions. Fig. 3.8 also shows that the BPMs within certain sextants of the machine have consistently larger rms noise than in the rest of the ring. This will be discussed in section 3.5.

3.4.4 SVD Analysis

The SVD of BPM data with oscillations above the peak-to-peak threshold can be processed to determine spatial modes with localized peaks as shown in

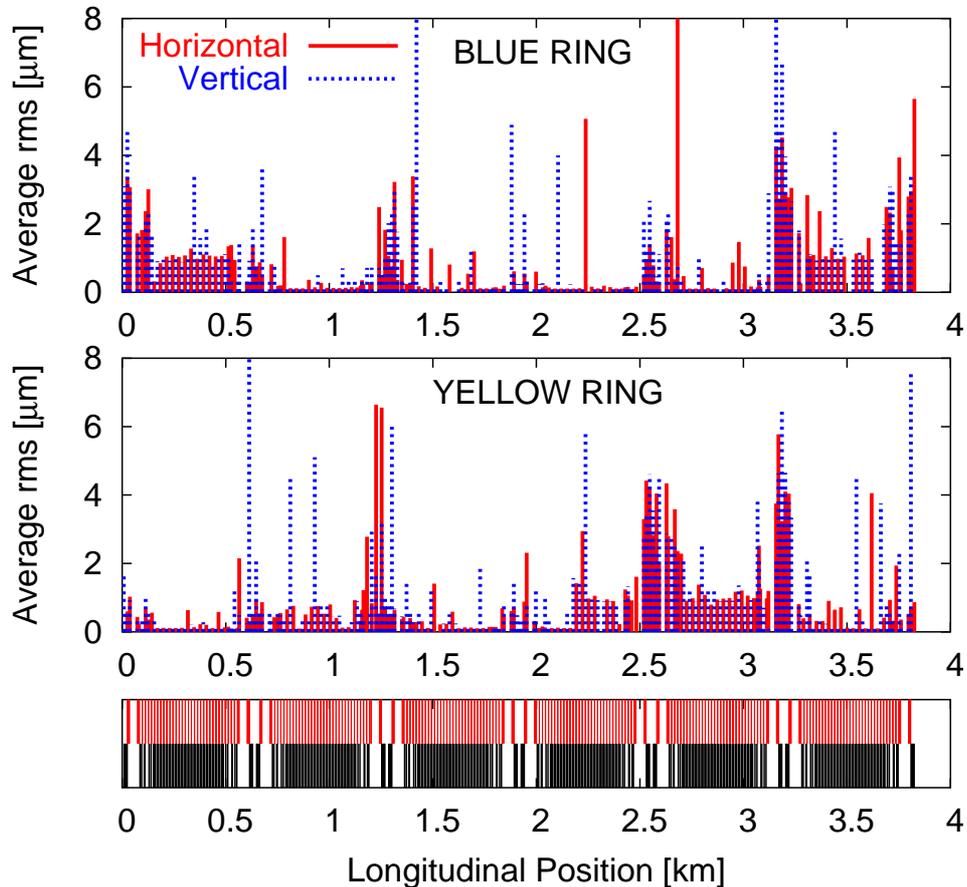


Figure 3.8: Average rms observable versus longitudinal position of the BPM.

Fig. 3.3. The largest peak values of all modes were determined for a large set of data and plotted in the histogram shown in Fig. 3.9. This figure shows a Gaussian-like distribution representing coherent signal in good BPMs and a sharp peak close to one indicating obviously faulty BPMs. In the absence of any correlations between noisy BPMs, the second peak unambiguously identifies the complete set of faulty BPMs. However, for machine data we observe multiple peaks in each mode due to correlations between noisy BPMs. There are many faulty BPMs between 0.6 and 1. Fig. 3.9 does not provide a clear way to determine a good threshold value. A large threshold might be too pessimistic, leaving unidentified faulty BPMs, while a smaller threshold might risk losing good data.

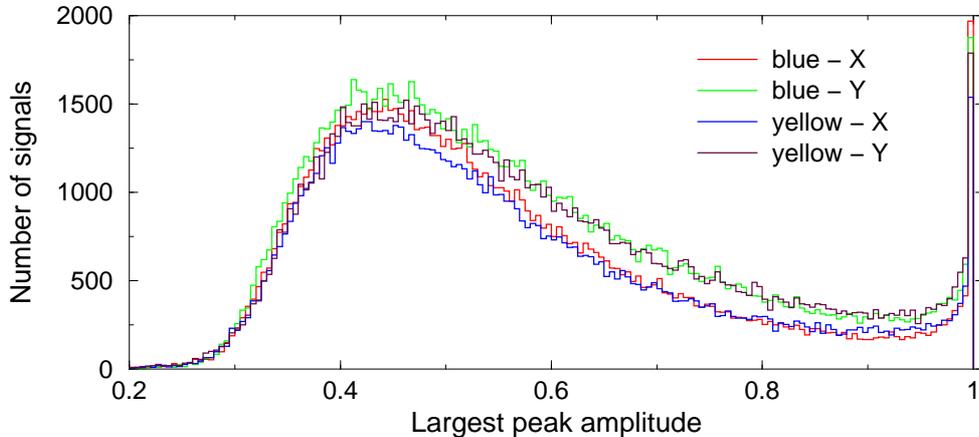


Figure 3.9: Histogram of largest peak amplitudes of all spatial modes for both planes in both rings. The Gaussian like distributions represent coherent signal from good BPMs and peaks at 1 represent faulty BPMs.

Since the singular vectors are normalized, we explore an alternate approach taking advantage of the unit norm of spatial vectors. In the presence of multiple peaks, the norm of the n largest peaks for noisy BPMs,

$$\|V\|_{1 \rightarrow n} = \sqrt{\sum_{j=1}^n v_{ij}^2}, \quad (3.1)$$

is approximately equal to 1, where $2 \geq n \geq 5$. However, a mode consisting of coherent signals usually have a norm < 0.85 for the same n . Thus, the range of the choice of threshold values is considerably smaller than in the largest peak approach. The norms of each mode using the n largest peaks were calculated for all data sets and plotted in the histograms shown in Fig. 3.10. As we increase n , a clear minimum becomes apparent, helping to determine the threshold values more accurately. Thus, this approach is more reliable than using the largest peak approach, where the optimum threshold value is ambiguous. One can determine a suitable threshold value from inspection. For the horizontal plane in the yellow ring, we find the threshold value lies between 0.85 and 1. It is important to note that signals from all four peaks contribute to fail a chosen cut, but we only tag the BPM with the largest amplitude in the norm as faulty to exclude degeneracies. The other correlated BPMs will appear dominant in subsequent modes and will be identified. A similar analysis for both planes and both rings was performed.

These threshold values and the peak-to-peak cut presented earlier are ap-

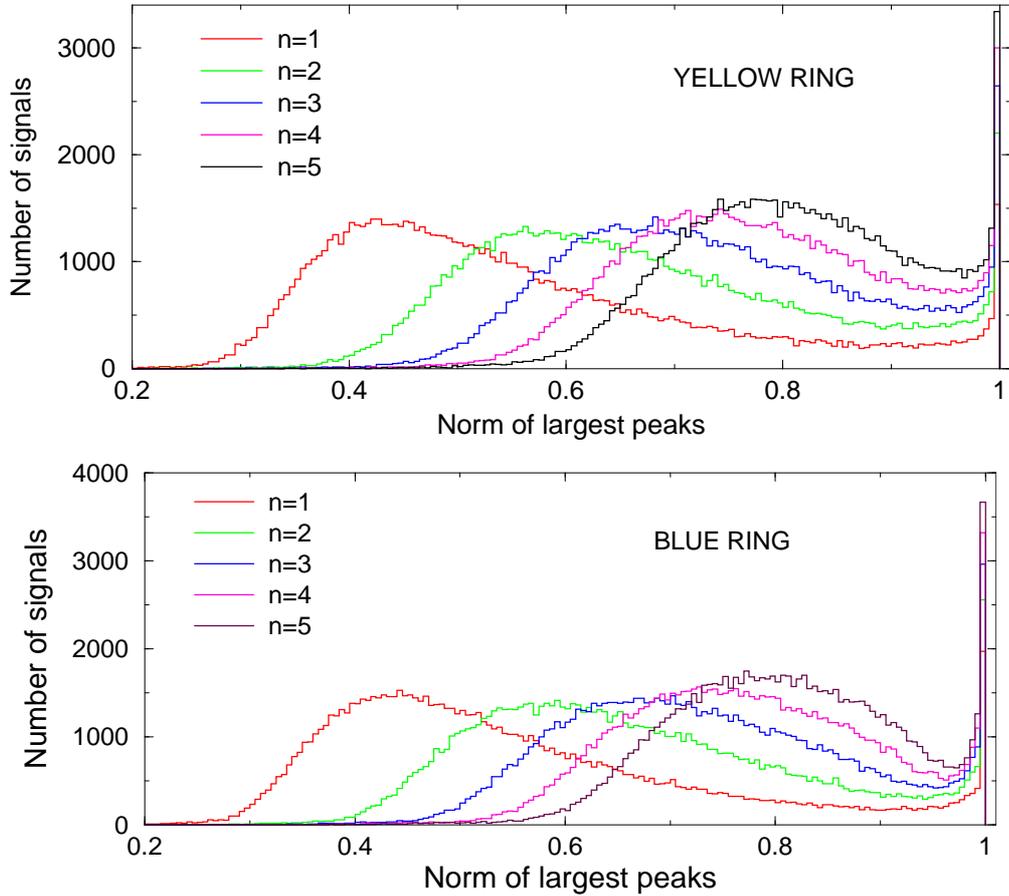


Figure 3.10: Histograms of the norm of n largest peak values in each mode of the spatial modes in the horizontal plane. Gaussian like distributions represent coherent signal from good BPMs and peaks at 1 represent faulty BPMs.

plied to data to construct a final histogram showing the number of occurrences of BPMs failing these cuts. Fig. 3.11 shows such a histogram with two different SVD thresholds (0.85 and 0.95) for the blue ring horizontal data. This comparison allows us to understand the sensitivity of the SVD cut and hence determine an optimum range. It is clear from Fig. 3.11 that the regions corresponding to more noisy BPMs (mostly interactions regions) are significantly enhanced with a lower threshold compared to that of the arc regions. However, certain arc regions show different behavior from the others. This disagreement will be discussed in section 3.5.

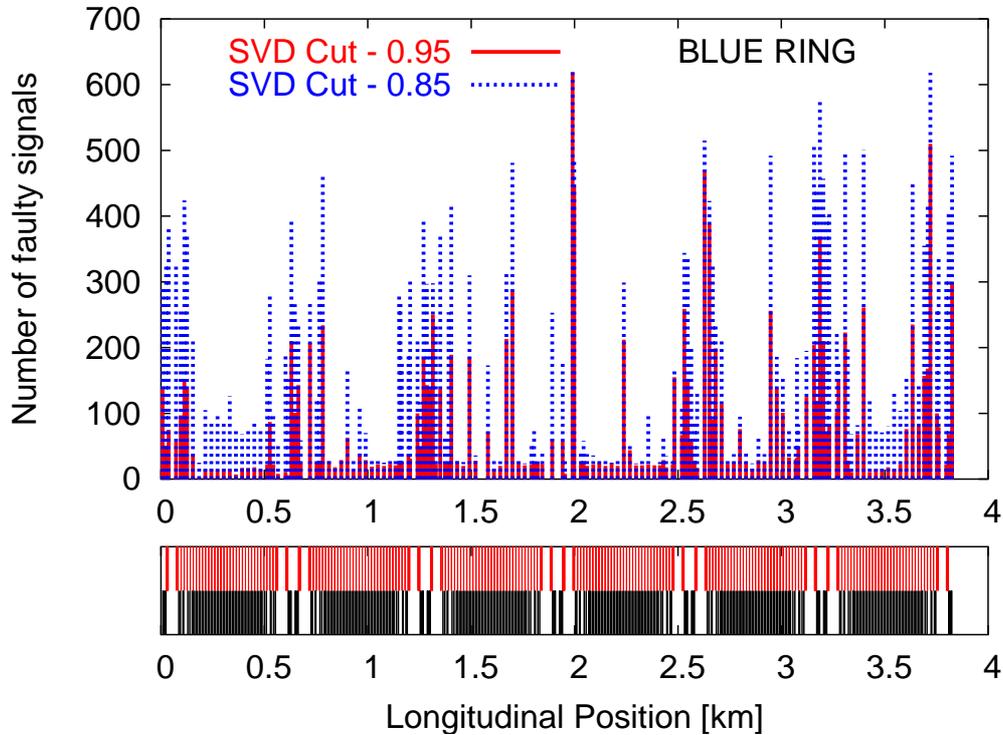


Figure 3.11: Comparison between two SVD thresholds (0.85 and 0.95) for Blue ring-horizontal plane.

3.4.5 SVD & FFT Comparison

To compare the SVD and FFT techniques, the number of identified faulty signals provided by each BPM is plotted versus its location, in Fig. 3.12. This plot contains the horizontal BPMs of the Yellow ring. The reader can also compare the results for horizontal BPMs of the Blue ring from Figs. 3.7 and 3.11. For a set of optimized cuts, the agreement is excellent between the two different techniques. This confirms the feasibility of identifying non-physical signals provided by beam position monitors.

3.5 Observation of System Improvements

As seen in Figs. 3.7, 3.8 and 3.11, BPMs in the arc regions between 0-0.5 km and 3.2-3.8 km have a strikingly larger background than the rest of the arcs. It was found that BPMs in these arcs were exhibiting sporadic noise (“hairs”) of ten to thousands of μm on BPM position data. This effect was caused due

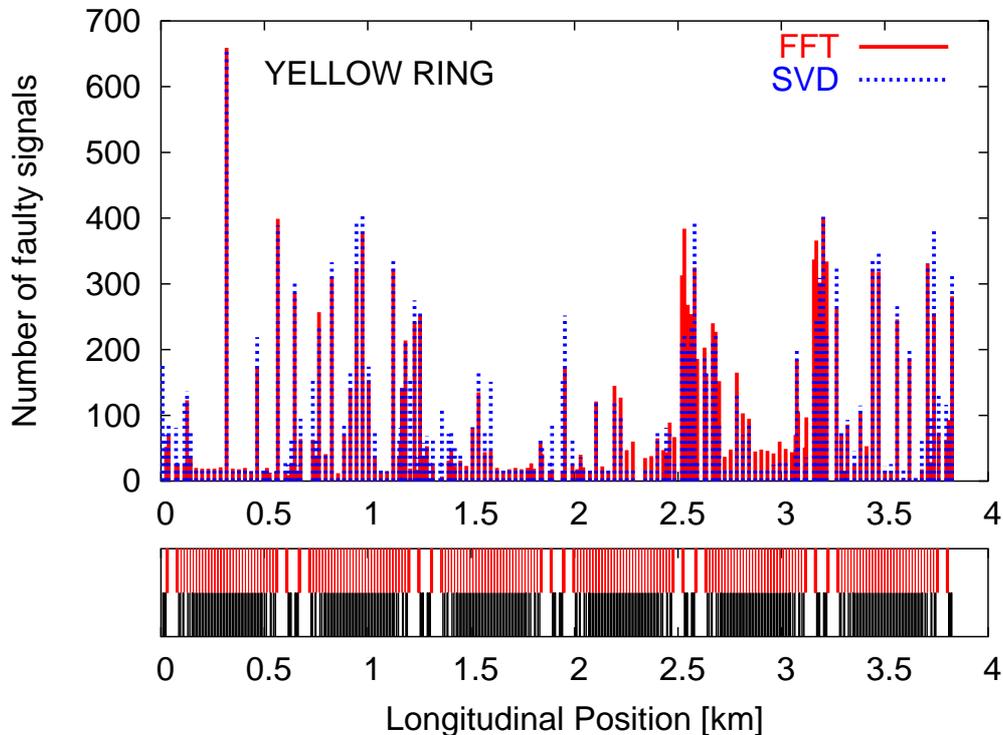


Figure 3.12: Comparison of the FFT and SVD techniques for 678 files from the yellow ring showing a good agreement. The number of times a BPM provides a faulty signal is plotted versus its location.

to unforeseen and untested conditions in low-level DSP code leading to noise jumps on some raw sampled data [38].

Low level software upgrades were implemented as a part of BPM system improvement during the commissioning period of Run 2004. A change of DSP timing parameters was particularly important in resolving the phenomena of “hairs”. Kicked data during Run 2004 were regularly recorded for analysis of BPM performance. A large set of data similar to the one from 2003 was analyzed using the above numerical techniques to understand the behavior of the BPMs after system improvements. Fig. 3.13 shows the percentage of occurrences of system failure per BPM versus longitudinal location for Run 2004. These percentages of system failure are in general larger than those observed in the previous year as seen in Fig. 3.4. This is partly due to the fact that the system improvements lead to a better recognition of system failures. The abnormal abundance of system failures at the location of 3 km is presently under investigation. A collective failure of several BPMs is

most likely due to a failure of the front-end computer system responsible for managing the data flow from those BPMs and not due to individual BPM hardware failure. Fig. 3.14 shows the average of the rms noise observables coming from each BPM plotted versus the longitudinal location of the BPM for Run 2004. This figure is to be compared to Fig. 3.8 from the previous year. We observe a clear improvement in BPM resolutions. The larger background at particular sextants has disappeared yielding a consistent BPM performance in all the arcs. This confirms that the problem of “hairs” has been resolved. Nevertheless BPMs located within the interaction regions still show a poorer performance than the rest. Run 2004. The background in PP run is smaller than in Au run, which seems to indicate failure rate of good BPMs are lower.

Peak-to-peak signals in PP run show a smaller occurrence close to zero amplitude than the Au-Au run indicating fewer BPMs with flat signals possibly contributing to the lower background. However, we find more faulty BPMs in IR8 (0.5 km) in PP run as seen in Fig. 3.15 which might point to possible damage of BPMs due to radiation or other operation conditions. Also, BPMs located within the interaction regions in all runs show poorer performance than the rest of the ring despite similar BPM electronics around the ring.

3.6 Conclusion

Two numerical algorithms, the Fourier transform and the Singular Value Decomposition, have been used to identify faulty BPM signals. Appropriate observables characterizing the noise content of a BPM signal have been defined for both techniques. The observables’ thresholds above which a signal is identified as faulty are obtained from statistics over a large set of RHIC BPM data. These cuts are used to assess the global performance of every BPM, thus identifying those BPMs that systematically provide faulty data. A comparison between the results from both independent techniques has been presented showing an excellent agreement. Analysis of run 2004 data shows clear improvements in the BPM system and successful elimination of “hairs”.

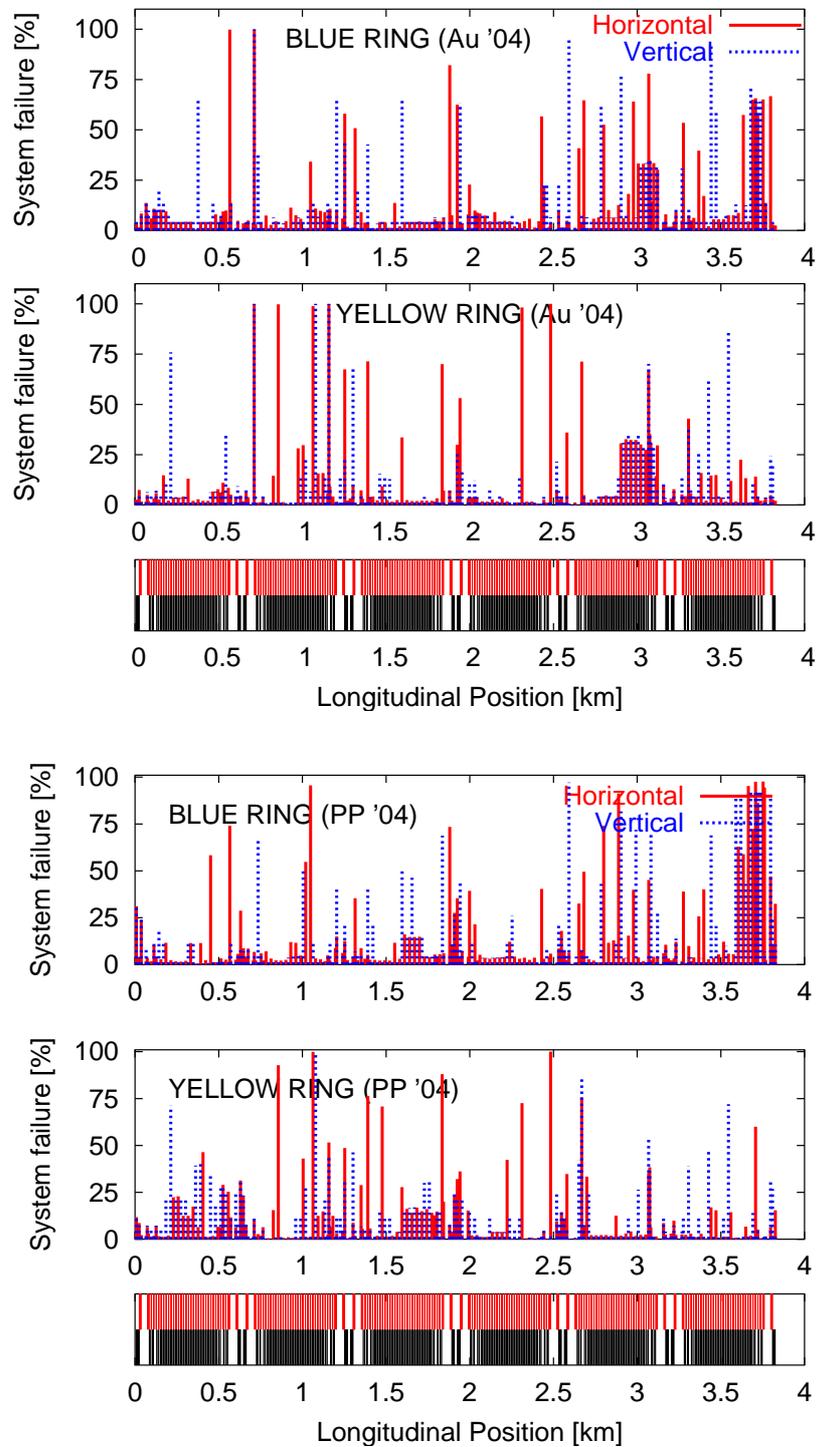


Figure 3.13: Percentage of occurrences of system failure per BPM versus longitudinal location. Run 2004 (Au-Au), and Run 2004 (PP). A representation of the lattice (dipoles in black and quadrupoles in red) in the bottom.

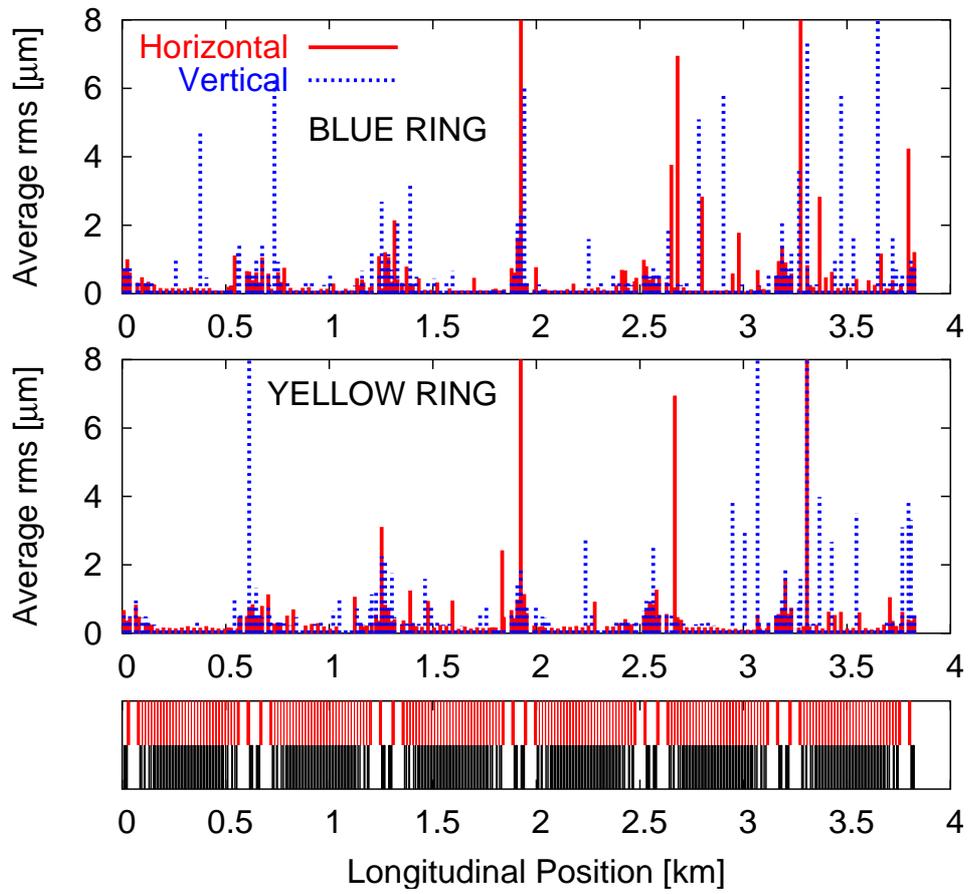


Figure 3.14: Average rms observable versus longitudinal position of the BPM for Run 2004.

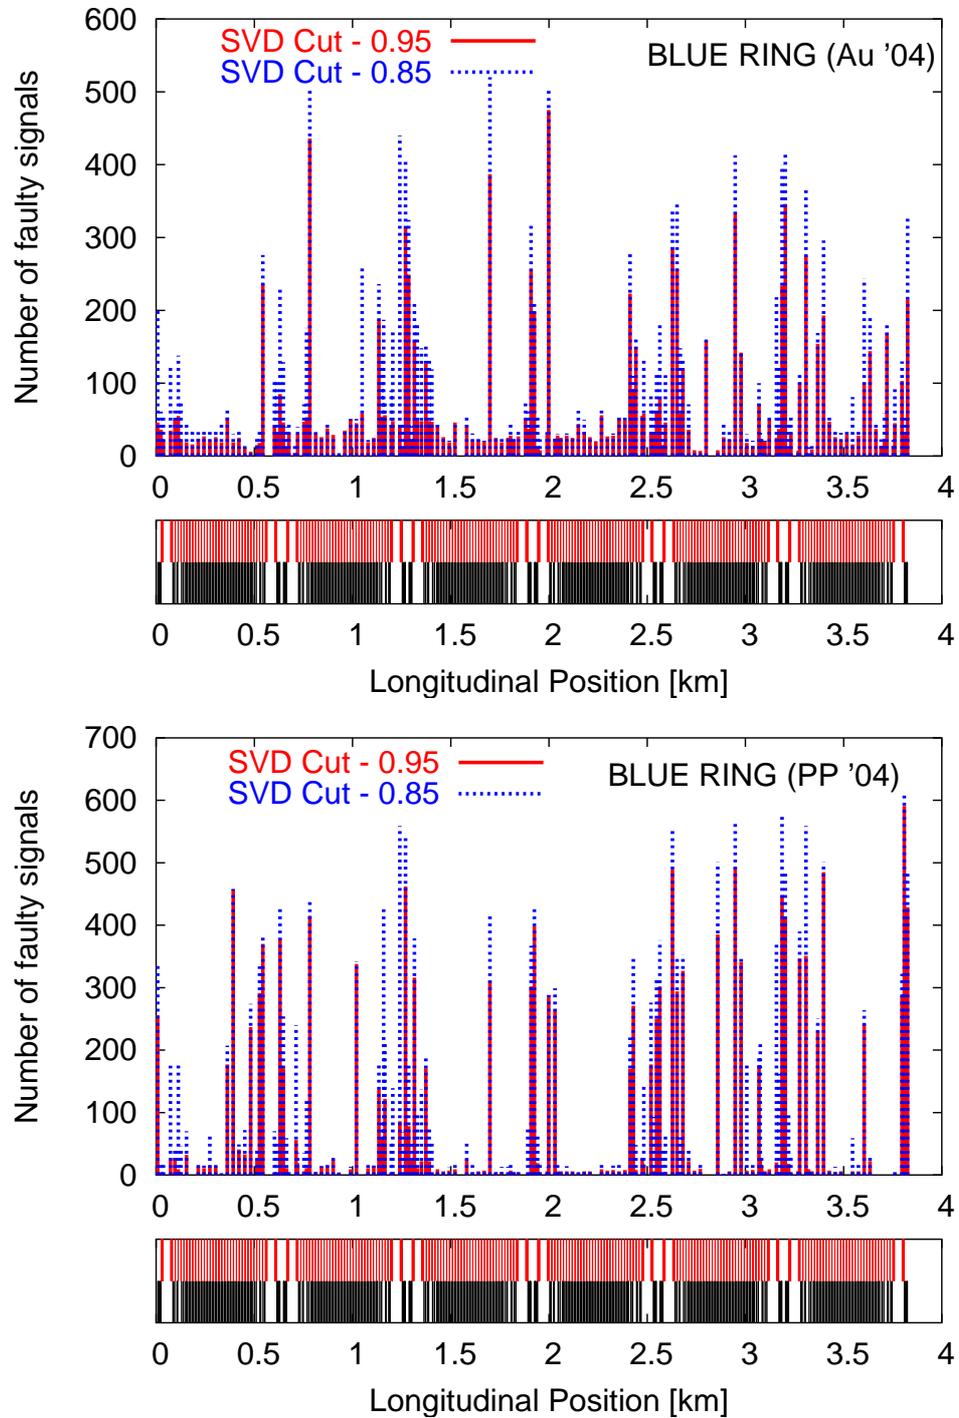


Figure 3.15: Histogram showing number of occurrences of faulty BPMs around the ring. Top to bottom: Run 2003, Run 2004 (Au-Au), and Run 2004 (PP). A representation of the lattice (dipoles in black and quadrupoles in red) in the bottom.